
A Stochastic Interpretation of Quantum Information

From Probabilistic Computing to Quantum Foundations

A. HARTMAN

October 25, 2026

Copyright © 2026 A. Hartman.

All rights reserved.

ISBN 978-1-291-55573-8

Contents

Introduction: The Magician's Trick	1
I Three Kinds of Computing	3
1 Reversible Computing	4
1.1 Bits and state vectors	4
1.2 Combining systems: the Kronecker product	5
1.3 Taking systems apart: factorization	6
1.4 Time symmetry and the conservation of information	7
1.5 Logic gates and truth tables	7
1.6 Operators as matrices	9
1.7 Evolving the state	10
1.8 Circuit notation and a first program	11
2 Probabilistic Computing	13
2.1 Probability vectors	13
2.2 Marginalization, and what it loses	14
2.3 Stochastic matrices	15
2.4 Factorizing stochastic matrices	16
2.5 Measuring uncertainty: information entropy	18
2.6 Programs that create correlation	19
2.7 Simulating a Markov chain: stochastic jumps	20
2.8 Local gates without giant matrices	21
2.9 Bayes' theorem: updating on what you learn	22
2.10 Sufficient statistics	24
3 Quantum Computing	28
3.1 A second degree of freedom	28
3.2 The quantum state	29
3.3 Unitary operators	31
3.4 Gates and observables	32
3.5 A quantum program: the entangled pair	36
3.6 The unistochastic bridge	36
3.7 Measurement: Bayes' theorem for amplitudes	38
3.8 Projectors	39

4	Density Matrices	42
4.1	Why the quantum state is not enough	42
4.2	The density matrix	43
4.3	Mixtures, and what a density matrix really describes	44
4.4	Purity	45
4.5	The partial trace	45
4.6	Kraus operators, quantum form	46
4.7	Von Neumann entropy	47
4.8	Quantum mutual information	48
4.9	When classical mutual information misleads	48
4.10	Decoherence	49
II	Quantum Information as Statistics	53
5	Non-Markovian Stochastic Dynamics	54
5.1	The same gate, twice	54
5.2	Why this is a genuine puzzle	55
5.3	The hidden-memory escape	57
5.4	The trouble with hidden memories	58
5.5	History instead of memory	59
5.6	The polar decomposition	60
5.7	The recursive expansion	62
5.8	The phase as a sufficient statistic	64
6	Unistochastic Dynamics	67
6.1	Complex numbers as bookkeeping	67
6.2	The multi-step shortcut	68
6.3	Base cases that appear mid-circuit	68
6.4	A worked example: the interrupted Hadamard pair	70
6.5	What unistochastic dynamics still lacks	71
6.6	The unique single-step transition matrix	73
6.7	Divisibility: where we part company	75
6.8	Divisibility and realism	77
6.9	Divisibility and minimality	81
6.10	Divisibility and counterfactuals	85
7	Quantum Algorithms	94
7.1	A picture of the phase: the hypergraph	94
7.2	Correlation hidden in the phase, revisited	96
7.3	Superdense coding: two bits from one	97
7.4	Phase kickback	99
7.5	When classical simulation is hard	100
7.6	Reading the answer out	103
7.7	Deutsch's algorithm	104
7.8	Why updating appears mandatory	107

7.9	Quantum cryptography	108
7.10	Quantum teleportation	112
8	Quantum Interference	122
8.1	The double slit	122
8.2	The Mach–Zehnder interferometer	123
8.3	The wave picture drawn out	125
8.4	Why the wave cannot be in the world	126
8.5	Interaction-free measurement	131
8.6	The hidden assumption	131
8.7	Where the base case comes from	133
8.8	Holism	135
8.9	When the statistics fix the law	142
8.10	Simulating the double slit	143
8.11	Wheeler’s delayed-choice experiment	149
8.12	Albert’s moving pointer	150
III	Restoring Realism	156
9	Realism	157
9.1	Theories that point	157
9.2	The need for systems	164
9.3	Wigner’s friend	169
9.4	Variant collapse is solipsism	170
9.5	Invariant collapse is a different theory	171
9.6	What remains	174
10	Quantum Contextuality	177
10.1	Four runs and a bookkeeping certainty	177
10.2	The experiment as circuits	179
10.3	Contextuality	180
10.4	Observables and emergeables	181
10.5	Measurement as a function of what is really there	182
10.6	The underdetermination problem	184
11	Quantum Non-Locality	189
11.1	A two-bit distillation	189
11.2	The inferences, and the contradiction	191
11.3	Why the inferences cannot be local	191
11.4	The time-ordering of events	193
11.5	The EPR paradox	195
11.6	Bell’s theorem	200
11.7	The collision with relativity	203
11.8	Apparent relativity	203
11.9	Non-locality without signals	205

11.10	The other horn: retrocausality	208
11.11	Both conclusions in one branch	214
11.12	Relational quantum mechanics	217
12	The PBR Theorem	226
12.1	Four preparations	226
12.2	The black box	227
12.3	Why no preparation-blind box can do this	228
12.4	The second premise	228
12.5	The law consults the history	229
	Conclusion: Grappling with Underdetermination	233
	References	238

Introduction: The Magician's Trick

Consider that you possess a probabilistic computer with a single bit of memory. Some programs on this computer stochastically flip that bit so that its final value is uniformly distributed: a 50% chance of being 0 and a 50% chance of being 1. Other programs place it into a *degenerate* distribution: they produce a 0 every single time you run them, or a 1 every single time. Call these the two categories of programs: the uniform ones and the degenerate ones.

A magician tells you to run a program from whichever category you like, copy the computer's final memory state, that single bit, onto a thumb drive, and hand it to him. You do. He pauses a moment, then tells you exactly which category of program you ran. You repeat this many times, choosing categories at random, and he is right every single time.

If you think about it, this should be impossible. A probability is never a property of an individual system. It is a property of an *ensemble*: the whole collection of runs you get by repeating the same experiment over and over under identical conditions. Any single run is one member of that collection and nothing more. "Uniform" and "degenerate" describe the collection, not the run: they tell you the proportions in which 0s and 1s come up across all the runs. About any one run they say nothing at all, because a single run produces a single bit, and one bit has no proportions in it.

A single bit value of 0 therefore tells you nothing about the ensemble it came from. It could have come from a program that always outputs 0, or from a program that outputs 0 half the time; the bit is just a 0 either way. One sample simply does not carry enough information to distinguish the two categories. Yet the magician distinguishes them, with certainty, every time.

Now consider a real system that behaves exactly like the magician. Take a quantum computer with a single bit of memory, initialized to 0, and apply to it the logic gate known as the Hadamard gate. Measure the bit, and the outcome is uniformly distributed: 50% for 0 and 50% for 1. Initialize the bit to 1 instead, apply the same gate, and the outcome is again uniform. So far, the Hadamard gate looks like a fair coin flip: whatever goes in, a uniform distribution comes out.

But now apply the Hadamard gate *twice in a row* to a bit initialized to 0. If the gate were a coin flip, the second application would receive a uniformly distributed bit and output a uniformly distributed bit, and we would see 0 half the time and 1 half the time. That is not what happens. The outcome is 0, deterministically, 100% of the time. The second Hadamard gate somehow undoes the first.

The second gate receives a single bit. That bit has a value of 0 or a value of 1, and, as we just said, a single bit value carries no information about the distribution it was drawn from. Yet the gate behaves one way when the incoming bit belongs to a degenerate ensemble and a

completely different way when it belongs to a uniform ensemble. The Hadamard gate is the magician.

This little puzzle is the heart of what makes quantum mechanics strange. The more famous paradoxes, wave-particle duality and Schrödinger's cat, are versions of it rather than separate difficulties. We set out to explain the trick.

The aim here is to demystify. Quantum mechanics is usually taught as a theory one can use but not understand, and that strangeness is blamed on nature herself: particles that exist nowhere until somebody looks, a world that splits at every measurement, facts that hold for one observer and not for another; a reality that refuses to be described in ordinary language. A century of that has left the impression that the theory is not merely surprising but beyond comprehension. It is not.

What we mean to restore is the ordinary picture of the world that physics had before it. The *ontology*, the furniture of the world, consists of objects, which occupy a definite place at a definite time and maintain their properties whether or not anyone is watching. The *nomology*, the body of law those objects answer to, consists of well-defined mathematical rules which drive them forward with the flow of time. There is one reality, with one past, and that past is precisely the one that actually happened.

The difficulty has never lain in nature herself, as if the world forbade any coherent account of itself. It has been a choice about how to describe it, made early, defended vigorously, and then handed on as something the experiments had settled. To quote the physicist John Bell, who proved the best known of the theorems later said to forbid all this: the “subjectivity of the orthodox version, the necessary reference to the ‘observer,’ could be eliminated,” and “vagueness, subjectivity, and indeterminism, are not forced on us by experimental facts, but by deliberate theoretical choice.”¹

A reader who has heard otherwise will object that this road is closed, and will have the roadblocks ready: the double slit, the Kochen–Specker theorem, Bell's own inequality, Frauchiger and Renner's paradox, the Pusey–Barrett–Rudolph theorem. Every one of them is taken up here in its turn, rebuilt as a circuit, and computed. Not one of them rules out what follows. To quote Bell again, “what is proved by impossibility proofs is lack of imagination.”²

This is not a traditional quantum mechanics text. The mathematics is kept as simple as the subject allows: everything is finite and discrete, done with bits, vectors and matrices. What is left when the machinery is pared back this far is the abstract structure of quantum information, which is the logical structure of the theory itself, and that structure is where every question here actually lives. Differential equations, infinite-dimensional spaces and particles in boxes contribute nothing to those questions; they raise the barrier to entry and settle none of the arguments. Our laboratory is accordingly the quantum *computer*: a finite system, evolving in discrete steps, whose every experiment can be written down as a small circuit and computed by hand, from the Mach–Zehnder interferometer to Wigner's friend.

The magician's trick has a banal explanation. Nature's version of it does too. It is just that the explanation lies where a long habit of thought has ruled out looking: not in the present state of the bit, and not in an invisible wave, but in *the past*, and in *the arrangement* of the world around it.

¹J. S. Bell, “On the impossible pilot wave,” *Foundations of Physics* 12 (1982), 989–999, at 989–990; reprinted in J. S. Bell, *Speakable and Unsayable in Quantum Mechanics* (Cambridge: Cambridge University Press, 1987). The second passage is Bell's, put in the original as a rhetorical question.

²Bell, “On the impossible pilot wave,” 997.

Part I

Three Kinds of Computing

Chapter 1

Reversible Computing

Quantum mechanics has a reputation for being unlike anything else. We are going to earn the right to disagree with that reputation, and the way to earn it is to start as far from quantum mechanics as possible: with ordinary, deterministic bits, doing ordinary, deterministic things. There is no physics and no probability here. What it contains is the complete mathematical skeleton of everything that follows: state vectors, the Kronecker product, operators as matrices, and circuit notation, each introduced where it is easiest to understand, in a setting where nothing is uncertain and everything can be checked by counting on fingers.

The setting is a very simple kind of computer: one whose memory evolves without ever destroying information. Such a computer is called *reversible*, and by the end of the chapter the reason for that restriction will be clear: physical law is time-symmetric, and a time-symmetric law destroys no information.

1.1 Bits and state vectors

The bit, the most fundamental unit of information in computing, exists in one of two discrete states: 0 or 1. Binary digits form the bedrock of digital systems, where collections of bits are represented through Boolean variables. A system of N bits can be modeled using N Boolean variables, each encoding the state of an individual bit. This binary abstraction underpins all computational processes, from the simplest arithmetic to the most complex algorithms.

That is the familiar way to describe a computer's memory, and for most of computer science it is the right one. We are, however, going to adopt a different representation, one that looks wasteful at first and will pay for itself many times over.

Rather than writing down the value of each bit, we will write down a vector with one entry for every possible *configuration* of the whole memory, and mark the configuration the memory is actually in. For a single bit there are two possible configurations, 0 and 1, so the vector has two entries. The entry corresponding to the current state is 1, and the other entry is 0:

$$\mathbf{e}_0 = \begin{pmatrix} 1 \\ 0 \end{pmatrix} \quad \mathbf{e}_1 = \begin{pmatrix} 0 \\ 1 \end{pmatrix}.$$

We will call these *bit vectors*, and write them in boldface. Think of the vector as a column of labeled pigeonholes, one per configuration, with a single token sitting in the pigeonhole the system occupies.

A complete set of these, one to each configuration, is called a *basis*. The word means only this: a collection of vectors from which every other vector in the space can be assembled as a weighted sum, in exactly one way. The two bit vectors above are a basis for a single bit, since any pair of numbers (a, b) is $a\mathbf{e}_0 + b\mathbf{e}_1$ and is that for no other choice of a and b . Writing a vector down at all means choosing a basis first, and a different choice gives different entries for the same object. Chapter 3 will find more than one basis available, which is where the choice starts to matter.

The encoding generalizes to systems of multiple bits. A system of N bits has 2^N possible configurations, and each configuration is identified by a distinct vector of length 2^N with a single 1 in the position indexing that configuration and 0 everywhere else. The positions are ordered by reading the configuration as a binary number. For example, a 3-bit system spans $2^3 = 8$ configurations, 000 through 111, represented by 8-element vectors:

$$\mathbf{e}_{000} = \begin{pmatrix} 1 \\ 0 \\ 0 \\ 0 \\ 0 \\ 0 \\ 0 \\ 0 \end{pmatrix} \quad \mathbf{e}_{001} = \begin{pmatrix} 0 \\ 1 \\ 0 \\ 0 \\ 0 \\ 0 \\ 0 \\ 0 \end{pmatrix} \quad \mathbf{e}_{010} = \begin{pmatrix} 0 \\ 0 \\ 1 \\ 0 \\ 0 \\ 0 \\ 0 \\ 0 \end{pmatrix} \quad \mathbf{e}_{011} = \begin{pmatrix} 0 \\ 0 \\ 0 \\ 1 \\ 0 \\ 0 \\ 0 \\ 0 \end{pmatrix} \quad \cdots$$

For a deterministic system this representation is plainly inefficient: we are using 2^N entries to store what N Boolean variables store, and only one entry is ever active. Hold that objection. Our entire plan is to let the other entries come alive, first as probabilities in Chapter 2, then as something stranger in Chapter 3, and the value of the representation is that *nothing else about the formalism will need to change* when they do. Whatever we build for definite bits here, we keep.

1.2 Combining systems: the Kronecker product

Suppose one computer's memory holds 2 bits, described by a state vector of size 4, and a second computer holds 1 bit, described by a state vector of size 2. We would like a single vector describing the memory of both computers at once: a 3-bit system, so a vector of size 8. The operation that combines the two vectors is called the *Kronecker product*, written $\otimes$.

The Kronecker product of two matrices is computed by taking each element of the left-hand matrix and multiplying it by the *entire* right-hand matrix, laying out the results in a grid:

$$A \otimes B = \begin{pmatrix} A_{11}B & A_{12}B & \cdots & A_{1n}B \\ A_{21}B & A_{22}B & \cdots & A_{2n}B \\ \vdots & \vdots & \ddots & \vdots \\ A_{m1}B & A_{m2}B & \cdots & A_{mn}B \end{pmatrix}.$$

The definition applies to column vectors as well, because a column vector of length N is just an $N \times 1$ matrix.

Let us compute one example in full. Take the first computer to be in the configuration 10 and the second in the configuration 1:

$$\mathbf{e}_{10} \otimes \mathbf{e}_1 = \begin{pmatrix} 0 \\ 0 \\ 1 \\ 0 \end{pmatrix} \otimes \begin{pmatrix} 0 \\ 1 \end{pmatrix} = \begin{pmatrix} 0 \cdot \begin{pmatrix} 0 \\ 1 \end{pmatrix} \\ 0 \cdot \begin{pmatrix} 0 \\ 1 \end{pmatrix} \\ 1 \cdot \begin{pmatrix} 0 \\ 1 \end{pmatrix} \\ 0 \cdot \begin{pmatrix} 0 \\ 1 \end{pmatrix} \end{pmatrix} = \begin{pmatrix} 0 \\ 0 \\ 0 \\ 0 \\ 0 \\ 1 \\ 0 \\ 0 \end{pmatrix} = \mathbf{e}_{101}.$$

The result is exactly the bit vector for the combined configuration 101: the bits of the left factor become the most significant bits of the result, and the bits of the right factor the least significant. In general,

$$\mathbf{e}_{\text{MSB,LSB}} = \mathbf{e}_{\text{MSB}} \otimes \mathbf{e}_{\text{LSB}},$$

where MSB and LSB stand for the most and least significant bits. Combining independent systems into one big system is, from now on, a single multiplication.

1.3 Taking systems apart: factorization

It is also possible to run this in reverse: to break a larger state vector into smaller ones. This is called *factorization*, since the smaller vectors, recombined with the Kronecker product, must reproduce the original.

To factorize, we first need the *identity matrix*. The identity matrix has a 1 in each of its diagonal elements (the elements whose row and column index are the same) and a 0 in each of its off-diagonal elements. Multiplying by it changes nothing; it is the matrix that means “leave this alone.”:

$$\mathbf{I} = \begin{pmatrix} 1 & 0 & 0 & \cdots & 0 \\ 0 & 1 & 0 & \cdots & 0 \\ 0 & 0 & 1 & \cdots & 0 \\ \vdots & \vdots & \vdots & \ddots & \vdots \\ 0 & 0 & 0 & \cdots & 1 \end{pmatrix}.$$

Identity matrices are often given a subscript to denote their size, so $\mathbf{I}_2$ is the 2×2 identity, $\mathbf{I}_4$ the 4×4 , and so on.

To remove unwanted bits from a state vector, we construct a *removal matrix* out of Kronecker products: in the position of each bit we wish to *keep*, we place the identity matrix $\mathbf{I}_2$; in the position of each bit we wish to *remove*, we place the row vector $(1 \ 1)$, which sums over that bit’s two possibilities and thereby forgets it:

$$\mathbf{e}_{\text{reduced}} = (\mathcal{R}_1 \otimes \mathcal{R}_2 \otimes \mathcal{R}_3 \otimes \cdots \otimes \mathcal{R}_n) \mathbf{e} \quad \text{where } \mathcal{R}_j \in \{\mathbf{I}_2, (1 \ 1)\}.$$

For a two-bit system, keeping only one bit or the other:

$$\mathbf{e}_{\text{MSB}} = (\mathbf{I}_2 \otimes (1 \ 1)) \mathbf{e}_{\text{MSB,LSB}}, \quad \mathbf{e}_{\text{LSB}} = ((1 \ 1) \otimes \mathbf{I}_2) \mathbf{e}_{\text{MSB,LSB}}.$$

The reader is encouraged to verify one case by hand: applying $I_2 \otimes \begin{pmatrix} 1 & 1 \\ 0 & 1 \end{pmatrix}$, which works out to

$$\begin{pmatrix} 1 & 1 & 0 & 0 \\ 0 & 0 & 1 & 1 \end{pmatrix},$$

to $\mathbf{e}_{10}$ returns $\mathbf{e}_1$: the most significant bit was 1, and the least significant bit has been forgotten.

For definite bit vectors, factorization always succeeds and never loses anything interesting: a configuration of bits is fully described by each bit's value. When we move to statistical distributions in Chapter 2, this will stop being true: taking a joint distribution apart and recombining the pieces can silently destroy the *correlations* between bits, and that failure will turn out to be one of the most important phenomena here. The machinery, however, is already in place.

1.4 Time symmetry and the conservation of information

Physical laws are, in most cases, time-symmetric: given complete information about the present configuration of a system, we can simulate how it will evolve forward in time, and we can equally run the clock backward and recover where it came from. If the natural world is time-symmetric in this sense, then the present configuration of a system must contain sufficient information to evolve it in either direction, which means, inherently, that no information can be lost as the system evolves forward. If information were lost, there would be no way to evolve the system backward unambiguously: multiple pasts would be compatible with the same present.

For a computer's memory, preserving information means the bits may only evolve according to *one-to-one maps*: mappings that send each possible configuration, as an input, to one and only one configuration as an output, with no two inputs sharing an output. A one-to-one map stands in contrast to a many-to-one mapping, which funnels several input configurations into a single output configuration (destroying the information of which input occurred), and to a one-to-many mapping, which sends one input to several possible outputs (and so is not deterministic at all; we will meet those in Chapter 2).

Deterministic computing in which information is never lost, because the bits evolve according to one-to-one mappings at all times, is called *reversible computing*.

1.5 Logic gates and truth tables

The deterministic evolution of bits is governed by *logic gates*, also referred to as *operators*. A logic gate can be specified by a truth table, which lists what output the gate produces for each possible input. Three of the most familiar gates in ordinary (non-reversible) computing are OR, XOR, and AND:

A	B	$A \text{ OR } B$	A	B	$A \text{ XOR } B$	A	B	$A \text{ AND } B$
0	0	0	0	0	0	0	0	0
0	1	1	0	1	1	0	1	0
1	0	1	1	0	1	1	0	0
1	1	1	1	1	0	1	1	1

Every one of these is many-to-one. The XOR gate maps the configurations 00 and 11 both to 0, and maps 01 and 10 both to 1: two bits go in, one bit comes out, and there is no way to tell from the output which input produced it. In reversible computing, the number of bits on the input must equal the number of bits on the output, since otherwise the map cannot possibly be one-to-one, so operators of this form are not permitted as they stand.

It is often possible to *extend* a non-reversible operator into a reversible one without adding any bits at all. Suppose we keep both bits: the result of the operation is stored in the least significant bit (the right-most bit, often denoted LSB), while the most significant bit (the left-most, MSB) is passed through unchanged. The truth tables become:

A	B	$A \text{ OR } B$	A	B	$A \text{ XOR } B$	A	B	$A \text{ AND } B$
0	0	00	0	0	00	0	0	00
0	1	01	0	1	01	0	1	00
1	0	11	1	0	11	1	0	10
1	1	11	1	1	10	1	1	11

This succeeds in making XOR a one-to-one map: the four outputs 00, 01, 11, 10 are all distinct, so nothing is lost. (This reversible XOR is important enough to have its own name, as we will see in a moment.) But the trick fails for OR and AND. For the extended OR, both 10 and 11 map to 11; for the extended AND, both 00 and 01 map to 00. These operators remain many-to-one no matter how the output convention is chosen, because their output column simply contains too little variety.

When extending the operator in this way does not yield a one-to-one mapping, reversibility can still be achieved by introducing an *ancilla bit*: an extra bit that is not part of the calculation itself but is added purely to give the lost information somewhere to live. With one ancilla bit X , the OR and AND operators can be implemented reversibly:

X	A	B	$A \text{ OR } B$	X	A	B	$A \text{ AND } B$
0	0	0	000	0	0	0	000
0	0	1	001	0	0	1	100
0	1	0	011	0	1	0	010
0	1	1	111	0	1	1	011
1	0	0	010	1	0	0	001
1	0	1	100	1	0	1	101
1	1	0	101	1	1	0	110
1	1	1	110	1	1	1	111

In each table, all eight outputs are distinct, so the maps are genuinely one-to-one. Typically, a program initializes its ancilla bits to 0, and so only the first half of the rows ever occur in practice; the second half matters only in that its outputs must be chosen to be unique so that the whole table remains one-to-one. Since every Boolean function can be built from gates like these, reversible computing can compute anything ordinary computing can; the cost is only some extra bits of memory.

Two additional operators are common in reversible computing though uncommon outside of it. The SWAP operator exchanges the values of two bits. The CX operator (“controlled-X,” also called controlled-NOT) flips the value of the least significant bit, called the *target*, if and

only if the most significant bit, called the *control*, has a value of 1:

A	B	A SW B	A	B	A CX B
0	0	00	0	0	00
0	1	10	0	1	01
1	0	01	1	0	11
1	1	11	1	1	10

Both are one-to-one. Note that CX is exactly the reversible XOR we constructed above: it computes the XOR of its two inputs into the target while carrying the control through unchanged. The reader may verify a useful fact about each: applying SWAP twice returns every configuration to where it started, and so does applying CX twice. Each of these gates is its own reverse.

The simplest reversible gate of all acts on a single bit: the NOT operator, which flips the bit's value. Its truth table sends 0 to 1 and 1 to 0, and it, too, is its own reverse.

1.6 Operators as matrices

Any operator whose input and output have the same number of bits can be translated directly into a matrix. The recipe: each *column* of the matrix is associated with one input configuration, in binary order: for a two-bit system, the first column with input 00, the second with 01, the third with 10, the fourth with 11. Each *row* is associated with an output configuration in the same order. To record that the operator maps input 01 to output 11, place a 1 in the second column at the fourth row. Every column gets exactly one 1 (each input produces exactly one output), and all other entries are 0.

Carrying this out for SWAP, CX, and the extended (two-bit, non-reversible) AND gives:

$$\text{SW} = \begin{pmatrix} 1 & 0 & 0 & 0 \\ 0 & 0 & 1 & 0 \\ 0 & 1 & 0 & 0 \\ 0 & 0 & 0 & 1 \end{pmatrix} \quad \text{CX} = \begin{pmatrix} 1 & 0 & 0 & 0 \\ 0 & 1 & 0 & 0 \\ 0 & 0 & 0 & 1 \\ 0 & 0 & 1 & 0 \end{pmatrix} \quad \text{AND} = \begin{pmatrix} 1 & 1 & 0 & 0 \\ 0 & 0 & 0 & 0 \\ 0 & 0 & 1 & 0 \\ 0 & 0 & 0 & 1 \end{pmatrix}$$

and for the one-bit NOT:

$$\text{NOT} = \begin{pmatrix} 0 & 1 \\ 1 & 0 \end{pmatrix}.$$

If the operator is one-to-one, then not only does each column contain a single 1, but each *row* does as well: no output is reachable from two different inputs. This is true for SW and CX above, but visibly false for AND, whose second row is empty (no input ever produces 01) and whose first row contains two 1s (both 00 and 01 produce 00). A matrix with exactly one 1 in every row and every column is called a *permutation matrix*, because all it does is shuffle, or permute, the configurations among themselves. For reversible computing, all valid operators are permutation matrices. We will denote a permutation matrix by the symbol $\mathcal{P}$.

Reversibility itself becomes a one-line calculation. An operator is reversed by *transposing* its matrix, swapping its rows and columns, which turns “input j goes to output k ” into “input k goes to output j .” The transpose of an operator is thus sometimes referred to as

its *time-reverse*. If an operator is multiplied by its own time-reverse then, if the operator is reversible, the two must cancel out exactly, yielding the identity matrix:

$$\mathcal{P}\mathcal{P}^\top = \mathbf{I}.$$

This equation is the algebraic signature of a permutation matrix: if multiplying a matrix (of 0s and 1s, one 1 per column) by its own transpose does not yield the identity, the operator was not one-to-one. The reader can check that $\text{SW SW}^\top = \mathbf{I}$ while $\text{AND AND}^\top \neq \mathbf{I}$. Keep the equation $\mathcal{P}\mathcal{P}^\top = \mathbf{I}$ in mind; in Chapter 3 an almost identical equation, $UU^\dagger = \mathbf{I}$, will define the operators of quantum computing, and the family resemblance is no accident.

1.7 Evolving the state

Recall that we are representing the memory of the computer as a state vector. In this representation, evolving the bits by a logic gate is a single matrix multiplication:

$$\mathbf{e}^{[t+1]} = \mathcal{P}^{[t+1 \leftarrow t]} \mathbf{e}^{[t]}.$$

The notation recurs constantly from here on, so it is set out carefully here. The square-bracketed superscript on a vector is a scalar identifying which discrete time interval the vector describes: $\mathbf{e}^{[t]}$ is the state at step t . The superscript on a matrix contains two scalars separated by an arrow: the right-hand scalar is the time interval of the state it expects as input, and the left-hand scalar is the time interval of the state it produces. So $\mathcal{P}^{[t+1 \leftarrow t]}$ carries the system from step t to step $t + 1$. Matrices adorned with transition times like this will be referred to as *transition matrices*, and one whose time difference is a single interval is a *single-step* transition matrix.

If the reader is new to linear algebra, the multiplication can also be expressed element-wise. Using one-based indices, with the first index of a matrix its row and the second its column,

$$e_j^{[t+1]} = \sum_{k=1}^N \mathcal{P}_{jk}^{[t+1 \leftarrow t]} e_k^{[t]} :$$

to get entry j of the new vector, run across row j of the matrix and down the old vector, multiplying entry by entry and summing. Since our state vector has a single 1 at the position of the current configuration, the sum simply picks out the column of the matrix at that configuration, and that column has its 1 at the output configuration. The matrix multiplication is doing nothing more than looking up a row in a truth table.

Multiplying two matrices uses the same rule, run once for each column of the right-hand matrix. Entry (j, k) of the product AB comes from running across row j of A and down column k of B , multiplying entry by entry and summing:

$$(AB)_{jk} = \sum_{l=1}^N A_{jl} B_{lk}.$$

A matrix applied to a vector is the special case where B has a single column. We will need this because a program is a chain of gates, and the whole chain can be collapsed into one matrix before it touches the state.

Matrix multiplication is associative: in a product of several matrices and a vector, it does not matter whether we apply each matrix to the vector in sequence, or first multiply the matrices together and then apply the result to the vector once:

$$\mathbf{e}^{[t+2]} = \mathcal{P}^{[t+2 \leftarrow t+1]} \left(\mathcal{P}^{[t+1 \leftarrow t]} \mathbf{e}^{[t]} \right) = \left(\mathcal{P}^{[t+2 \leftarrow t+1]} \mathcal{P}^{[t+1 \leftarrow t]} \right) \mathbf{e}^{[t]}.$$

This means that, rather than evolving the state vector gate by gate, we can always compose the product of every gate in an entire program into a single transition matrix, and evolving the state by this single matrix yields the same result as evolving it sequentially. We can compose arbitrary multi-step transition matrices from products of single-step ones:

$$\mathcal{P}^{[b \leftarrow a]} = \prod_{t=a}^{b-1} \mathcal{P}^{[t+1 \leftarrow t]},$$

where the product is understood to be taken in the order that places later times to the left. (Order matters: matrix multiplication is associative but not commutative.)

An operator may act on fewer bits than the memory contains: a two-bit gate is a 4×4 matrix, but a three-bit memory's state vector has size 8×1 , and a 4×4 matrix cannot multiply an 8×1 vector, as the former matrix's column count must match the latter's row count. The solution is once again the Kronecker product. In the same way that state vectors of subsystems combine into a state vector of the whole, operators on subsystems combine into an operator on the whole:

$$\mathcal{P}_{\text{MSB,LSB}} = \mathcal{P}_{\text{MSB}} \otimes \mathcal{P}_{\text{LSB}},$$

which has the effect of running the two operators in parallel on their respective bits. If we wish to apply a two-bit operator to just two of three bits and leave the third alone, we take its Kronecker product with the identity matrix in the untouched bit's position. "Do nothing" is an operator too, and it is the identity.

1.8 Circuit notation and a first program

We can now introduce a visual notation for computer programs. Each bit is written as a row, and each column is a discrete time step. The left-most column gives each bit's initial state. An operator written in a bit's row means that operator is applied to that bit at that time step. A multi-bit operator is written once in every row it acts on, aligned in the same column; the shared column is what says the entries belong to one operator rather than two. The final column, marked $\perp$, represents a measurement: the values of the bits are read out. Here is a two-bit program:

$$\begin{array}{cccccc} \mathbf{e}_0 & \text{NOT} & \text{SWAP} & \text{NOT} & \perp \\ \mathbf{e}_0 & & \text{SWAP} & & \perp \end{array}$$

Reading left to right: both bits start at 0; the NOT operator is applied to the first (most significant) bit; then the SWAP operator is applied between the first and second bits; then NOT is applied to the first bit again; then both bits are measured.

The program is computed by translating each column into a matrix (using Kronecker products with the identity where a bit is left alone) and multiplying, right-most first since the initial state stands on the right:

$$(\text{NOT} \otimes \text{I}_2) \text{SW} (\text{NOT} \otimes \text{I}_2) (\mathbf{e}_0 \otimes \mathbf{e}_0) = \cdots$$

Expanding every factor:

$$\begin{aligned} & \left(\begin{pmatrix} 0 & 1 \\ 1 & 0 \end{pmatrix} \otimes \begin{pmatrix} 1 & 0 \\ 0 & 1 \end{pmatrix} \right) \begin{pmatrix} 1 & 0 & 0 & 0 \\ 0 & 0 & 1 & 0 \\ 0 & 1 & 0 & 0 \\ 0 & 0 & 0 & 1 \end{pmatrix} \left(\begin{pmatrix} 0 & 1 \\ 1 & 0 \end{pmatrix} \otimes \begin{pmatrix} 1 & 0 \\ 0 & 1 \end{pmatrix} \right) \left(\begin{pmatrix} 1 \\ 0 \end{pmatrix} \otimes \begin{pmatrix} 1 \\ 0 \end{pmatrix} \right) = \cdots \\ & \begin{pmatrix} 0 & 0 & 1 & 0 \\ 0 & 0 & 0 & 1 \\ 1 & 0 & 0 & 0 \\ 0 & 1 & 0 & 0 \end{pmatrix} \begin{pmatrix} 1 & 0 & 0 & 0 \\ 0 & 0 & 1 & 0 \\ 0 & 1 & 0 & 0 \\ 0 & 0 & 0 & 1 \end{pmatrix} \begin{pmatrix} 0 & 0 & 1 & 0 \\ 0 & 0 & 0 & 1 \\ 1 & 0 & 0 & 0 \\ 0 & 1 & 0 & 0 \end{pmatrix} \begin{pmatrix} 1 \\ 0 \\ 0 \\ 0 \end{pmatrix} = \begin{pmatrix} 0 \\ 0 \\ 0 \\ 1 \end{pmatrix} = \mathbf{e}_{11} = \mathbf{e}_1 \otimes \mathbf{e}_1. \end{aligned}$$

In words: both bits begin with a value of 0. The most significant bit is flipped to 1, giving 10. Its value is then swapped with the least significant bit, giving 01. The most significant bit is flipped a second time, and the final configuration of the two-bit system is 11, which is what the measurement at the end of the circuit reveals. The reader is encouraged to trace the same program through the truth tables directly and confirm that the matrices are just doing the bookkeeping.

Chapter 2

Probabilistic Computing

Probabilistic computing differs from deterministic computing in that logic gates may act as *stochastic perturbations* upon the bits: they can change the values of the computer's memory at random. The bits still always *have* values (at every moment, each bit is a definite 0 or a definite 1), but it becomes impossible for us to track those definite values with certainty. The best we can do is track a statistical distribution over the values they might have.

This distinction, between what the bits are and what we can say about them, will do a great deal of work here. In each individual run of a program, the memory follows some definite trajectory of configurations. A statistical distribution describes something else: the proportions we would find over an *ensemble* of runs, meaning many repetitions of the same experiment. The distribution is our description of the system, not a property carried around by any individual system. A coin about to be flipped has a 50% chance of landing heads, but no individual coin ever lands half-heads-half-tails.

2.1 Probability vectors

Here is where the “wasteful” state vector of Chapter 1 justifies itself. We kept a pigeonhole for every possible configuration of the memory; a statistical distribution is described by simply filling every pigeonhole with the probability of its configuration. The resulting column vector is called a *probability vector*, which we will write $\mathbf{p}$, and it is formed as a linear combination, a weighted sum, of the bit vectors of each possible configuration, each weighted by the probability of that configuration occurring:

$$\mathbf{p} = p_1\mathbf{e}_{\dots 0} + p_2\mathbf{e}_{\dots 1} + p_3\mathbf{e}_{\dots 10} + p_4\mathbf{e}_{\dots 11} + \cdots + p_n\mathbf{e}_{\dots} = \begin{pmatrix} p_1 \\ p_2 \\ p_3 \\ p_4 \\ \vdots \\ p_n \end{pmatrix}.$$

To superimpose vectors in this way merely means to add them. For a probabilistic computer with a single bit of memory:

$$\mathbf{p} = \begin{pmatrix} p_1 \\ p_2 \end{pmatrix} = p_1 \begin{pmatrix} 1 \\ 0 \end{pmatrix} + p_2 \begin{pmatrix} 0 \\ 1 \end{pmatrix}.$$

Naturally, since the elements are probabilities, each must range from 0 to 1 and all of them must sum to 1:

$$\sum_{j=1}^n p_j = 1.$$

The bit vectors of Chapter 1 are themselves valid probability vectors: the special case of complete certainty, with 100% probability on one configuration. A distribution of this form is called *degenerate*. At the other extreme, a distribution giving every configuration equal probability is called *uniform*, and it represents complete uncertainty. Everything else lives in between, and one of our tasks here is to make “in between” quantitative.

To extract a single probability from the vector, multiply by the transpose of the corresponding bit vector, which converts it into a row vector and selects the entry:

$$P(\mathbf{e}_k) = \mathbf{e}_k^T \mathbf{p}.$$

Probability vectors for independent systems combine exactly as bit vectors do, with the Kronecker product:

$$\mathbf{p}_{\text{MSB,LSB}} = \mathbf{p}_{\text{MSB}} \otimes \mathbf{p}_{\text{LSB}}.$$

The reader can check that the arithmetic of the Kronecker product does precisely the right thing: the probability of the joint configuration ab comes out as the product $P(a)P(b)$, which is the correct rule for independent events, and only for independent events. That small caveat is about to become a large theme.

2.2 Marginalization, and what it loses

Probability vectors can also be taken apart in the same way that bit vectors can, using the removal matrices of Chapter 1. When a probability vector is factorized in this way it is referred to specifically as *marginalization*, and the bits removed from the vector are said to be *marginalized over*:

$$\mathbf{p}_{\text{reduced}} = (\mathcal{R}_1 \otimes \mathcal{R}_2 \otimes \cdots \otimes \mathcal{R}_n) \mathbf{p} \quad \text{where } \mathcal{R}_j \in \{\mathbf{I}_2, \begin{pmatrix} 1 & 1 \end{pmatrix}\}.$$

The row vector $\begin{pmatrix} 1 & 1 \end{pmatrix}$ in a bit’s position sums over that bit’s possibilities: “whatever this bit was, add up the probabilities.” The result, called a *marginal* distribution, is the statistical description of the remaining bits considered on their own.

One must be careful when marginalizing, because something can be lost that is invisible in the marginals. Consider the following three probability vectors, each representing a distribution over the configurations of a two-bit memory:

$$\mathbf{p}_{\text{uncorrelated}} = \begin{pmatrix} 0.25 \\ 0.25 \\ 0.25 \\ 0.25 \end{pmatrix} \quad \mathbf{p}_{\text{correlated}} = \begin{pmatrix} 0.5 \\ 0 \\ 0 \\ 0.5 \end{pmatrix} \quad \mathbf{p}_{\text{anti-correlated}} = \begin{pmatrix} 0 \\ 0.5 \\ 0.5 \\ 0 \end{pmatrix}.$$

Now marginalize each of them over the most significant bit to get a reduced vector for the least significant bit, and likewise over the least significant bit to get a reduced vector for the

most significant bit. In all three cases (correlated, anti-correlated, and uncorrelated) we get identical results:

$$\mathbf{p}_{\text{MSB}} = \begin{pmatrix} 0.5 \\ 0.5 \end{pmatrix} \quad \mathbf{p}_{\text{LSB}} = \begin{pmatrix} 0.5 \\ 0.5 \end{pmatrix}.$$

Each bit, considered in isolation, has a 50% chance of being 0 and a 50% chance of being 1 in all three distributions. But the three joint distributions describe radically different situations. In the correlated case, the two bits are guaranteed to agree: the outcome is always 00 or 11. In the anti-correlated case they are guaranteed to disagree: always 01 or 10. In the uncorrelated case, knowing one bit tells you nothing about the other.

The marginals do not record any of this. If we factorize the bits when there are statistical correlations between them, we lose the information about those correlations, and recombining the marginals with the Kronecker product does not give it back: the Kronecker product of the two $(0.5, 0.5)$ marginals yields the uncorrelated joint distribution, whichever of the three we started from:

$$\mathbf{p}_{\text{MSB,LSB}} \begin{cases} = \mathbf{p}_{\text{MSB}} \otimes \mathbf{p}_{\text{LSB}} & \text{if } \mathbf{p}_{\text{MSB,LSB}} = \mathbf{p}_{\text{uncorrelated}} \\ \neq \mathbf{p}_{\text{MSB}} \otimes \mathbf{p}_{\text{LSB}} & \text{otherwise.} \end{cases}$$

A probability vector that cannot be decomposed into a Kronecker product of smaller probability vectors is said to be *non-factorizable*. Bit vectors are always factorizable: a definite configuration is fully described by each bit's definite value, and it is meaningless to speak of correlations between definite values. Statistical distributions are different: correlation is a property of the *joint* distribution that simply does not live anywhere in the individual bits' descriptions. This is why the vector representation earns its exponential size. Tracking each bit's probabilities separately, two numbers per bit, would lose track of the correlative information, and the correlative information is usually the point.

2.3 Stochastic matrices

The logic gates of a probabilistic computer can also be represented by matrices. Suppose that at some step of a program, one of a set of reversible operators is applied at random: permutation matrix $\mathcal{P}_1$ with probability $P(\mathcal{P}_1)$, or $\mathcal{P}_2$ with probability $P(\mathcal{P}_2)$, and so on. Over the ensemble, the effect on the probability vector is described by the weighted sum

$$\Gamma = P(\mathcal{P}_1) \mathcal{P}_1 + P(\mathcal{P}_2) \mathcal{P}_2 + P(\mathcal{P}_3) \mathcal{P}_3 + \cdots + P(\mathcal{P}_n) \mathcal{P}_n.$$

Matrices of this general kind, with non-negative entries describing how probabilities flow from input configurations to output configurations, are referred to as *stochastic matrices*, or Markov matrices. The columns play the same role they did in Chapter 1, but statistically: a stochastic matrix is a statistical truth table, whose column k holds the probability distribution of outputs produced when the input configuration is k .

For example, consider a two-bit gate with a 25% probability of acting as the CX operator and a 75% probability of acting as the SWAP operator. The stochastic matrix representing

this is formed by weighing those matrices by their probabilities and summing:

$$\Gamma = 0.25 \text{ CX} + 0.75 \text{ SW} = \begin{pmatrix} 1 & 0 & 0 & 0 \\ 0 & 0.25 & 0.75 & 0 \\ 0 & 0.75 & 0 & 0.25 \\ 0 & 0 & 0.25 & 0.75 \end{pmatrix}.$$

Notice that each column sums to one, because probability is conserved: whatever goes in must come out somewhere. A matrix with this property is called a *column-stochastic* matrix, and column-stochasticity is the defining requirement for a valid probabilistic operator. Notice that in this example each *row* also sums to one. When both the rows and columns sum to one, the matrix is called *doubly stochastic*. If our linear combination is formed entirely from genuine permutation matrices, that is, from one-to-one maps, the result is always doubly stochastic; indeed, a doubly stochastic matrix can be interpreted as a random selection among reversible operations. If the probabilistic computer is information-preserving in each individual run, its valid operators are exactly the permutation matrices and their mixtures, the doubly stochastic matrices.

Evolution works exactly as before. Using the transition notation of Chapter 1,

$$\mathbf{p}^{[t+1]} = \Gamma^{[t+1 \leftarrow t]} \mathbf{p}^{[t]}, \quad \Gamma^{[b \leftarrow a]} = \prod_{t=a}^{b-1} \Gamma^{[t+1 \leftarrow t]},$$

with multi-step transition matrices composed by matrix products, later times to the left, exactly as with permutation matrices. And smaller stochastic matrices combine into larger ones with the Kronecker product, representing gates applied to different bits in parallel:

$$\Gamma_{\text{MSB,LSB}} = \Gamma_{\text{MSB}} \otimes \Gamma_{\text{LSB}}.$$

Flipping a fair coin can be read two ways, and the second is the useful one here: rather than producing a fresh random face, the flip leaves the coin as it was half the time and turns it over the other half. Written that way it is an operation on the bit already present, and it has a matrix. The COIN operator applies, with equal probability, either the identity or the NOT operator to a single bit:

$$\text{COIN} = 0.5 \text{ I}_2 + 0.5 \text{ NOT} = \begin{pmatrix} 0.5 & 0.5 \\ 0.5 & 0.5 \end{pmatrix}.$$

Whatever distribution goes in, the uniform distribution comes out.

2.4 Factorizing stochastic matrices

We saw how to marginalize a probability vector. Can we do the same to a stochastic matrix? That is: if a gate acts jointly on bits A and B, is there a reduced stochastic matrix describing what the gate does to bit A alone?

There is, and the construction is set out carefully here because its quantum generalization (Chapter 4) is one of the workhorses of quantum information theory.

Start with the most extreme example available: a gate that is not random at all. Let XC denote the controlled-NOT with the roles reversed from CX: the *least* significant bit is the

control, and the most significant bit is the target. Its truth table sends $00 \rightarrow 00$, $01 \rightarrow 11$, $10 \rightarrow 10$, $11 \rightarrow 01$, which as a matrix is

$$XC = \begin{pmatrix} 1 & 0 & 0 & 0 \\ 0 & 0 & 0 & 1 \\ 0 & 0 & 1 & 0 \\ 0 & 1 & 0 & 0 \end{pmatrix} :$$

perfectly deterministic, perfectly reversible.

Now ask: what does XC do to bit A (the most significant bit) considered on its own? The correct answer is that it depends on bit B. If $B = 0$, bit A is left alone. If $B = 1$, bit A is flipped. So partition the columns of the joint truth table by the value of B on the input, and in each partition, marginalize over B on the output: record what happens to A and forget what happened to B. This produces one small matrix per value of the removed bit:

$$K_0 = \begin{pmatrix} 1 & 0 \\ 0 & 1 \end{pmatrix} = I_2, \quad K_1 = \begin{pmatrix} 0 & 1 \\ 1 & 0 \end{pmatrix} = \text{NOT}.$$

Element-wise, for the removal of bit B in general:

$$(K_b)_{a'a} = \sum_{b'} \Gamma_{(a'b'),(ab)},$$

where $(a'b')$ indexes the joint output configuration and (ab) the joint input configuration. Each K_b is itself a valid column-stochastic matrix: its columns sum to one because the columns of the joint matrix did.

If bit B is statistically independent of bit A, with its own distribution $\mathbf{p}_B$, then the reduced dynamics of A is the weighted mixture of these component matrices:

$$\Gamma_A = \sum_b P(B=b) K_b.$$

The joint gate, marginalized down to one bit, becomes a probability distribution over smaller gates, precisely the form of stochastic matrix we introduced in the previous section, with the weights supplied by the discarded bit's statistics rather than by a random number generator.

Try it on XC with bit B uniformly distributed:

$$\Gamma_A = 0.5 I_2 + 0.5 \text{NOT} = \text{COIN}.$$

The result is small and the consequence is not. The COIN operator, the paradigm of randomness, is the *shadow* of a perfectly deterministic, perfectly reversible two-bit gate, seen by an observer who can only watch one of the two bits. Randomness in a subsystem is entirely compatible with determinism in the whole system; ignorance of the rest of the system is indistinguishable, from the inside, from genuine chance. We will meet this moral again when quantum systems shed their randomness into environments in Chapter 4, where it goes by the name of decoherence.

The construction comes with the same warning label as the marginalization of probability vectors: it is only valid when the discarded bit is uncorrelated with the bit that is kept. If A and B are correlated, no reduced stochastic matrix for A alone exists, and it is easy to see

why with XC. Feed it the correlated distribution $0.5\mathbf{e}_{00} + 0.5\mathbf{e}_{11}$: the input 00 leaves A at 0, and the input 11 flips A to 0, so A ends up 0 with certainty. Now feed it the *uncorrelated* distribution with the same marginals (each bit uniform, independent): the four configurations are equally likely, and A ends up uniform. Same marginal input for A in both cases; different marginal output for A. No matrix acting on A's marginal alone can reproduce both, because the joint correlation, which A's marginal does not record, makes a difference to A's future. The dynamics of a subsystem correlated with something we are not tracking is not, in general, expressible as any stochastic matrix on the subsystem; this observation will return with force in Part II.

The component matrices K_b are the classical ancestors of what are called *Kraus operators* in quantum information theory. In Chapter 4 we will meet their quantum form, where the same idea (the reduced dynamics of a subsystem, written as a collection of component operators indexed by the discarded environment) is carried out at the level of amplitudes rather than probabilities.

2.5 Measuring uncertainty: information entropy

If a person were to guess the values of the bits, they would guess the most likely configuration, and their guess would be right more often the farther the distribution is from uniform. “Certainty,” in this sense, is distance from uniformity: a degenerate distribution represents maximal certainty, a uniform distribution maximal uncertainty.

The standard way to quantify this is *Shannon information entropy*. Information entropy can be thought of as the amount of new information a person learns, on average, from measuring the values of the bits. If the bits are in a degenerate distribution, the person already knows the values with certainty, so measuring teaches them nothing: the entropy is 0. If N bits are in a uniform distribution, the person knows nothing in advance, and measuring teaches them everything: the entropy is N , one full bit of information per bit of memory. Anything between degeneracy and uniformity yields an entropy between 0 and N :

$$H(\mathbf{p}) = - \sum_j p_j \log_2 p_j,$$

with the convention that a term with $p_j = 0$ contributes 0. The logarithm base 2 is what denominates the answer in bits.

Entropy is often normalized to a value between 0 and 1 by dividing by N . The normalized entropy is referred to as *efficiency*:

$$\eta(\mathbf{p}) = \frac{- \sum_{j=1}^{2^N} p_j \log_2 p_j}{N}.$$

As it ranges from 0 to 1 it can be read as a percentage. If two people communicate through a channel and the statistical distribution of each bit received has an efficiency below 100%, the recipient is learning less than one bit of information per bit received: there is redundancy in the channel, which implies the channel could be used more efficiently by compression, such as through a Huffman encoding. We will not need compression here, but we will need efficiency as a diagnostic: it is the natural measure of how much of a system's capacity is actually carrying information.

Entropy also interacts with dynamics. Despite doubly stochastic matrices being linear combinations of permutation matrices, each of which is reversible, a probabilistic computer evolving by doubly stochastic operators is *not* reversible. If a computer evolves solely by doubly stochastic operators, its entropy can only increase or stay the same over time; it can never decrease, and so there is a general tendency toward the uniform distribution. The COIN operator shows why: it maps *every* input distribution to the uniform output, a many-to-one mapping on distributions. Algebraically, a doubly stochastic matrix that is not itself a permutation fails the reversibility test of Chapter 1:

$$\Gamma\Gamma^\top \neq \mathbf{I}.$$

Individual runs remain reversible, since some definite permutation happened each time, but which permutation happened is information we do not have, and our *description* degrades accordingly. Entropy increase, in a probabilistic computer, is a fact about descriptions, not about any loss suffered by the bits themselves.

Finally, entropy gives us the tool we have been circling: a measure of correlation. If a system is correlated with another system, they are said to possess *mutual information* about one another. Mutual information is computed by summing the entropies of the marginal (factorized) distributions and subtracting the entropy of the joint (non-factorized) distribution:

$$I(X; Y) = H(X) + H(Y) - H(X, Y).$$

If the joint distribution factorizes, the marginal entropies add up to the joint entropy exactly and the mutual information is 0. Any correlation makes the joint distribution more predictable than its marginals suggest, so $H(X, Y)$ falls short of $H(X) + H(Y)$, and the shortfall is precisely the number of bits of information that knowing one system provides about the other.

For our three two-bit distributions: all have $H(\text{MSB}) = H(\text{LSB}) = 1$. The uncorrelated case has joint entropy 2, so $I = 1 + 1 - 2 = 0$: measuring one bit reveals nothing about the other. The correlated and anti-correlated cases have joint entropy 1, so $I = 1 + 1 - 1 = 1$: measuring one bit reveals the other completely. The correlations that marginalization silently destroyed are now something we can count.

2.6 Programs that create correlation

Let us put the pieces together in a short program for a two-bit probabilistic computer. Recall the circuit notation: rows are bits, columns are time steps, and an overline marks the control bit of a controlled operation, with the target on the unmarked row. This program applies COIN to the most significant bit, then CX between the bits:

$$\begin{array}{cccc} \mathbf{e}_0 & \text{COIN} & \overline{\text{CX}} & \perp \\ \mathbf{e}_0 & & \text{CX} & \perp \end{array}$$

Computing it:

$$\text{CX}(\text{COIN} \otimes \mathbf{I}_2)(\mathbf{e}_0 \otimes \mathbf{e}_0) = \begin{pmatrix} 1 & 0 & 0 & 0 \\ 0 & 1 & 0 & 0 \\ 0 & 0 & 0 & 1 \\ 0 & 0 & 1 & 0 \end{pmatrix} \begin{pmatrix} 0.5 & 0 & 0.5 & 0 \\ 0 & 0.5 & 0 & 0.5 \\ 0.5 & 0 & 0.5 & 0 \\ 0 & 0.5 & 0 & 0.5 \end{pmatrix} \begin{pmatrix} 1 \\ 0 \\ 0 \\ 0 \end{pmatrix} = \begin{pmatrix} 0.5 \\ 0 \\ 0 \\ 0.5 \end{pmatrix} = \mathbf{p}_{\text{correlated}}.$$

The result is a statistical distribution with 50% for 00 and 50% for 11: both bits have a 50% chance of being 0 or 1, but they are guaranteed to have the same value. This should be rather intuitive: we randomly flipped the most significant bit, and then flipped the least significant bit if and only if the most significant bit came up 1. The CX gate *records* the first bit's random outcome onto the second, and recording is correlating: the program manufactures one full bit of mutual information out of one coin flip.

If we negate the least significant bit at the start, the same program yields the anti-correlated distribution instead; if we apply COIN to both bits, the result is uniform and uncorrelated, as each bit's fresh coin flip destroys any relation to the other. Each variant is a one-line change to the matrices above.

In between the perfectly correlated and perfectly uncorrelated cases, mutual information ranges continuously between 1 and 0, and measuring one bit of the pair improves one's certainty about the other. We will make "improves one's certainty" precise before the chapter ends. But first, having assembled programs out of stochastic matrices, we should ask what it takes to actually *run* one.

2.7 Simulating a Markov chain: stochastic jumps

Everything so far has evolved the probability vector: multiply by a stochastic matrix, read off the new distribution. That is the ensemble description, and we have insisted from the start that it is only that: in each individual run of a probabilistic computer, the bits hold definite values at every moment. A program of the kind we have been writing, a sequence of stochastic updates in which each step needs nothing beyond the current configuration, is called a *Markov chain* (Chapter 5 will have a great deal to say about this property). What we have not yet done is simulate one properly: produce not the distribution over runs but a single run, one definite history the machine could actually live through.

The whole simulation is one mechanism, repeated. Start the bits in a definite configuration, sampled from the initial distribution. Now suppose the bits occupy the definite configuration j when a gate arrives. Column j of the gate's stochastic matrix is the probability distribution over successor configurations, since a stochastic matrix is a statistical truth table and we are looking up its j -th entry. To sample from that column, form its *cumulative distribution function*: the running total of its entries,

$$F(0) = 0, \quad F(i) = \sum_{k=1}^i \Gamma_{kj},$$


which climbs from 0 to 1 as i runs over the configurations. Then draw a single random number γ , uniform on $[0, 1)$, and select the configuration whose slice of the running total contains it:

$$x' = i \iff F(i-1) \leq \gamma < F(i).$$

An intuitive way to conceive of this is an analog clock whose face has been divided into 2^N zones, one per configuration, each zone sized in proportion to that configuration's probability in the column. The value γ is where the hand points; whichever zone it lands in is the system's next definite configuration. One uniform random number per gate, one lookup, one hop.

Run the correlation program of the previous section this way. The bits start at 00, definitely. The COIN step's column for 00 splits its probability evenly between 00 and 10, and γ decides; suppose it lands on 10. The CX step's column for 10 is degenerate, all of its probability on 11, so the hop is forced. This run of the program printed 11. Another run, whose γ fell the other way, prints 00. No run ever prints 01 or 10, and over many runs 00 and 11 appear in equal proportion: by the law of large numbers, the ensemble of histories converges to $\mathbf{p}_{\text{correlated}}$, the same answer the linear algebra delivered in one pass. The two descriptions are one theory. And the correlation itself, which no marginal can see, is plainly visible in every single history: whatever the first bit did, the second did the same.

All the histories can be drawn at once as a tree. Each branch is a path the bit can take, labeled with the probability of taking it, and the probability of ending somewhere is the sum over the paths that arrive there. Take the COIN operator applied twice to a bit that starts at 0:



Each of the four paths carries $\frac{1}{2} \times \frac{1}{2} = \frac{1}{4}$. Two of them end at 0 and two at 1, so each outcome has probability $\frac{1}{4} + \frac{1}{4} = \frac{1}{2}$: the bit is still a fair coin, as COIN COIN = COIN already told us. The point to carry forward is the arithmetic at the bottom. Path weights are probabilities, probabilities are never negative, and so the contributions arriving at a configuration can only *add*. Nothing a probabilistic computer does can make two paths meet and destroy one another.

2.8 Local gates without giant matrices

There is one more piece of craft to collect before returning to probability proper, and it is the piece that makes everything here computable at scale. The vector evolution looks innocent on two bits and becomes alarming on twenty. The vector for N bits has 2^N entries, which is bad enough. But write a one-bit gate the way we have been writing it,

$$\mathbf{I}_2 \otimes \cdots \otimes S \otimes \cdots \otimes \mathbf{I}_2,$$

and the resulting matrix has $2^N \times 2^N = 4^N$ entries. At $N = 20$ the vector occupies a few megabytes while the matrix, built out literally, occupies terabytes; and all of it describes an operation that touches one bit and, by construction, leaves the other nineteen alone.

The structure of the gate is also the cure. A gate acting on one bit cannot move probability between configurations that differ anywhere else, so each configuration can trade probability with exactly one partner: its *twin*, the configuration identical to it except at the target bit.

The giant matrix, whatever it looks like written out, is secretly the little matrix S stamped once onto each twin pair: probability flows within pairs and never between them. So never build the matrix. Sweep through the 2^{N-1} twin pairs, and within each pair (x_0, x_1) , named for the target bit's value in each, apply S directly:

$$\begin{pmatrix} p'(x_0) \\ p'(x_1) \end{pmatrix} = S \begin{pmatrix} p(x_0) \\ p(x_1) \end{pmatrix}.$$

Count the arithmetic: 2^{N-1} pairs, one 2×2 multiplication each, roughly 2^{N+1} multiplications in total, where building and applying the full matrix costs on the order of 4^N . The same idea handles a gate on k bits: the configurations organize into 2^{N-k} families of 2^k mutual twins, and applying the small matrix within each family costs 2^{N+k} operations. The vector remains exponentially large; nothing removes that. The matrix, however, never has to exist. Two further courtesies come free. The pairs do not interact, so they can be updated in place, one pair at a time, or handed out to as many processors as are available, since no pair ever needs another's entries.

For the jump simulation of the previous section, the simplification goes further still. A hop needs only one column of the gate's matrix, and locality makes that column almost entirely empty: from a definite configuration, only the target bit is in play, so the hop reduces to a single biased coin flip, the column of S selected by the target bit's current value. No 2^N appears anywhere. A single run of a probabilistic program therefore costs scarcely more than a run of the deterministic computer of Chapter 1 that it generalizes. The moral, stated carefully: *individual histories are easy; what is exponential is the ensemble description*, the wish to track every alternative at once. The world, which runs one history at a time, never faces the exponential burden; only the bookkeeper who insists on all the alternatives simultaneously does. Part II will return to this asymmetry at the exact point where quantum circuits appear to refuse the simplification.

One item of machinery now remains: the rule for updating a probability distribution on new information. It is what makes “improves one's certainty” precise, and everything that follows leans on it.

2.9 Bayes' theorem: updating on what you learn

Each entry of the probability vector can be conceived of as the probability assigned to a *hypothesis* about the computer's configuration: the first entry to the hypothesis “the memory is all 0s,” the second to “the memory is all 0s except the least significant bit,” and so forth:

$$\mathbf{p} = \begin{pmatrix} P(H_0) \\ P(H_1) \\ P(H_2) \\ P(H_3) \\ \vdots \end{pmatrix}.$$

When a measurement reveals something about the configuration, our certainty should improve, and Bayes' theorem tells us how to update the distribution to reflect exactly what was learned

and nothing more:

$$P(H_k | E) = \frac{P(E | H_k) P(H_k)}{P(E)}, \quad P(E) = \sum_k P(E | H_k) P(H_k),$$

where E is the evidence acquired. (The letter H here stands for “hypothesis” and has nothing to do with the H of entropy; the collision of notation is regrettable but standard.)

For measurements of bit values, the theorem takes an especially simple form, because the probability of the evidence given a hypothesis is always either 0% or 100%: either the hypothesized configuration is compatible with what was observed, or it is not. Suppose a two-bit computer is in the distribution

$$\mathbf{p} = \begin{pmatrix} 0.4 \\ 0.1 \\ 0.1 \\ 0.4 \end{pmatrix}$$

and the most significant bit is measured to have a value of 1. The evidence is incompatible with the hypotheses 00 and 01, so $P(E | H_{00}) = P(E | H_{01}) = 0$, and compatible with 10 and 11, so $P(E | H_{10}) = P(E | H_{11}) = 1$. The probability of the evidence is $P(E) = 0+0+0.1+0.4 = 0.5$, and Bayes’ theorem, applied entry by entry, gives

$$P(H_k | E) = \frac{1}{0.5} \begin{pmatrix} 0 \\ 0 \\ 0.1 \\ 0.4 \end{pmatrix} = \begin{pmatrix} 0 \\ 0 \\ 0.2 \\ 0.8 \end{pmatrix}.$$

Look at what the calculation actually did. The numerator zeroed out the entries incompatible with the observation. The denominator, being exactly the sum of the surviving entries, divided the vector by its own sum. Dividing a vector by its own sum is called *normalization*,

$$\|\mathbf{p}\| = \frac{\mathbf{p}}{\sum_{j=1}^n p_j},$$

and it guarantees the entries again sum to 1. So for measurements of bits, Bayes’ theorem is simply: *zero out the configurations incompatible with what you observed, then normalize*. Nothing about this rule is mysterious, physical, or dynamical. It is what any accurate description must do when new information arrives. We belabor this point because, two chapters from now, the same zero-and-normalize rule applied to a different vector will be called “the collapse of the wavefunction,” and it will be billed by many authors as the deepest mystery in physics.

Entropy lets us audit the update. The entropy of $\mathbf{p}$ before the measurement is

$$H(\mathbf{p}) = -(0.4 \log_2 0.4 + 0.1 \log_2 0.1 + 0.1 \log_2 0.1 + 0.4 \log_2 0.4) \approx 1.7219.$$

After the update it is 0.7219 (rounded): a decrease of exactly 1. That is no accident. The marginal distribution of the measured bit was (0.5, 0.5), one full bit of uncertainty, and measuring it delivered one full bit of information. Information gained is entropy lost, and a measurement of a single bit can decrease entropy by at most 1.

The audit extends to the bit that was not measured. Before the measurement, its marginal entropy was $H(0.5, 0.5) = 1$; after the measurement, its marginal distribution is $(0.2, 0.8)$, with entropy 0.7219 (rounded). Measuring the first bit improved our certainty about the second by $1 - 0.7219 = 0.2781$ bits. And if we compute the mutual information between the two bits in the original distribution,

$$I(\text{MSB}; \text{LSB}) = 1 + 1 - 1.7219 \approx 0.2781,$$

it is the same number. Mutual information is exactly the amount by which measuring one system sharpens our description of the other.

2.10 Sufficient statistics

When can a *summary* of data stand in for the data?

Start with the classic setting. A coin of unknown bias is flipped 1,000 times, and the full data set is the exact sequence of heads and tails, in order. Suppose what we care about is the coin's bias. Then the sequence is absurdly more information than we need: the single number "603 heads" serves every purpose the sequence serves. Any inference about the bias that could be drawn from the full sequence can be drawn from the count alone, because for a fixed number of heads, *which* flips came up heads carries no information about the bias: given the count, the ordering is just shuffling. A function of the data with this property, here the count, is called a *sufficient statistic*. The formal definition says exactly this: a statistic T is sufficient (for some question) when the data's remaining detail, once T is known, is irrelevant to that question.

Three features of the concept matter to us.

First, sufficiency is always *relative to a question*. The count of heads is sufficient for estimating the bias; it is useless for the question "did the coin alternate?" A summary discards detail, and whether the discarded detail mattered depends on what you ask.

Second, a sufficient statistic is a *compression*, and usually an enormous one. The sequence of 1,000 flips is 1,000 bits; the count fits in ten. Nothing is lost, for the question at hand, yet almost everything is thrown away.

Third, and most importantly for us: a sufficient statistic is *not a new physical object*. The count of heads is not some additional entity that travels alongside the coin. It is a function of the history, computed by whoever finds it convenient, discardable by anyone willing to lug the full history around instead.

The version of sufficiency we will lean on is the *predictive* one, where the "question" is the future of a process. Consider a walker on a long line of sites, taking one random step at a time, left or right with even odds. The walker's history is the full sequence of steps; but to predict the distribution of its location after the next step, or after any number of future steps, the only thing anyone needs from that history is a single number: the current position. Given the position, the path that led to it is irrelevant to everything that comes next; step histories that arrive at the same site are interchangeable as far as the future is concerned. The position is a sufficient statistic of the history for prediction.

Control theory and engineering have a name for a summary with exactly this property: they call it the *state* of the system. In the engineer's usage, the state is, by definition, whatever function of the past is sufficient to predict the future: no more, no less. It is a

superbly useful concept, and it carries no ontological commitment whatsoever: the engineer's "state" is a bookkeeping construct, an interface between past and future, not a substance. When the process is Markovian, the present configuration already is such a summary, and no history-keeping is needed at all. When a process is *not* Markovian, meaning its future genuinely depends on more of its past than the present configuration reveals, then predicting it requires either carrying the relevant history, or compressing that history into a richer sufficient statistic and carrying that.

Keep this distinction in view; Part II depends upon it. There, we will meet a dynamics whose future is not fixed by the present probability distribution alone. It will turn out that a certain compact object, famous, complex-valued, and widely believed to be the physical furniture of the universe, is a sufficient statistic of the system's statistical history, in precisely the walker's sense: a compression of the past, maintained because it is convenient, eliminable by anyone willing to carry the history instead. Before we can meet it, we need to build the computer it lives in.

Code: this chapter in Octave

Everything in this chapter can be verified in a few lines of GNU Octave.

```
% --- the gates ---
e0 = [1;0]; I2 = eye(2); NOT = [0 1; 1 0];
COIN = 0.5*I2 + 0.5*NOT;
CX = [1 0 0 0; 0 1 0 0; 0 0 0 1; 0 0 1 0];

% --- Section 2.6: the correlation program ---
p = CX * kron(COIN, I2) * kron(e0, e0) % -> [0.5; 0; 0; 0.5]
pMSB = kron(I2, [1 1]) * p % marginalize out the LSB
pLSB = kron([1 1], I2) * p % marginalize out the MSB

% --- Section 2.7: stochastic jumps through the same program ---
hop = @(G, x) find(rand() < cumsum(G(:,x)), 1); % one gamma, one hop

M = 100000; counts = zeros(4,1);
for run = 1:M
    x = hop(kron(COIN, I2), 1); % start at 00, coin the MSB
    x = hop(CX, x); % record it onto the LSB
    counts(x) = counts(x) + 1;
end
counts / M % -> [0.5; 0; 0; 0.5]

% --- Section 2.8: twin-pair update, no giant matrix ---
p = kron(e0, e0); q = 1; % COIN on target bit q (the MSB)
for x0 = 0:3
    if bitand(bitshift(x0, -q), 1) == 0
        x1 = bitxor(x0, 2^q); % the twin: bit q flipped
        pr = COIN * [p(x0+1); p(x1+1)]; % the 2x2 gate on the pair
        p(x0+1) = pr(1); p(x1+1) = pr(2);
    end
end
p % = kron(COIN, I2) * kron(e0, e0)

% --- Sections 2.5 and 2.9: entropy, mutual information, Bayes ---
H = @(p) -sum(p(p>0) .* log2(p(p>0)));

p = [0.4; 0.1; 0.1; 0.4];
Imi = H(kron(I2,[1 1])*p) + H(kron([1 1],I2)*p) - H(p) % 0.2781

keep = [0; 0; 1; 1]; % observed: MSB = 1
q = p .* keep; q = q / sum(q) % -> [0; 0; 0.2; 0.8]
bits_learned = H(p) - H(q) % 1 bit

% --- Section 2.4: classical Kraus decomposition ---
```

```
XC = [1 0 0 0; 0 0 0 1; 0 0 1 0; 0 1 0 0]; % LSB controls, MSB target
sumB = kron(I2, [1 1]); % forget B on the output
K0 = sumB * XC * kron(I2, [1;0]) % B input = 0 -> identity
K1 = sumB * XC * kron(I2, [0;1]) % B input = 1 -> NOT
pB = [0.5; 0.5];
GA = pB(1)*K0 + pB(2)*K1 % -> COIN
```

Chapter 3

Quantum Computing

Quantum computing differs from probabilistic computing in a mathematically insignificant way. Yet, despite the difference being trivial mathematically, it is non-trivial ontologically. Quantum systems behave in a way that defies basic intuition to such a degree that decades of books and papers have been written trying to make sense of it on a philosophical level.

The entire difference is in what each entry of the state vector holds. The probability vector of Chapter 2 assigned one number to each configuration of the memory. The corresponding vector in quantum computing assigns to each configuration a quantity with *two* degrees of freedom instead of one. That is all. Every other piece of structure (the Kronecker product, the transition matrices, the zero-and-normalize update rule) carries over almost unchanged. Our work here is to lay out the formalism plainly, compute with it until it feels routine, and mark precisely the spots where something strange peeks through. Explaining the strangeness is the business of Part II; our job is only to make sure that when the strangeness arrives, we will know exactly which equation it lives in.

3.1 A second degree of freedom

A quantity with two degrees of freedom is most conveniently represented as a *complex number*. A reader comfortable with complex numbers can skim to the next section.

A complex number is a pair of real numbers packaged as

$$z = a + bi,$$

where the symbol i is defined by the property $i^2 = -1$. The number a is called the real part and b the imaginary part; the names are historical accidents and carry no metaphysics: z is simply a point with coordinates (a, b) on a plane, called the complex plane. Addition works coordinate-wise, and multiplication follows from expanding brackets and replacing i^2 with -1 :

$$(a + bi)(c + di) = (ac - bd) + (ad + bc)i.$$

The *conjugate* of z , written $\bar{z}$, flips the sign of the imaginary part: $\bar{z} = a - bi$. Multiplying a number by its own conjugate always yields a non-negative real number,

$$z\bar{z} = (a + bi)(a - bi) = a^2 + b^2,$$

and its square root, $|z| = \sqrt{a^2 + b^2}$, is called the *magnitude* of z : the distance from the origin to the point (a, b) . We will constantly need the *square magnitude*, $|z|^2 = z\bar{z} = a^2 + b^2$.

The same point on the plane can be described by polar coordinates instead: its distance r from the origin and the angle φ of the line from the origin to it,

$$z = r(\cos \varphi + i \sin \varphi) = re^{i\varphi},$$

where the last equality, Euler's formula, can be taken for our purposes as a definition of the notation $e^{i\varphi}$: the point on the unit circle at angle φ . In this representation, $r = |z|$ is the magnitude and the angle $\varphi = \arg(z)$ is called the *argument*, or, in the word we will use throughout, the *phase*. Multiplication is beautifully simple in polar form: magnitudes multiply, and phases add:

$$(r_1 e^{i\varphi_1}) (r_2 e^{i\varphi_2}) = r_1 r_2 e^{i(\varphi_1 + \varphi_2)}.$$

A complex number, then, is a magnitude and a direction on a plane; and the one genuinely new capability this gives us over the non-negative real numbers of Chapter 2 is the possibility of *cancellation*. Two positive probabilities can only ever add. Two complex numbers of equal magnitude and opposite phase sum to zero. Cancellation between amplitudes is what the word *interference* will name throughout this book, and it is the seed of every interference phenomenon in it. Section 3.4 draws the first instance: two paths of equal size and opposite phase, arriving at the same outcome and adding to nothing.

3.2 The quantum state

The major distinction between quantum computing and probabilistic computing is that the vector representing the system's statistical state has two degrees of freedom per entry rather than one. This vector is called the state vector, the quantum state, or the wavefunction; we will just refer to it as the quantum state. As its entries have two degrees of freedom each, it is almost always represented as a complex-valued vector:

$$\psi = \begin{pmatrix} a_1 \\ a_2 \\ a_3 \\ \vdots \\ a_n \end{pmatrix} \quad \text{where } \psi \in \mathbb{C}^n.$$

The complex entries are referred to as *amplitudes*, or sometimes as probability amplitudes. Since an amplitude is a complex number, it has the polar form of Section 3.1: a magnitude and a phase. These are not new quantities invented for quantum mechanics. The magnitude is the same magnitude, the phase is the same phase, and the cancellation Section 3.1 described in the abstract, two numbers of equal magnitude and opposite phase summing to zero, is about to happen between amplitudes attached to actual configurations of actual bits. The second degree of freedom per entry, the one a probability vector does not have, is exactly this phase.

Recall the interpretive ladder so far. The entries of a bit vector marked the definite configuration of the bits. The entries of a probability vector gave the probability of each

configuration. What do the entries of the quantum state refer to? What even *is* an amplitude? Notice that the question does not answer itself, and the formalism alone will not answer it either. The reading of the situation we will defend must be earned, so for now we simply take the formalism as given and learn its rules.

The first rule connects amplitudes to something observable. If we multiply the transpose of a bit vector by a probability vector, we extract the probability of that configuration. The rule for the quantum state is the same, except we must take the square magnitude of the result:

$$P(\mathbf{e}_k) = \left| \mathbf{e}_k^T \psi \right|^2.$$

This rule is called the *Born rule*: the square magnitudes of the amplitudes are probabilities.¹ It follows that where the probability vector was constrained to have entries summing to 1, the quantum state is constrained to have *square magnitudes* summing to 1:

$$\sum_{k=1}^n |\psi_k|^2 = 1 \quad \text{where } \psi \in \mathbb{C}^n.$$

The rule has an immediate consequence about what the quantum state does *not* fix. Multiply every entry by one and the same unit complex number $e^{i\theta}$. Each square magnitude is untouched, since $|e^{i\theta}\psi_k|^2 = |\psi_k|^2$, so every Born probability is untouched; and it stays untouched however long a circuit runs afterwards, because a common factor passes straight through matrix multiplication: $U(e^{i\theta}\psi) = e^{i\theta}(U\psi)$. Two states differing by such a factor therefore agree in every prediction anyone can ever test. That factor is called a *global phase*, and it carries no physical information at all; ψ and $e^{i\theta}\psi$ are two names for one state of affairs. What can matter is the difference between the phases of two entries, called their *relative phase*, which is why the cancellation of Section 3.1 needed two amplitudes and not one. A global phase shifts every entry alike, so it leaves every relative phase exactly where it was.

A redundancy of this shape, where a whole family of descriptions answers to one physical situation, is what physicists call a *gauge* freedom, and this one has a name. The numbers $e^{i\theta}$ are the complex numbers of magnitude 1, and under multiplication they form a group written $U(1)$; since a single θ multiplies every entry of ψ at once, the quantum state is said to carry a *global* $U(1)$ freedom. That word *global* matters: letting θ vary from place to place instead would make the freedom *local*, and local $U(1)$ is a different and much larger subject that nothing here will need.

Every bit vector is a valid quantum state: a single entry of 1, square magnitude 1, everything else 0. A quantum state of this form (allowing also a single nonzero entry of magnitude 1 with any phase, such as -1 or i) plays the role the degenerate distribution played in Chapter 2: it represents a configuration the Born rule makes certain. Such states are called *eigenstates* of the computational basis. (The word comes from a longer story about observables and linear operators, which Section 3.4 tells; for now it is just a name.) A quantum state with more than one nonzero amplitude is called a *superposition* of configurations, literally a weighted sum of basis vectors, exactly as a probability vector was:

$$\psi = a_1\mathbf{e}_{\dots 0} + a_2\mathbf{e}_{\dots 1} + a_3\mathbf{e}_{\dots 10} + \dots$$

¹M. Born, “Zur Quantenmechanik der Stoßvorgänge,” *Zeitschrift für Physik* 37 (1926), 863–867, and “Quantenmechanik der Stoßvorgänge,” *Zeitschrift für Physik* 38 (1926), 803–827.

We flag now, and will argue at length later, that “superposition” is a term for a mathematical shape (a vector with several nonzero entries) and not, on its own, a physical claim about a system being in several configurations at once. Chapter 2 taught us to read a spread-out vector as a description of an ensemble; whether that reading survives in the quantum case is exactly what Parts II and III are about.

Quantum states of independent systems combine with the Kronecker product, as everything here does:

$$\psi_{\text{MSB,LSB}} = \psi_{\text{MSB}} \otimes \psi_{\text{LSB}},$$

and a joint quantum state that cannot be written as a Kronecker product of smaller quantum states is called *entangled*: the amplitude-level analogue of the non-factorizable, correlated probability vectors of Chapter 2. We will construct one shortly.

3.3 Unitary operators

The operators of quantum computing are also complex-valued matrices,

$$U = \begin{pmatrix} a_{11} & a_{12} & \cdots & a_{1n} \\ a_{21} & a_{22} & \cdots & a_{2n} \\ \vdots & \vdots & \ddots & \vdots \\ a_{n1} & a_{n2} & \cdots & a_{nn} \end{pmatrix} \quad \text{where } U \in \mathbb{C}^{n \times n},$$

and the valid ones are called *unitary* operators. Unitarity is defined by direct analogy with the reversibility test of Chapter 1. There, an operator was reversible when multiplying it by its own transpose, its time-reverse, yielded the identity: $\mathcal{P}\mathcal{P}^\top = \text{I}$. For a complex matrix, the time-reverse is the *Hermitian transpose* $U^\dagger$: the ordinary transpose followed by conjugating every entry (flipping the sign of its imaginary part). A matrix is unitary exactly when

$$UU^\dagger = \text{I}.$$

Place this alongside what we know. Permutation matrices satisfy $\mathcal{P}\mathcal{P}^\top = \text{I}$: deterministic and reversible. Doubly stochastic matrices are mixtures of permutations, but themselves fail the test, $\Gamma\Gamma^\top \neq \text{I}$: probabilistic and irreversible, with entropy increasing toward uniformity. Unitary matrices pass the test. So quantum operators produce statistical behavior (the Born rule hands us probabilities, not certainties) while remaining, as operators, perfectly reversible: every quantum program can be run backwards by taking the Hermitian transpose of each gate and reversing their order, and no information is ever lost. A randomness that loses nothing is not the randomness of Chapter 2, where doubly stochastic evolution ratcheted entropy upward. Something is off, and the plain statement of it is this: *unitary matrices are the complex generalization of permutation matrices, not of stochastic matrices*. They shuffle; they do not smear. What exactly they shuffle, given that the observable outcomes are statistical, is the puzzle we are assembling piece by piece.

Beyond this one change of test, everything is as before. Unitary operators combine with the Kronecker product,

$$U_{\text{MSB,LSB}} = U_{\text{MSB}} \otimes U_{\text{LSB}},$$

identity matrices pad out the bits an operator leaves alone, and the discrete time-evolution rule is the same equation we have written twice already:

$$\psi^{[t+1]} = U^{[t+1 \leftarrow t]} \psi^{[t]}, \quad U^{[b \leftarrow a]} = \prod_{t=a}^{b-1} U^{[t+1 \leftarrow t]},$$

with multi-step operators composed by matrix products, later times to the left. Every permutation matrix is unitary (its entries are real, so the Hermitian transpose is just the transpose), so all of reversible computing (NOT, SWAP, CX, the ancilla constructions) imports into quantum computing wholesale.

3.4 Gates and observables

Below are the most common single-bit unitary operators. First the three *Pauli operators*:

$$X = \begin{pmatrix} 0 & 1 \\ 1 & 0 \end{pmatrix} \quad Y = \begin{pmatrix} 0 & -i \\ i & 0 \end{pmatrix} \quad Z = \begin{pmatrix} 1 & 0 \\ 0 & -1 \end{pmatrix}.$$

The Pauli-X operator is simply the NOT gate under a new name. The Pauli-Z operator does nothing to the probabilities at all: it multiplies the amplitude of configuration 1 by -1 , leaving every square magnitude untouched. The Pauli-Y does both at once. That an operator which “does nothing observable” is included among the fundamental gates is itself something to keep in mind.

These three have a second job. Each is its own Hermitian transpose, $X^\dagger = X$ and likewise for the other two, and a matrix of that kind is called an *observable*: a quantity a measurement can return a value for.

Which values it can return is fixed by the matrix. A state that a matrix merely multiplies by a number, $A\psi = \lambda\psi$, is an *eigenstate* of that matrix in the sense Section 3.2 promised to explain later, and the number λ is its *eigenvalue*. The eigenvalues are the values the measurement can return. Each Pauli operator squares to the identity, $X^2 = I$ and likewise for the others, and that leaves only two numbers available, since a number whose square is 1 is either $+1$ or -1 . So measuring any of the three returns $+1$ or -1 and nothing else. That is how a bit’s two outcomes are written when an observable rather than a configuration is being discussed.

Three observables rather than one, for a thing with two outcomes, has a physical reason. Take the quantity these operators were first written for, the spin of an electron, measured with a device called a *Stern–Gerlach apparatus*: the electron passes through and is deflected either one way or the other relative to how the device is turned. Two outcomes, so one bit, and we assign 0 to one deflection and 1 to the other. The device can then be physically rotated to measure the deflection along a different axis. Space is three-dimensional, so there are three mutually perpendicular directions it can be turned to, and the three measurements they give are what the three Pauli operators stand for: the X basis, the Y basis, and the Z basis. Reading a bit on a quantum computer is the same situation without the magnet. Which of the three a measurement asks for is a setting of the instrument, chosen by whoever performs it, and the question it asks changes with the setting.

The *expected value* of an observable, written $\langle A \rangle$, is the sum of its two outcomes with each weighted by how likely it is to occur. It runs from -1 to $+1$, and a value of 0 means the two are equally likely, since $0.5(+1) + 0.5(-1) = 0$. The state supplies it directly:


$$\langle A \rangle = \psi^\dagger A \psi.$$

For Z it is the difference of the two Born probabilities, $\langle Z \rangle = P(0) - P(1)$: $+1$ on $\mathbf{e}_0$, -1 on $\mathbf{e}_1$, and 0 on any state whose two square magnitudes are equal. An expected value need not be one of the outcomes, and 0 never is.

The three expected values together locate the bit. For any state of a single bit,

$$\langle X \rangle^2 + \langle Y \rangle^2 + \langle Z \rangle^2 = 1,$$

so the triple $(\langle X \rangle, \langle Y \rangle, \langle Z \rangle)$ is a point on a sphere of radius one. That sphere is the standard picture of one bit, with an axis for each observable, and it is called the *Bloch sphere* after Felix Bloch, who put the states of a two-level system on it in the course of working out nuclear induction.²



The poles of each axis are the states that return $+1$ and -1 with certainty for the observable on that axis. On the vertical axis those are $\mathbf{e}_0$ and $\mathbf{e}_1$. The equator is everything with $\langle Z \rangle = 0$: the states whose two square magnitudes are equal, differing from one another only in phase. The basis a state is written on is a choice of which observables it reports. By convention an unadorned measurement, the $\perp$ ending every circuit so far, is a measurement on the Z basis, which for a quantum computer is the *computational basis* of Section 3.2. So a bit certainly at 1 has $\langle Z \rangle = -1$, while $\langle X \rangle$ and $\langle Y \rangle$ are both 0 : the bottom pole, and nowhere along the other two axes. Turning the instrument to X or Y instead is the rotated measurement Chapter 10 builds its experiment from.

Two states are *orthogonal* when their overlap $\psi^\dagger \phi$ is zero, which is to say that a measurement certain to report one never reports the other. The two states a given observable is certain about, its eigenstates, are always orthogonal, and two orthogonal vectors are enough to assemble any state of one bit, so they are a basis in the sense of Section 1.1. That is the connection between an observable and a basis, and why a bit has three of the latter rather than one: a basis apiece for X , for Y and for Z .

²F. Bloch, “Nuclear Induction,” *Physical Review* 70 (1946), 460–474.

Two observables need not be certain together, and matrix multiplication is where that shows. Operators *commute* when $AB = BA$, and the Pauli operators do not: $XZ - ZX \neq 0$. Observables that fail to commute share no eigenstate, so no state is certain about both, which is why fixing a bit's value on Z leaves X and Y at 0 rather than at some other definite pair. Nothing has been hidden from view; the two questions cannot both have answers at once.

That restriction is looser than it first appears once there is more than one bit. An observable of several bits is assembled the way operators have been since Section 1.2, by Kronecker product: choose one of I , X , Y or Z for each bit and multiply the choices together. Such an operator is a *Pauli string*, and we write it by its letters, so that XX means $X \otimes X$ and XI means X on the first bit and nothing asked of the second. What a string returns is a parity: read each bit on the basis its letter names, skip the bits carrying I , and multiply the ± 1 outcomes.

Now XX and ZZ do commute, even though X and Z do not, because swapping the order flips a sign in each of the two positions and the two flips cancel. Strings that disagree in an even number of places commute; those disagreeing in an odd number, like XI against ZZ , do not. So a pair of bits can be certain about XX and about ZZ at once while neither bit is certain about anything on its own: a string can report a parity while reporting nothing about the readings inside it. Section 7.5 makes that the whole basis of a bookkeeping, and Chapter 10 makes it the heart of an argument.

The picture is drawn in three dimensions but does not carry three degrees of freedom. Every point on it lies at distance 1 from the center, which uses up one of the three, and locating a point on the surface takes the two that remain: two angles. A single bit's quantum state has two degrees of freedom however one chooses to write it. Its two complex amplitudes are four real numbers, normalization fixes one of them, and the global phase of Section 3.2 changes nothing observable, which leaves two. The sphere fits a bit exactly because the counts agree.

They stop agreeing at the second bit, and the reason will matter later. Spheres grow one per bit while the state's degrees of freedom double with each bit added. Three bits drawn as three spheres would offer six, while the three-bit state carries fourteen: its eight complex amplitudes are sixteen real numbers, less the same two. Six numbers cannot hold fourteen, and no arrangement of three spheres will do it.

A sphere apiece suffices only while the bits stay independent of one another. That is the same restriction Section 2.2 found for ordinary probability, where separate descriptions of correlated bits cannot reproduce the joint one and the correlative information lives nowhere in the parts. A Bloch sphere for each bit is the quantum form of a probability for each bit, and it fails in the same place and for the same reason. What this counting means for the ontology of the quantum state, for what it could be a description of, becomes relevant in Section 8.4.

Next the Hadamard operator and the two phase gates:

$$H = \frac{1}{\sqrt{2}} \begin{pmatrix} 1 & 1 \\ 1 & -1 \end{pmatrix} \quad S = \begin{pmatrix} 1 & 0 \\ 0 & i \end{pmatrix} \quad T = \begin{pmatrix} 1 & 0 \\ 0 & \frac{1+i}{\sqrt{2}} \end{pmatrix}.$$

And finally the rotation operators, parameterized families that interpolate continuously where the gates above jump:

$$R_X(\theta) = \begin{pmatrix} \cos(\theta/2) & -i \sin(\theta/2) \\ -i \sin(\theta/2) & \cos(\theta/2) \end{pmatrix} \quad R_Y(\theta) = \begin{pmatrix} \cos(\theta/2) & -\sin(\theta/2) \\ \sin(\theta/2) & \cos(\theta/2) \end{pmatrix}$$

$$R_Z(\theta) = \begin{pmatrix} e^{-i\theta/2} & 0 \\ 0 & e^{i\theta/2} \end{pmatrix}.$$

The Hadamard operator is the smallest gate that does something no permutation and no stochastic matrix can do. Apply it to each eigenstate and take the Born rule of the result:


$$H\mathbf{e}_0 = \frac{1}{\sqrt{2}} \begin{pmatrix} 1 \\ 1 \end{pmatrix} \longrightarrow |\mathbf{He}_0|^2 = \begin{pmatrix} 0.5 \\ 0.5 \end{pmatrix} \quad H\mathbf{e}_1 = \frac{1}{\sqrt{2}} \begin{pmatrix} 1 \\ -1 \end{pmatrix} \longrightarrow |\mathbf{He}_1|^2 = \begin{pmatrix} 0.5 \\ 0.5 \end{pmatrix}.$$

Measured after one Hadamard, the bit is uniformly distributed regardless of its input: statistically, the gate looks exactly like the COIN operator of Chapter 2. But the two output states are not the same state. They have identical Born distributions and different phases (the second carries a -1 where the first carries a $+1$), and the difference, invisible to a measurement made right now, is sitting there waiting for the next operator. Apply the Hadamard a second time:

$$HH = \frac{1}{2} \begin{pmatrix} 1 & 1 \\ 1 & -1 \end{pmatrix} \begin{pmatrix} 1 & 1 \\ 1 & -1 \end{pmatrix} = \begin{pmatrix} 1 & 0 \\ 0 & 1 \end{pmatrix} = I_2,$$

so $HH\mathbf{e}_0 = \mathbf{e}_0$, deterministically. In the top-right entry of the product, the $+1$ and the -1 met and canceled: the possibility of cancellation, promised when we introduced complex numbers, doing its first day of work.

The probability tree of Section 2.7 makes the difference visible. Draw the same picture for two Hadamards, but label each branch with its *amplitude* rather than a probability:



Every path still carries the same size, $\frac{1}{\sqrt{2}} \times \frac{1}{\sqrt{2}} = \frac{1}{2}$, and there are still two paths to each outcome. The single dashed branch is the -1 in the corner of H : a flip of the phase of Section 3.1, applied to a path rather than to an abstract complex number. Now add the paths up. The two arriving at 0 give $+\frac{1}{2} + \frac{1}{2} = 1$. The two arriving at 1 give $+\frac{1}{2} - \frac{1}{2} = 0$: they are equal in size and opposite in phase, so they annihilate, and the outcome 1 never occurs.

This is interference, and the tree shows it happening. Both paths to 1 are still there, each the same size as either path to 0, and what cancels is the pair of amplitudes, $+\frac{1}{2}$ against $-\frac{1}{2}$, meeting at the same outcome.

Compare the classical tree, where four paths of weight $\frac{1}{4}$ produced a fair coin. The path weights here have the same magnitudes throughout; the only difference is a sign on one branch, and it is enough to turn a coin flip into a certainty.

The Hadamard gate is therefore its own time-reverse, and a “coin flip” that undoes itself when repeated is no coin flip at all. This is the two-Hadamard puzzle of the Introduction,

now sitting concretely in front of us as four matrix entries. We will not analyze it here (that is the opening of Chapter 5), but the reader should pause long enough to feel the trap close: *if* the statistical reading of Chapter 2 is right, the uniform distribution after the first gate is a complete statistical description of the bit, and the second gate has nothing left to read the difference from. Either the statistical reading is wrong, or a statistical description of the *present* is not all there is to read.

3.5 A quantum program: the entangled pair

Let us write a first genuinely quantum program. It is the program from Chapter 2 that manufactured correlation, with the COIN operator replaced by its quantum counterpart:

$$\begin{array}{ccccc} \mathbf{e}_0 & \text{H} & \overline{\text{CX}} & \perp \\ \mathbf{e}_0 & & \text{CX} & \perp \end{array}$$

Computing it:

$$\text{CX}(\text{H} \otimes \text{I}_2)(\mathbf{e}_0 \otimes \mathbf{e}_0) = \begin{pmatrix} 1 & 0 & 0 & 0 \\ 0 & 1 & 0 & 0 \\ 0 & 0 & 0 & 1 \\ 0 & 0 & 1 & 0 \end{pmatrix} \frac{1}{\sqrt{2}} \begin{pmatrix} 1 \\ 0 \\ 1 \\ 0 \end{pmatrix} = \frac{1}{\sqrt{2}} \begin{pmatrix} 1 \\ 0 \\ 0 \\ 1 \end{pmatrix}.$$

This quantum state is called a *Bell state*, and it is entangled: no Kronecker product of two single-bit states can produce it. (Any product state has amplitudes $a_1b_1, a_1b_2, a_2b_1, a_2b_2$; making the middle two zero forces one of the outer two to be zero as well.) Its Born rule distribution is

$$\left| \frac{1}{\sqrt{2}}(\mathbf{e}_{00} + \mathbf{e}_{11}) \right|^2 = \begin{pmatrix} 0.5 \\ 0 \\ 0 \\ 0.5 \end{pmatrix} = \mathbf{p}_{\text{correlated}},$$

precisely the correlated distribution our probabilistic program produced: both bits uniform on their own, guaranteed to agree with each other. Measured here and now, the Bell pair and the classically correlated pair are indistinguishable. As with the Hadamard gate, the difference between them is not in the present statistics; it is in what *further operations* can do: a Bell pair fed into more gates can produce statistics the classical pair cannot, and quantifying that gap will occupy us from Chapter 4 onward. Entanglement, whatever else it turns out to be, enters the formalism as nothing more exotic than *non-factorizability*, which Chapter 2 taught us is just what correlation looks like when written as a vector.

3.6 The unistochastic bridge

How, in general, do the quantum formalism and the probabilistic formalism of Chapter 2 relate? The Born rule connects them, and the connection has a precise matrix form.

Take any unitary matrix and compute the square magnitude of every entry individually. The resulting real matrix,

$$\Gamma = |U|^2 \quad (\text{element-wise}),$$

is doubly stochastic: each of its columns and each of its rows sums to 1, because the columns and rows of a unitary matrix all have square magnitudes summing to 1. A doubly stochastic matrix obtained this way is called *unistochastic*. For example,

$$|H|^2 = \begin{pmatrix} 0.5 & 0.5 \\ 0.5 & 0.5 \end{pmatrix} = \text{COIN},$$

which makes precise the sense in which the Hadamard gate “looks like” a coin flip.

Now suppose an entire quantum program, a whole circuit, is composed into a single multi-step unitary, as associativity always permits. If the initial state is an eigenstate $\mathbf{e}_k$, the Born distribution of the output is

$$\mathbf{p}^{[t]} = \left| U^{[t \leftarrow 0]} \mathbf{e}_k \right|^2 = \left| U^{[t \leftarrow 0]} \right|^2 \mathbf{e}_k,$$

where the second equality holds because multiplying by $\mathbf{e}_k$ merely selects a column, and the square magnitude of a column is the column of square magnitudes. In words: *the statistics of any quantum circuit, start to finish, are exactly those of a single stochastic perturbation, given by the unistochastic matrix of the whole circuit’s unitary*. Run as one big step, a quantum computer is statistically indistinguishable from a probabilistic computer.

The square magnitudes of a unitary sum to 1 down every column, because the columns are unit vectors, and also along every row, because the rows are unit vectors too. So $|U|^2$ is not merely stochastic but *doubly* stochastic, and Section 2.3 showed what follows: a doubly stochastic matrix is a mixture of permutations. The Born statistics of an entire circuit are therefore the statistics of choosing one permutation of the configurations at random and applying it.

This is a real constraint on what a quantum computer can do to a distribution, and it is easy to overstate, so here is exactly what it says. A doubly stochastic matrix leaves the uniform distribution fixed, and it never lowers entropy. So no circuit, however long, can take a distribution that is spread over configurations and sharpen it into a narrow one. What the double stochasticity does *not* give us is reversibility of the statistics. The qualification Section 2.5 attached to the classical case attaches here unchanged: a permutation destroys nothing, but which permutation acted is not something we are told, so the description still degrades even though each individual run is reversible. A single Hadamard is already the counterexample to any stronger reading. It is a whole circuit, its unistochastic matrix is COIN, and it carries the definite configuration 0 to an even split: one bit of entropy where there was none. Unitarity gives reversibility at the level of ψ , and not at the level of the distribution over configurations.

The catch, and by now the reader can guess there is one, is the word *single*. The unistochastic construction does not respect composition:

$$\left| \prod_{k=1}^n U_k \right|^2 \neq \prod_{k=1}^n |U_k|^2 \quad \text{in general.}$$

The Hadamard pair is the minimal witness. Converting each gate separately and composing gives $|H|^2 |H|^2 = \text{COIN} \cdot \text{COIN} = \text{COIN}$: uniform output, the straightforward Markov-chain prediction. Composing first and then converting gives $|HH|^2 = |I_2|^2 = I_2$: deterministic

output, which is what the Born rule actually predicts. The two orders of operations disagree, and the disagreement is entirely produced by the cancellation in the top-right entry of HH: a cancellation that squaring each gate separately destroys before it can happen. Step-by-step statistics and whole-circuit statistics refuse to factor through one another. A theory whose per-step description composes into the wrong multi-step answer is exactly what a *non-Markovian* process looks like from the inside, and making that diagnosis precise is where Part II begins.

3.7 Measurement: Bayes' theorem for amplitudes

If our probability vector is not degenerate and we measure a bit, Chapter 2 told us what to do: zero out the configurations incompatible with the observation, then normalize. There is a rule of exactly the same shape for updating the quantum state on new information.

Consider the following quantum state of a two-bit memory:

$$\psi = \sqrt{\begin{pmatrix} 0.4 \\ 0.1 \\ 0.1 \\ 0.4 \end{pmatrix}} e^{i \begin{pmatrix} \pi/4 \\ \pi/4 \\ \pi/4 \\ \pi/4 \end{pmatrix}} = \begin{pmatrix} 0.4472 + 0.4472i \\ 0.2236 + 0.2236i \\ 0.2236 + 0.2236i \\ 0.4472 + 0.4472i \end{pmatrix},$$

where the root and the exponential act entry by entry: each component is $\sqrt{p_j} e^{im_j}$, its Born probability under the root and its phase in the exponent. The distribution is (0.4, 0.1, 0.1, 0.4), deliberately the same one as the Bayes example of Chapter 2, and every phase here happens to be $\pi/4$. Suppose the least significant bit is measured to have a value of 0. The configurations 01 and 11 are now ruled out, so we zero out their amplitudes:

$$\begin{pmatrix} 0.4472 + 0.4472i \\ 0 \\ 0.2236 + 0.2236i \\ 0 \end{pmatrix},$$

and then we normalize. Normalizing a quantum state is not the same operation as normalizing a probability vector: the constraint to restore is that the *square magnitudes* sum to 1, so we divide the vector by the square root of the sum of its square magnitudes,

$$\|\psi\| = \frac{\psi}{\sqrt{\sum_{j=1}^n |\psi_j|^2}} \quad \text{where } \psi \in \mathbb{C}^n.$$

The surviving square magnitudes sum to $0.4 + 0.1 = 0.5$, so we divide by $\sqrt{0.5}$:

$$\psi_{\text{after}} = \sqrt{\begin{pmatrix} 0.8 \\ 0 \\ 0.2 \\ 0 \end{pmatrix}} e^{i \begin{pmatrix} \pi/4 \\ \pi/4 \\ \pi/4 \\ \pi/4 \end{pmatrix}} = \begin{pmatrix} 0.6325 + 0.6325i \\ 0 \\ 0.3162 + 0.3162i \\ 0 \end{pmatrix},$$

whose Born distribution is $(0.8, 0, 0.2, 0)$. Compare Chapter 2: same prior distribution, a one-bit observation, and Bayes' theorem there produced the posterior by the same zeroing followed by the appropriate normalization. At the level of the probabilities, the amplitude update *is* the Bayesian update.

Set the two polar forms side by side and notice what the update touched. The phase vector is identical in both: the same $\pi/4$ in every slot, before and after. Everything that happened, happened to the distribution under the root. That is the whole of the measurement update, written in these coordinates.

Notice also what the two zeroed slots now hold. Their phases are still $\pi/4$, formally, but they are multiplied by $\sqrt{0}$, so the entries are zero regardless. Any other value could stand in those slots and the state would be exactly the same vector. A phase enters only through the product $\sqrt{p_j} e^{im_j}$, and when p_j is zero that product vanishes for every m_j whatever, so there is no way to write a non-zero phase onto a configuration of zero probability. The phase of a ruled-out configuration has not been stored somewhere for later recovery; it is not expressible, and Chapter 5 will show why that settles the matter: the interference terms that steer the dynamics each carry a factor $\sqrt{p_k p_l}$, so a configuration with no probability contributes nothing to any of them, now or at any later step.

This process of updating the quantum state after new information is acquired is often referred to as the *collapse of the wavefunction*, or the reduction of the quantum state. The reader has now performed one, and is entitled to an observation: nothing happened. No dynamical process was invoked, no mechanism, no speed of propagation, no moment at which the collapse “occurred.” A vector describing predictions was replaced by a vector describing predictions conditioned on one more fact, by the same zero-and-normalize arithmetic that Chapter 2 showed is just Bayes' theorem for measurements of bits.

Whether wavefunction collapse is a physical event or a bookkeeping update is, of course, the central battleground of quantum foundations, and we do not get to settle it by notation. If the quantum state is a physical object, then collapse is a physical process with deeply troubling properties (instantaneous, non-linear, and observer-dependent), and entire literatures attempt to domesticate it. If the quantum state is a statistical description, collapse is as mysterious as crossing options off a list. What we are entitled to say now, with the mathematics in front of us, is only this: *the update rule itself contains nothing that forces the physical reading*. The formalism of collapse is Bayes' formalism. The argument about what it means is deferred, on purpose, until we have the tools to conduct it properly: Chapter 4 builds the last of those tools, and Chapter 9 has the argument.

3.8 Projectors

The zero-and-normalize update has a standard name for its first half. Zeroing every component except the one for configuration $\mathbf{e}_j$ is what the matrix $\mathbf{e}_j \mathbf{e}_j^\dagger$ does, and a matrix of that form is called a *projector*. On one bit there are two of them:

$$\Pi_0 = \mathbf{e}_0 \mathbf{e}_0^\dagger = \begin{pmatrix} 1 & 0 \\ 0 & 0 \end{pmatrix}, \quad \Pi_1 = \mathbf{e}_1 \mathbf{e}_1^\dagger = \begin{pmatrix} 0 & 0 \\ 0 & 1 \end{pmatrix}.$$

Two properties follow by inspection. Applying a projector twice does no more than applying it once, $\Pi^2 = \Pi$, since what has already been zeroed cannot be zeroed again. And each is its

own Hermitian transpose, $\Pi^\dagger = \Pi$. Summing over every configuration keeps everything and zeroes nothing, so the projectors of a basis add to the identity: $\Pi_0 + \Pi_1 = \mathbf{I}$.

A projector need not single out one configuration. To keep a whole set of them, add their projectors together. On two bits, “the least significant bit is 0” keeps the configurations 00 and 10, so its projector is

$$\Pi = \mathbf{e}_{00}\mathbf{e}_{00}^\dagger + \mathbf{e}_{10}\mathbf{e}_{10}^\dagger = \text{diag}(1, 0, 1, 0),$$

which is the matrix form of the removal that Section 3.7 carried out by hand. Two quantities come straight off it. The probability of the outcome is the surviving weight,

$$P = |\Pi\psi|^2 = \psi^\dagger \Pi \psi,$$

the second form holding because $\Pi^2 = \Pi$ and $\Pi^\dagger = \Pi$; and the updated state is what survives, restored to unit length,

$$\psi \longmapsto \frac{\Pi\psi}{|\Pi\psi|}.$$

On the state of Section 3.7 this Π gives $P = 0.5$ and returns the same updated vector computed there.

Projectors also give observables a compact form. An observable with the two outcomes ± 1 is the difference of the projectors onto the configurations returning each, and for the ontic basis of a single bit that difference is a gate already met in Section 3.4:

$$\Pi_0 - \Pi_1 = \begin{pmatrix} 1 & 0 \\ 0 & -1 \end{pmatrix} = \mathbf{Z}.$$

Code: this chapter in Octave

```

% --- the gates ---
e0 = [1;0]; I2 = eye(2);
X = [0 1; 1 0]; Z = [1 0; 0 -1]; Y = [0 -1i; 1i 0];
H = (1/sqrt(2)) * [1 1; 1 -1];
CX = [1 0 0 0; 0 1 0 0; 0 0 0 1; 0 0 1 0];

% --- Section 3.5: the Bell-state program ---
psi = CX * kron(H, I2) * kron(e0, e0)    % -> (1/sqrt(2))[1;0;0;1]
born = abs(psi).^2                        % -> [0.5; 0; 0; 0.5]

% --- Section 3.6: the unistochastic mismatch ---
abs(H*H).^2          % composing first: identity -> deterministic
(abs(H).^2)^2        % converting first: COIN*COIN -> still COIN

% --- Section 3.7: measurement update (Bayes for amplitudes) ---
a = sqrt(0.2)+sqrt(0.2)*1i; b = sqrt(0.05)+sqrt(0.05)*1i;
psi = [a; b; b; a];          % Born distribution [0.4 0.1 0.1 0.4]
keep = [1; 0; 1; 0];         % observed: LSB = 0
q = psi .* keep;
q = q / sqrt(sum(abs(q).^2)) % renormalize by the root of the sum
abs(q).^2                   % -> [0.8; 0; 0.2; 0]

```

Chapter 4

Density Matrices

In our discussion of quantum computing, we did not discuss how to factorize the quantum state, nor how to compute Shannon information or mutual information upon it. These operations are not generally well-defined on the quantum state at all, and seeing why they fail is the fastest way to understand the object that repairs them: the density matrix. With it, the toolkit of Chapter 2 (marginals, entropy, mutual information) finally extends to quantum systems, and we will close the chapter by answering a question left open in Chapter 3: saying exactly how much more a Bell pair holds than a classically correlated pair of coins.

4.1 Why the quantum state is not enough

Two things we could do freely with probability vectors break when attempted on quantum states.

Marginalization breaks. In Chapter 2 we discarded a bit by applying a removal matrix: $\begin{pmatrix} 1 & 1 \end{pmatrix}$ in the discarded position, summing the probabilities over that bit's possibilities. Try the same trick on amplitudes. Take the two-bit quantum state

$$\psi = \frac{1}{2} \begin{pmatrix} 1 \\ 1 \\ 1 \\ -1 \end{pmatrix},$$

whose Born distribution is uniform, $(0.25, 0.25, 0.25, 0.25)$; the correct marginal for the most significant bit is therefore $(0.5, 0.5)$. But applying the removal matrix to the amplitudes gives

$$(\mathbf{I}_2 \otimes \begin{pmatrix} 1 & 1 \end{pmatrix}) \psi = \frac{1}{2} \begin{pmatrix} 2 \\ 0 \end{pmatrix} = \begin{pmatrix} 1 \\ 0 \end{pmatrix},$$

whose Born distribution is $(1, 0)$: wrong. The culprit is the one new capability complex numbers gave us. Summing over a discarded bit is legitimate for probabilities, which can only accumulate; amplitudes of opposite phase *cancel* in the sum, and the cancellation corrupts the statistics of the bits we kept. Marginalization must happen at the level of probabilities, after the squaring has already taken place. The quantum state performs its squaring too late to allow it.

Classical uncertainty breaks. A quantum state has no room to express ordinary, classical ignorance. Suppose a machine prepares a bit in the state $\mathbf{H}\mathbf{e}_0$ on heads and $\mathbf{H}\mathbf{e}_1$ on tails, and we do not know which came up. Our description should be a 50/50 statistical mixture of *quantum states*. No single quantum state expresses this: averaging the two vectors is meaningless, since adding amplitudes lets their phases cancel, which is not what an ignorant fifty-fifty is meant to say, and picking either one asserts knowledge we do not have. In Chapter 2, statistical mixing was the whole content of the formalism; in Chapter 3 it has no home. A quantum description of a system should be able to carry both kinds of statistical structure at once: the Born-rule structure internal to a quantum state, and garden-variety ignorance about which quantum state it is. The quantum state alone carries only the first.

4.2 The density matrix

Both problems have the same fix: instead of the vector ψ , work with the matrix formed from ψ and its own Hermitian transpose,

$$\rho = \psi\psi^\dagger.$$

This is called a *density matrix*. Where ψ is a column of n amplitudes, ρ is an $n \times n$ grid: entry (j, k) holds the product $\psi_j\overline{\psi_k}$.

On the diagonal, $j = k$, the entries are $\psi_j\overline{\psi_j} = |\psi_j|^2$: the Born-rule probabilities, sitting in plain view with the squaring already performed. The probability of a configuration is read off the diagonal,

$$P(\mathbf{e}_k) = \mathbf{e}_k^\top \text{diag}(\rho),$$

and since the diagonal is an ordinary probability vector, everything from Chapter 2 can be applied to it, including proper marginalization. Off the diagonal, the entries $\psi_j\overline{\psi_k}$ record the relative phases between pairs of amplitudes: the information that squaring destroys, kept in a form we will learn to trace and transform. The density matrix is the bookkeeping device that carries the probabilities and the phase relations side by side without letting either corrupt the other.

Two examples we will reuse all chapter. The Bell state of Chapter 3 gives

$$\psi = \frac{1}{\sqrt{2}} \begin{pmatrix} 1 \\ 0 \\ 0 \\ 1 \end{pmatrix} \longrightarrow \rho_{\text{Bell}} = \psi\psi^\dagger = \frac{1}{2} \begin{pmatrix} 1 & 0 & 0 & 1 \\ 0 & 0 & 0 & 0 \\ 0 & 0 & 0 & 0 \\ 1 & 0 & 0 & 1 \end{pmatrix}.$$

And the classically correlated pair from Chapter 2 (50% both-zero, 50% both-one) is a statistical mixture of two definite configurations, which the next section will teach us to write as

$$\rho_{\text{classical}} = 0.5 \mathbf{e}_{00}\mathbf{e}_{00}^\dagger + 0.5 \mathbf{e}_{11}\mathbf{e}_{11}^\dagger = \frac{1}{2} \begin{pmatrix} 1 & 0 & 0 & 0 \\ 0 & 0 & 0 & 0 \\ 0 & 0 & 0 & 0 \\ 0 & 0 & 0 & 1 \end{pmatrix}.$$

Set them side by side. The diagonals are identical, since both systems answer every immediate measurement with the correlated distribution $(0.5, 0, 0, 0.5)$, and the *only* difference is the pair of $\frac{1}{2}$ s in the corners of ρ_{Bell} . Chapter 3 promised that the difference between a Bell pair and a

correlated pair of coins is invisible now and decisive later; here is that difference, localized to two matrix entries.

4.3 Mixtures, and what a density matrix really describes

The equation $\rho = \psi\psi^\dagger$ is not the general definition. A density matrix can be formed by combining any number of quantum states, each weighted by a classical probability of being the actual state:

$$\rho = p_1\psi_1\psi_1^\dagger + p_2\psi_2\psi_2^\dagger + p_3\psi_3\psi_3^\dagger + \cdots + p_n\psi_n\psi_n^\dagger.$$

This solves the coin-flipping preparation machine: its output is $0.5(\mathbf{He}_0)(\mathbf{He}_0)^\dagger + 0.5(\mathbf{He}_1)(\mathbf{He}_1)^\dagger$. Both the quantum states and their classical weights contribute to the total probability of each outcome, and the diagonal automatically delivers the correctly averaged statistics.

Evolution carries over from Chapter 3 by direct substitution. Since ψ evolves to $U\psi$, the matrix $\psi\psi^\dagger$ evolves to $(U\psi)(U\psi)^\dagger$, giving

$$\rho_{\text{after}} = U\rho_{\text{before}}U^\dagger,$$

and it is also possible to evolve a density matrix by a probability distribution over unitary matrices, giving a stochastic mixture of quantum gates, in the spirit of Chapter 2:

$$\rho_{\text{after}} = p_1U_1\rho_{\text{before}}U_1^\dagger + p_2U_2\rho_{\text{before}}U_2^\dagger + \cdots$$

Now a fact with more philosophy packed into it than is obvious at first glance. Compute the density matrix for the coin-flipped mixture of $\mathbf{He}_0$ and $\mathbf{He}_1$:

$$0.5 \cdot \frac{1}{2} \begin{pmatrix} 1 & 1 \\ 1 & 1 \end{pmatrix} + 0.5 \cdot \frac{1}{2} \begin{pmatrix} 1 & -1 \\ -1 & 1 \end{pmatrix} = \frac{1}{2} \begin{pmatrix} 1 & 0 \\ 0 & 1 \end{pmatrix},$$

and compare it with the 50/50 mixture of the eigenstates $\mathbf{e}_0$ and $\mathbf{e}_1$:

$$0.5 \begin{pmatrix} 1 & 0 \\ 0 & 0 \end{pmatrix} + 0.5 \begin{pmatrix} 0 & 0 \\ 0 & 1 \end{pmatrix} = \frac{1}{2} \begin{pmatrix} 1 & 0 \\ 0 & 1 \end{pmatrix}.$$

Two entirely different preparation procedures (different ingredient states, different stories about what happened) yield the *same* density matrix. And it is a theorem of the formalism (we will not prove it, but we will lean on it) that all predictions of quantum theory depend on a system's preparation only through ρ : every probability, for every measurement, after every further evolution. Ensembles with the same density matrix are empirically identical, forever, no matter what is done to them.

This can be read as a deflation or as a definition. Deflation: the density matrix does not record everything about how a system was prepared, so it is not a complete history. Definition: the density matrix records *exactly* what the preparation makes a difference to, because it is the sufficient statistic of the preparation for all future predictions, in precisely the Chapter 2 sense of the term. Detail that two preparations share, the formalism keeps; detail that cannot matter to any future measurement, the formalism silently discards. We will exploit this property repeatedly: here it is a convenience, and in Chapter 9, where the question is whether a “collapsed” and an “uncollapsed” laboratory are distinguishable, it carries the whole argument.

4.4 Purity

Density matrices have a property known as *purity*. The matrix, again, represents a classical probability distribution over quantum states. If that distribution is degenerate, meaning the matrix is built from a single quantum state, $\rho = \psi\psi^\dagger$, the density matrix is said to be *pure*. Purity can be tested by a simple computation, the trace of the square of the matrix, where the trace $\text{Tr}(M)$ of a matrix M is the sum of its diagonal entries:

$$\gamma = \text{Tr}(\rho^2) \quad \text{where } \frac{1}{n} \leq \gamma \leq 1, \rho \in \mathbb{C}^{n \times n}.$$

The purity is 1 exactly for pure states. Anything less is referred to as a *mixed* state, and the floor $\gamma = 1/n$ is reached by one matrix only: $\rho = \mathbf{I}/n$, the uniform classical mixture we computed twice in the last section, carrying no phase information whatsoever: a distribution with nothing quantum-mechanical left in it at all. Our two running examples separate immediately:

$$\text{Tr}(\rho_{\text{Bell}}^2) = 1 \quad \text{Tr}(\rho_{\text{classical}}^2) = 0.25 + 0.25 = 0.5.$$

The Bell pair, for all its internal statistics, is a *single, exactly known quantum state*; the classical pair is genuinely a mixture. Purity is the diagnostic that tells them apart from the inside.

4.5 The partial trace

We can now do properly what Section 4.1 showed cannot be done on amplitudes: discard a subsystem. The operation on density matrices analogous to marginalization is known as the *partial trace*. Similar to marginalization, it is constructed by combining, with the Kronecker product, identity matrices in the positions of the bits we wish to preserve and bit vectors in the positions of the bits we wish to remove. Where marginalization summed once over the removed bit's possibilities, though, the partial trace sandwiches the density matrix from both sides and sums:

$$\text{Tr}_{\text{out}}(\rho) = \sum_{k=0}^{n-1} (\mathcal{R}_1 \otimes \mathcal{R}_2 \otimes \cdots \otimes \mathcal{R}_n) \rho (\mathcal{R}_1 \otimes \mathcal{R}_2 \otimes \cdots \otimes \mathcal{R}_n)^\dagger \quad \text{where } \mathcal{R}_j \in \{\mathbf{I}_2, \mathbf{e}_k^\top\},$$

the sum running over the configurations k of the removed bits. It is also common to write a partial trace by naming the subsystem removed:

$$\rho^A = \text{Tr}_B \rho^{AB},$$

read as tracing out subsystem B from the joint system AB , leaving a *reduced density matrix* for subsystem A alone. On the diagonal, the partial trace is exactly Chapter 2 marginalization of the probabilities; off the diagonal, it discards the phase relations that involved the removed bits, keeping those purely internal to the kept ones. The sandwiching by $\mathbf{e}_k$ and $\mathbf{e}_k^\top$, rather than a single sum over an all-ones row vector, is precisely what prevents amplitudes belonging to *different* configurations of the discarded system from meeting and canceling, which is the failure that broke Section 4.1.

Apply it to the Bell state, tracing out the second bit. The two terms of the sum pick out the sub-blocks of ρ_{Bell} where the removed bit's input and output configurations agree, and yield

$$\rho^A = \text{Tr}_B \rho_{\text{Bell}} = \frac{1}{2} \begin{pmatrix} 1 & 0 \\ 0 & 0 \end{pmatrix} + \frac{1}{2} \begin{pmatrix} 0 & 0 \\ 0 & 1 \end{pmatrix} = \frac{1}{2} \begin{pmatrix} 1 & 0 \\ 0 & 1 \end{pmatrix}.$$

Something remarkable has happened. Performing a partial trace over a density matrix with *maximal* purity has yielded a reduced density matrix with *minimal* purity: we began with $\text{Tr}(\rho_{\text{Bell}}^2) = 1$, the joint state exactly known, and ended with $\text{Tr}((\rho^A)^2) = 0.5$, the reduced state a maximally mixed, phase-free, perfectly classical coin. The corners of ρ_{Bell} (the phase relations) were attached to *both* bits jointly, and tracing out either bit takes them away entirely.

4.6 Kraus operators, quantum form

Chapter 2 wrote the reduced dynamics of a bit coupled to a discarded bit as a mixture of component matrices, $\Gamma_A = \sum_b P(b) K_b$, and said that the quantum form of the same construction would follow. Here it is.

Let a joint unitary U act on our system A together with an auxiliary system B (an *environment*) prepared in the quantum state ψ_B . The reduced evolution of A alone is obtained by evolving the joint state and tracing out B , and it can always be organized as follows: define one component operator per configuration k of the environment's *output*,

$$K_k = (\text{I} \otimes \mathbf{e}_k^T) U (\text{I} \otimes \psi_B),$$

read right to left: feed the environment its initial state, apply the joint unitary, and select the branch in which the environment ends in configuration k . Then

$$\rho_{\text{after}}^A = \sum_k K_k \rho_{\text{before}}^A K_k^\dagger, \quad \sum_k K_k^\dagger K_k = \text{I},$$

the second equation being the guarantee that probability is conserved. The K_k are called *Kraus operators*, and the parallel with Chapter 2 is exact in form: there, the classical component matrices were indexed by the discarded bit and mixed with its probabilities; here, the component operators are indexed by the environment's final configuration and sandwich the density matrix one branch at a time. The difference is where the construction lives: the classical K_b acted on probabilities, while the quantum K_k act on amplitudes, each branch dragging its phase information along with it.

Run the Chapter 2 example again, now at the level of amplitudes: let $U = \text{XC}$ be the controlled-NOT with our bit A as target and the environment bit B as control, and prepare the environment in the uniform superposition $\psi_B = \text{H}\mathbf{e}_0$. The two Kraus operators come out as

$$K_0 = \frac{1}{\sqrt{2}} \text{I}_2, \quad K_1 = \frac{1}{\sqrt{2}} \text{X},$$

so the reduced dynamics of A is $\rho \mapsto \frac{1}{2}\rho + \frac{1}{2}\text{X}\rho\text{X}$: half “leave it alone,” half “flip it,” which is the COIN operator's quantum silhouette, exactly as in Chapter 2. But watch it act on an eigenstate: $\mathbf{e}_0\mathbf{e}_0^\dagger$ becomes $\frac{1}{2}\mathbf{e}_0\mathbf{e}_0^\dagger + \frac{1}{2}\mathbf{e}_1\mathbf{e}_1^\dagger = \text{I}/2$, purity $1 \rightarrow \frac{1}{2}$. A single deterministic, reversible two-bit gate has *decohered* bit A. Bit A became entangled with bit B, and bit B

was then traced away. Kraus operators are decoherence made algorithmic, and they are the standard language in which the noise of real quantum hardware is described.

A reader who has met decoherence before may find this arrangement back to front. The familiar account starts the *system* in a superposition, couples it to an environment that began definite, and traces the environment away, and what the trace destroys is a coherence the system brought in with it. The arrangement above is the other way round: the superposition is in bit B, the bit about to be discarded, and bit A starts in the definite configuration $\mathbf{e}_0$. Two things depend on that choice.

The first is the parallel with Chapter 2. There the component matrices K_b were indexed by the discarded bit, and the discarded bit was the one holding the probabilities: COIN came out of XC because the control was a coin and the target was the bit being watched. Keeping the parallel means putting the superposition where Chapter 2 put the randomness, which is the environment. Exchanging the roles breaks it, and the break is easy to see. Give A the superposition $\mathbf{H}\mathbf{e}_0$, start B in $\mathbf{e}_0$, and let A be the control and B the target, so that B records which configuration A is in. The component operators are then the projectors $K_0 = \mathbf{e}_0\mathbf{e}_0^\dagger$ and $K_1 = \mathbf{e}_1\mathbf{e}_1^\dagger$, and the reduced dynamics they generate is $\rho \mapsto \mathbf{e}_0\mathbf{e}_0^\dagger\rho\mathbf{e}_0\mathbf{e}_0^\dagger + \mathbf{e}_1\mathbf{e}_1^\dagger\rho\mathbf{e}_1\mathbf{e}_1^\dagger$, which zeroes the off-diagonal entries and leaves the diagonal untouched. No half chance of X appears anywhere in it, so nothing resembling COIN appears either. Applied to bit A in $\mathbf{e}_0$ that operation does nothing whatever, and the purity stays at 1.

The second concerns where the purity drop comes from. In the arrangement as written, bit A enters at purity 1, a definite configuration with no superposition in it, and leaves at $\frac{1}{2}$. The mixedness therefore cannot be a spread that bit A already carried, since it carried none; it is produced by the coupling to B and by the discarding of B, and by nothing else. Start A in a superposition instead and the same final I/2 admits a second reading, that a spread present from the beginning has merely been laid bare, and that reading is awkward to rule out. Keeping A definite at the start rules it out at no cost, which leaves the plain account: a bit that was somewhere definite throughout, correlated with a bit nobody tracked, and described afterwards by a matrix that records the correlation rather than anything about bit A on its own.

4.7 Von Neumann entropy

The quantum equivalent to Shannon entropy is referred to as *von Neumann entropy*, computed as

$$S(\rho) = -\text{Tr}(\rho \log_2 \rho).$$

Keep in mind that this is a *matrix* logarithm, not an element-wise one, and the practical route to it runs through a fact we will state without proof: every density matrix can be rotated by some unitary into a diagonal form, and its diagonal in that form is known as its *spectrum*: a list of non-negative numbers summing to 1, an ordinary probability distribution. The von Neumann entropy of ρ is just the Shannon entropy of its spectrum. Intuitively, the spectrum is the mixture the density matrix “really” is: the most economical classical uncertainty about quantum states consistent with ρ . And $S(\rho)$ measures that irreducible uncertainty in bits. For a pure state the spectrum is degenerate $(1, 0, 0, \dots)$, so $S = 0$: a pure state, whatever its internal Born statistics, is a completely known quantum state, and von Neumann entropy does

not count Born-rule randomness. For the maximally mixed bit $I/2$ the spectrum is uniform, so $S = 1$.

Our running examples, assembled:

$$S(\rho_{\text{Bell}}) = 0 \quad S(\rho_{\text{classical}}) = 1 \quad S(\text{Tr}_B \rho_{\text{Bell}}) = S(I/2) = 1.$$

Note the inversion in the first and third entries: the Bell pair has *less* entropy than either of its own halves. Classically this is impossible, because Shannon entropy of a joint system is never smaller than that of a part. Its possibility here is the entire subject of the next section.

4.8 Quantum mutual information

It is then possible to define mutual information for quantum systems in exactly the Chapter 2 shape, with von Neumann entropy in place of Shannon entropy:

$$I(A : B) = S(\rho^A) + S(\rho^B) - S(\rho^{AB}).$$

We will call this *quantum mutual information* (QMI). When contrast is needed, we will call the Chapter 2 quantity, computed from the diagonal (the measured statistics), *classical mutual information* (CMI).

We can now make good on the promise from the close of Chapter 3. For the classically correlated pair,

$$I(A : B) = 1 + 1 - 1 = 1 :$$

one bit of correlation, exactly as Chapter 2 computed, and exactly what the story warrants, since one coin flip was recorded onto the second bit. For the Bell pair,

$$I(A : B) = 1 + 1 - 0 = 2.$$

Two bits of mutual information between two bits that agree on a coin flip. The first bit of it is the same bit the classical pair has: measure either half and the other half's value is determined. The second bit is new, stored in nothing but the corners of ρ_{Bell} (the joint phase relations), and no measurement of the bits' *values* can see it. For a single pair of bits, CMI has a maximum value of 1, but QMI has a maximum value of 2. This is not a bookkeeping curiosity: it is a real, usable resource, and Chapter 7 will use it. The superdense coding protocol transmits two full classical bits by sending one quantum bit, retrieving exactly the extra bit of QMI stored in a shared Bell pair.

4.9 When classical mutual information misleads

Classical mutual information can be actively deceptive when applied to quantum systems. Consider the quantum state

$$\psi = \sqrt{\begin{pmatrix} 0.25 \\ 0.25 \\ 0.25 \\ 0.25 \end{pmatrix}} e^{i \begin{pmatrix} 0 \\ -\pi/2 \\ -\pi/2 \\ 0 \end{pmatrix}} = \begin{pmatrix} 0.5 \\ -0.5i \\ -0.5i \\ 0.5 \end{pmatrix},$$

written here in polar form: the first factor is the element-wise square root of the Born distribution, the second the element-wise phases. This state's Born distribution is uniform and uncorrelated, so the CMI between the two bits is 0: measure both, and the outcomes are two independent fair coins, teaching each other nothing. One might conclude the bits have nothing to do with one another.

The conclusion would be wrong. Apply a Pauli-X rotation to the second bit (a local operation, touching nothing but that bit) before measuring:

$$(\mathbf{I}_2 \otimes \mathbf{R}_X(-\pi/2))\psi = \begin{pmatrix} 1/\sqrt{2} \\ 0 \\ 0 \\ 1/\sqrt{2} \end{pmatrix},$$

the Bell state, whose measured statistics are perfectly correlated. The correlation was there all along. It was stored entirely in the phases, where value-measurements cannot see it, and one local rotation moves it back into the values. Had we rotated before measuring, each bit's outcome would have told us the other's with certainty.

The QMI records it. This ψ is pure, so $S(\rho^{AB}) = 0$, and its reduced density matrices both work out to $\mathbf{I}/2$, so

$$I(A : B) = 1 + 1 - 0 = 2 :$$

maximal, while the CMI is minimal. Thus the CMI can be misleading for quantum systems, as bits of correlation can be implicitly recoverable through a set of operations even when the present measurement statistics show none. QMI is the complete count: it registers correlation wherever it is stored (in the values, in the phases, or split between them), and it is invariant under local rotations like the one above, which merely slide correlation from the values to the phases and back.

For our argument, this section is more than a warning label. Chapter 3 left us with a suspicion: that a statistical description of the *present values* is not all there is to a quantum system. The QMI/CMI gap makes the suspicion quantitative, because there is information in quantum systems, real enough to be extracted by local operations, that the statistics of present values do not carry. Where is it kept? The phases hold it in the formalism; but what are the phases *of*? Part II answers: they are a compressed record of where the system's statistics have *been*. The information that present statistics cannot carry, a history can.

4.10 Decoherence

The Bell case of Section 4.5 is not special. If a bit becomes perfectly entangled with another bit, or with any other collection of bits, and we perform a partial trace to remove what it is entangled with, the reduced density matrix we obtain is always of minimal purity. A bit which is perfectly entangled with something else, considered in complete isolation from what it is entangled with, behaves exactly as if it were classical: its off-diagonal terms vanish, its statistics are an ordinary probability distribution, and no operation on it alone can reveal any phase. This phenomenon is known as *decoherence*.

What is left can be written in the form Section 4.3 gave for a mixture. A decohered bit has $\rho = p \mathbf{e}_0 \mathbf{e}_0^\dagger + q \mathbf{e}_1 \mathbf{e}_1^\dagger$, a weighted combination of its own eigenstates, each of which is a

definite configuration. That is a classical probability distribution over definite values, written as a density matrix, and it is the form a probability vector takes in this formalism.

The reader has seen a simpler version of this already. In Chapter 2, the paradigmatically random COIN operator turned out to be the shadow of the deterministic two-bit XC gate, watched from one bit while ignoring the other. The quantum statement is the stronger one: *classicality itself* is the shadow of a pure entangled state, watched from a subsystem while ignoring what it is entangled with.

Nothing about the joint system became mixed, random, or classical. The joint state remains pure, its purity 1, its evolution unitary and reversible. Mixedness appears only in the description of a subsystem.

The ignored remainder is usually large. An environment of stray photons and air molecules interacts with everything and records as it goes, and the everyday world of definite-looking, phase-free statistics is the result. This is why quantum computers are built in shielded refrigerators: their entire task is to keep phase relations *inside* the computation, and every bit of unintended entanglement with the outside world traces out, from the computer's point of view, as pure classical noise.

Perfect entanglement is the endpoint of a continuum rather than a case on its own. Between a product state and a maximally entangled one lies a whole range of partial entanglement, and the purity of the subsystem tracks it exactly. Take the entangled-pair program of Section 3.5 and put a rotation through θ where its Hadamard was:


$$\begin{array}{ccccc} \mathbf{e}_0 & R_Y(\theta) & \overline{CX} & \perp & \\ \mathbf{e}_0 & & CX & \perp & \end{array} \quad \psi = \cos(\theta/2)\mathbf{e}_{00} + \sin(\theta/2)\mathbf{e}_{11}.$$

At $\theta = 0$ the two bits are unentangled, and at a right angle this is the Bell state. Tracing out the second leaves the first in

$$\rho^A = \begin{pmatrix} \cos^2(\theta/2) & 0 \\ 0 & \sin^2(\theta/2) \end{pmatrix}, \quad P = 1 - \frac{1}{2} \sin^2 \theta,$$

The joint state is pure, so $S(\rho^{AB}) = 0$ and the two marginals share an entropy, which makes the quantum mutual information of Section 4.8 twice that entropy. Increasing θ moves the two quantities together:

θ	0	$\pi/6$	$\pi/4$	$\pi/3$	$\pi/2$
$I(A : B)$	0	0.709	1.202	1.623	2
P	1	0.875	0.75	0.625	0.5



The two ends are the cases already computed. A product state carries no QMI and leaves the subsystem pure; the Bell state carries the maximum of 2 and leaves it at $P = 1/2$, the minimum available to a single bit. Every value between is reachable, and the more a bit shares with what it is entangled with, the less pure its own description becomes.

So decoherence works as a dial rather than a switch, and its setting is how much the subsystem has come to share with whatever we are ignoring. A bit slightly entangled with an environment is slightly decohered: its off-diagonal terms are reduced rather than removed, and its phase is partly, not wholly, out of reach of operations on it alone. Everyday decoherence is this dial turned all the way up. Each collision correlates a system a little further with its environment, and no one can track the quintillions of particles a chaotic environment contains, so any subsystem carved out of the universe tends to look decohered unless it has been deliberately isolated.

Code: this chapter in Octave

```

% --- the gates ---
e0 = [1;0]; e1 = [0;1]; I2 = eye(2);
H = (1/sqrt(2))*[1 1; 1 -1];
CX = [1 0 0 0; 0 1 0 0; 0 0 0 1; 0 0 1 0];

% --- Sections 4.2-4.5: density matrices, purity, partial trace ---
psi = CX * kron(H, I2) * kron(e0, e0);
rhoB = psi * psi'; % Bell density matrix
rhoC = 0.5 * kron(e0,e0)*kron(e0,e0)' ...
      + 0.5 * kron(e1,e1)*kron(e1,e1)'; % classical correlated pair

purity = @(r) real(trace(r^2));
purity(rhoB), purity(rhoC) % 1 and 0.5

trB = @(r) kron(I2,e0') * r * kron(I2,e0) ...
      + kron(I2,e1') * r * kron(I2,e1);
rhoA = trB(rhoB) % -> I2/2, purity 0.5

% --- Section 4.6: Kraus operators (XC gate, environment in H|0>) ---
XC = [1 0 0 0; 0 0 0 1; 0 0 1 0; 0 1 0 0]; % B (LSB) controls, A target
psiB = H*e0;
K0 = kron(I2, e0') * XC * kron(I2, psiB) % -> I2/sqrt(2)
K1 = kron(I2, e1') * XC * kron(I2, psiB) % -> X/sqrt(2)
K0'*K0 + K1'*K1 % completeness: identity
K0*(e0*e0')*K0' + K1*(e0*e0')*K1' % -> I2/2: decohered

% --- Sections 4.7-4.9: entropy, QMI vs CMI ---
S = @(r) (@(v) -sum(v(v>1e-12).*log2(v(v>1e-12))))(real(eig(r)));
trA = @(r) kron(e0',I2) * r * kron(e0,I2) ...
      + kron(e1',I2) * r * kron(e1,I2);
QMI = @(r) S(trB(r)) + S(trA(r)) - S(r);
QMI(rhoB), QMI(rhoC) % 2 and 1

psi = [0.5; -0.5i; -0.5i; 0.5];
abs(psi).^2 % uniform: CMI = 0
QMI(psi*psi') % 2
Rx = @(t) [cos(t/2) -1i*sin(t/2); -1i*sin(t/2) cos(t/2)];
kron(I2, Rx(-pi/2)) * psi % -> Bell state

% --- Section 4.10: partial entanglement, QMI against purity ---
RY = @(t) [cos(t/2) -sin(t/2); sin(t/2) cos(t/2)];
for th = [0 pi/6 pi/4 pi/3 pi/2] % rotate, then CX
    v = CX * kron(RY(th)*e0, e0);
    [th, QMI(v*v'), purity(trB(v*v')), 1 - sin(th)^2/2]
end % QMI 0 -> 2 as purity falls 1 -> 0.5

```

Part II

Quantum Information as Statistics

Chapter 5

Non-Markovian Stochastic Dynamics

Part I built three computers and a toolkit. Part II uses them. Our subject from here to the end of Part II is a single question, put off since the Introduction: *what kind of statistical process is a quantum circuit?* We pose the question in its sharpest form, survey the standard answer and its problems, and then do the central technical work: splitting the quantum state into statistics and phase, and showing that the phase can be expanded away into the system's statistical history. What remains, when the expansion is done, is a statistical theory whose laws refuse to be Markovian.

5.1 The same gate, twice

Consider the quantum circuit below: a computer with a single bit of memory, one Hadamard operator applied to a bit with an initial value of 0, then a measurement.

$$\mathbf{e}_0 \xrightarrow{H} \perp \quad |\mathbf{H}\mathbf{e}_0|^2 = \left| \frac{1}{\sqrt{2}} \begin{pmatrix} 1 & 1 \\ 1 & -1 \end{pmatrix} \begin{pmatrix} 1 \\ 0 \end{pmatrix} \right|^2 = \begin{pmatrix} 0.5 \\ 0.5 \end{pmatrix}.$$

The result is a uniform distribution over the bit's value. Run the same circuit with the bit initialized to 1 instead, and the outcome is also a uniform distribution. So, regardless of whether the input is 0 or 1, the output is uniform. Knowing that tells us exactly which stochastic matrix to assign to the gate's behavior. Both columns of the statistical truth table are the uniform distribution:

$$\Gamma_{\text{coin}} = \begin{pmatrix} 0.5 & 0.5 \\ 0.5 & 0.5 \end{pmatrix}.$$

The Hadamard gate, statistically, is the COIN operator. Chapter 3 said as much: $|\mathbf{H}|^2 = \text{COIN}$.

Now the third circuit: the Hadamard applied twice in a row to a bit with an initial value of 0.

$$\mathbf{e}_0 \xrightarrow{H \ H} \perp \quad |\mathbf{H} \mathbf{H} \mathbf{e}_0|^2 = \begin{pmatrix} 1 \\ 0 \end{pmatrix}.$$

The output is a degenerate distribution: a deterministic 0. We know the output of the first Hadamard gate was a uniform distribution over 0 and 1, and we know the outcome after the

second is always 0, which means the second gate maps *both* of its possible inputs to 0. In both possible cases, its output is 0. That, too, fixes a stochastic matrix, and it is not the one we assigned a moment ago:

$$\Gamma_{\text{zero}} = \begin{pmatrix} 1 & 1 \\ 0 & 0 \end{pmatrix}.$$

In probabilistic computing, a logic gate imposes one and the same stochastic perturbation wherever it appears: that is what the Chapter 2 formalism *means* by a gate. The Hadamard gate does not fit that description. The same physical operation, applied to the same physical system, exhibits one statistical behavior in the first slot of the circuit and a different statistical behavior in the second. What is particularly interesting is what the choice tracks: if the bit arrives in a degenerate distribution, the gate takes it to a uniform one, and if the bit arrives in the uniform distribution manufactured by a previous Hadamard, the gate takes it back to a degenerate one,

$$\Gamma_{\text{H}} = \begin{cases} \Gamma_{\text{coin}} & \text{if } \mathbf{p} = \begin{pmatrix} 1 \\ 0 \end{pmatrix} \text{ or } \begin{pmatrix} 0 \\ 1 \end{pmatrix} \\ \Gamma_{\text{zero}} & \text{if } \mathbf{p} = \begin{pmatrix} 0.5 \\ 0.5 \end{pmatrix} \text{ arrived via H,} \end{cases}$$

as if the gate were choosing its behavior based upon the probability distribution of the incoming bit.

All of that was for a bit that began at 0. Now the fourth circuit: the same two Hadamards, on a bit with an initial value of 1.

$$\mathbf{e}_1 \quad \text{H} \quad \text{H} \quad \perp \quad |\text{H H } \mathbf{e}_1|^2 = \begin{pmatrix} 0 \\ 1 \end{pmatrix}.$$

The first Hadamard gate leaves a uniform distribution here too, and the outcome after the second is always 1, so this time the second gate maps both of its possible inputs to 1:

$$\Gamma_{\text{one}} = \begin{pmatrix} 0 & 0 \\ 1 & 1 \end{pmatrix}.$$

The gate's behavior tracks the distribution its input arrives in here just as it did before, a degenerate distribution going to a uniform one and a uniform one going back to a degenerate one, and the case statement has the same shape with its second row changed:

$$\Gamma_{\text{H}} = \begin{cases} \Gamma_{\text{coin}} & \text{if } \mathbf{p} = \begin{pmatrix} 1 \\ 0 \end{pmatrix} \text{ or } \begin{pmatrix} 0 \\ 1 \end{pmatrix} \\ \Gamma_{\text{one}} & \text{if } \mathbf{p} = \begin{pmatrix} 0.5 \\ 0.5 \end{pmatrix} \text{ arrived via H.} \end{cases}$$

5.2 Why this is a genuine puzzle

Conceptually, there is a serious difficulty here, and rushing past it would lose the whole point.

A statistical distribution describes the behavior of a system over an ensemble. It is not a physical property of the individual system. If a flipped coin has a 50% chance of landing

heads and a 50% chance of landing tails, then in the real world, in each individual case, it lands on one or the other. The Hadamard gate, in an individual run, receives a single bit, and that bit has a single value: 0 or 1. It is not possible to distinguish, from the bit value alone, whether that bit belongs to a degenerate ensemble or a uniform one. A 0 drawn from a fair coin is the same 0 as a 0 that was certain all along. One sample carries no information about its ensemble. Yet somehow, in every individual run, the second Hadamard gate behaves in the way appropriate to the ensemble its input belongs to.

This is the magician of the Introduction, no longer an analogy but a matrix identity. And notice precisely what makes it strange: nothing about the *outcome statistics* is strange: a uniform distribution, then a degenerate one; no probability is negative, nothing oscillates, no cat is involved. What is strange is that the *dynamics* (which stochastic perturbation happens) depends on information that the present physical state of the bit, on any ordinary reading, does not contain.

There is a name in probability theory for the assumption being violated. A stochastic process is *Markovian* when the statistics of its next step depend on the system's history only through the system's present state. In matrix terms: each instruction gets one stochastic matrix, fixed by the definition of the instruction itself, and evolution is iterated matrix multiplication: $\Gamma\Gamma'\Gamma''\dots$ applied to the present distribution, step after step. (The matrices are allowed to differ from step to step. A scheduled sequence of different gates is still perfectly Markovian; what matters is that no step's matrix depends on the *history* of the process.) Every process in Chapter 2 was Markovian by construction. The Hadamard pair is our proof, by exhibition, that quantum circuits are not: no assignment of one stochastic matrix per instruction can reproduce both behaviors of the same instruction, because the second gate's behavior depends on *how the present distribution came about*. A non-Markovian stochastic process may follow more complicated rules, such as the one given below, which the reader can check reproduces both Hadamard behaviors:

$$\mathbf{p}^{[t+1]} = \left| U \sqrt{\mathbf{p}^{[t]}} \right|^2,$$

with the square root taken element-wise. Feed it (1, 0) and it returns (0.5, 0.5); feed it (0.5, 0.5) and it returns (1, 0). This little rule is not Markovian, because it is not a stochastic matrix acting on $\mathbf{p}$ (the square root makes it non-linear in the probabilities), and it is not yet general quantum mechanics either (it silently assumes all phases are zero, an assumption whose repair occupies Section 5.6). But it already demonstrates the shape of the resolution: the statistics can be lawful, exact, and computable without the per-step behavior being expressible as a fixed matrix on the present distribution.

Notice where the phases are in it, because that is why it is here. The state it evolves is a probability distribution and carries no phase at all, so whatever the second Hadamard responds to, it is not a phase the bit brought with it. The interference survives regardless: feed the rule a uniform distribution and a degenerate one comes back. What explains it is the only thing left that could, which is U . For the Hadamard pair that is the single minus sign in the corner of H , a phase of π .

Section 5.7 will show that the phases can be taken out of the state in general, and not only in this special case. The gain is real: it retires the picture of a wave with no location in the world flowing through the circuit, and leaves bits with definite values and the definitions

of the gates. It does not explain the magician’s trick. The gates still have phases, and the dependence on the ensemble statistics is still there because of them. The ontology gets simpler and the puzzle stays where it was.

The shape of the answer is already implied, though. A statistical law that fails to be Markovian is ordinarily put down to a hidden memory the description failed to account for, and the word *memory* carries the suggestion that what matters is the past being taken into account. Section 5.5 will argue that no memory is needed for that, a dependence on the past being enough on its own, and Section 8.8 will argue that a dependence of that kind implies a holistic character in the laws.

5.3 The hidden-memory escape


Traditionally, when a stochastic process is found to be non-Markovian, it is assumed that there exists a hidden memory: extra unobserved state which, if taken into account, restores the Markov property. The system looked non-Markovian only because we were watching a subsystem of it. Two earlier results were versions of this move seen from the other side. The XC gate of Chapter 2 acts deterministically on two bits and looks like a fair coin on one of them, and the decoherence of Chapter 4 leaves a joint state pure while the subsystem we keep looks classical. In both, the oddity in the subsystem’s behavior is an artifact of the subsystem, and restoring the discarded part restores the ordinary description. Models built on this principle are known as *hidden Markov models*.

For our Hadamard pair, a hidden memory is easy to build. Give the bit a concealed companion bit, initialized to 0, which flips to 1 when a Hadamard acts on a fresh bit and back to 0 when a Hadamard acts the second time. As a stochastic matrix on the joint (memory, bit) system, taking memory as the most significant bit, the behavior of “the Hadamard gate plus its bookkeeping” becomes a single, fixed, perfectly Markovian matrix:

$$\Gamma = \begin{pmatrix} 0 & 0 & 1 & 1 \\ 0 & 0 & 0 & 0 \\ 0.5 & 0.5 & 0 & 0 \\ 0.5 & 0.5 & 0 & 0 \end{pmatrix}.$$

Column 00 (memory 0, bit 0): output uniform over {10, 11}. In effect, flip the memory on and randomize the bit. Column 10 (memory 1, bit 0): output 00 deterministically. In effect, reset the memory and zero the bit. One matrix, no piecewise cases, Markov property restored. The reader can verify that $\mathbf{e}_{00}$ evolves to a uniform distribution over the bit and then back to $\mathbf{e}_{00}$, matching the quantum statistics exactly.

This toy works only for this specific circuit, but the strategy generalizes, and the generalization is orthodox quantum mechanics itself. Picture any quantum program split into two parts: Program A, everything up to some time interval, and Program B, the very next instruction. The hidden Markov model says: Program A does not merely hand Program B a bit; it hands over a bit *and a hidden memory*, both updated at every step, and Program B reads both.



And in the orthodox formulation of quantum information, we can say precisely what the hidden memory is: *the quantum state*. The quantum state is evolved through the circuit alongside the bits, updated by every instruction, and the statistical distribution of the bits at each moment is given by the Born rule applied to it. The evolution rules of the hidden Markov model that is textbook quantum mechanics are simply

$$\mathbf{p}^{[t+1]} = \left| U^{[t+1 \leftarrow t]} \psi^{[t]} \right|^2, \quad \psi^{[t+1]} = U^{[t+1 \leftarrow t]} \psi^{[t]} :$$

one rule for the statistics of the bits, one rule for the hidden memory that makes those statistics Markov-computable. The second Hadamard gate is no magician on this picture: it reads the memory, in which $\psi = \frac{1}{\sqrt{2}}(1, 1)$ and $\psi = \frac{1}{\sqrt{2}}(1, -1)$ are different entries, even though their Born distributions agree.

5.4 The trouble with hidden memories

The hidden Markov picture explains the trick. Now examine what it commits us to.

First: the memory is, indeed, hidden. Not hidden as in unmeasured, but hidden as in *unmeasurable*. Whenever the memory state of a quantum computer is observed, what is observed are bits with definite values. Never, under any circumstances, is the quantum state itself read out of the hardware; it is reconstructible only by tomography over an ensemble of many runs, which is exactly how one reconstructs a *statistical distribution*, not how one reads a physical register. An entity posited to exist in every individual run, which no individual run can ever exhibit, has a legitimacy problem.

Second: the memory does not fit anywhere. The quantum state grows as 2^n in the number of bits. For a 3-bit system it has 8 complex entries. Where do these 8 entries physically reside? They cannot be divided evenly among the 3 bits, since 8 does not divide by 3, and no assignment of pieces of ψ to regions of the hardware survives inspection, because entangled states are exactly the ones that do not factor into local pieces. The hidden memory, taken as ontology, is a physical object with no physical address.

And 2^n is the easy case, finite only because the system is a finite collection of bits. Let any parameter of the world be continuous, as position in space is ordinarily taken to be, and the configurations stop forming a finite list: the memory would then hold an entry for every point of a continuum. In that case it would have to be infinitely large.

Third, there are those who would resolve the discomfort by boldness: declare that *only* the hidden memory exists. Only the quantum state flows through the circuit; the bits, with their definite values, are appearances to be explained away. This position runs into a contradiction. A hidden Markov memory is, by definition, hidden, whereas when the memory of a quantum computer is measured, it is precisely bits which are measured: every observation, every readout, every datum ever collected in support of quantum mechanics has consisted of definite values of observables.

If only the quantum state exists, then what is the thing measurement reveals, and why does the formalism's own Born rule speak of nothing else? A theory whose sole ontology is the one object that never appears, and whose every appearance is of objects the theory says do not exist, has made measurement impossible. This paradox is, of course, the measurement problem, and we take it at face value: it shows by contradiction that the ψ -only interpretation is not a viable resting place. (Its sophisticated descendants, Many Worlds chief among them, earn a fuller hearing, and a fuller refusal, in Chapter 9.)

The complaint is older than quantum computers. Einstein asks us to take a radioactive atom of known average decay time and wait. After an hour, has it decayed? The quantum state spreading out past the barrier gives the probability of finding the emitted particle in one place or another, and, as he puts it, “does not imply any assertion concerning the time instant of the disintegration of the radioactive atom.” One is inclined to assume that the individual atom decays at a definite time; no such time appears in the description.¹ The question has two possible answers, and the quantum state resembles neither of them, which is what makes it hard to read as a description of the single atom. Einstein draws the conclusion as a conditional. If the individual atom has a definite disintegration time, then as regards that atom the description “must be interpreted as an incomplete description,” and the quantum state is to be taken as “the description, not of a singular system, but of an ideal ensemble of systems.”

Einstein then credits Schrödinger with the step that makes this impossible to set aside. Couple the atom to something that records what it did. A counter trips a hammer, the hammer shatters a flask of acid, and a cat is in the box. The chain is ordinary physics throughout, so whatever blurring the theory assigns the nucleus is inherited by the cat, and the quantum state of the whole, in Schrödinger's words, has “the living and the dead cat (pardon the expression) mixed or smeared out in equal parts.” His own gloss on what this shows is the one to keep: “an indeterminacy originally restricted to the atomic domain becomes transformed into macroscopic indeterminacy, which can then be resolved by direct observation. That prevents us from so naively accepting as valid a ‘blurred model’ for representing reality.”² Cats are met alive or dead, which is the same report the bits give at every readout.

So the hidden-memory escape saves the Markov property by positing an unmeasurable, unlocatable object, and makes no sense at all once that object is promoted to sole reality, since what it describes is then nothing anybody perceives. Chapter 9 takes the matter up at length. What exactly were we so determined to save? The Markov property is an assumption about the *shape* of statistical laws. Perhaps the lesson of the Hadamard pair is simply that the assumption is false.

5.5 History instead of memory

Here is the alternative we develop. Assume the system is genuinely memoryless (the bits are all there is), but allow the *laws* to be non-Markovian: let the evolution law simply take into account the history of the system, not merely its immediately previous time interval. The

¹Einstein, “Remarks Concerning the Essays Brought together in this Co-operative Volume,” pp. 667–669.


²Schrödinger, “Die gegenwärtige Situation in der Quantenmechanik,” *Naturwissenschaften* 23 (1935), §5 and §9, translated by J. D. Trimmer as “The Present Situation in Quantum Mechanics,” *Proceedings of the American Philosophical Society* 124 (1980), 323–338.

history-dependence lives at the level of the nomology, not the ontology. The behavior of an operator is a function of everything that comes before it in the circuit.

Two clarifications before the mathematics, because this proposal reliably attracts two misreadings.

It is not retrocausality. A non-Markovian law is non-temporal only in the sense that the entire past influences the present, rather than the instantaneous past alone. Nothing about the future acts backward; causation still runs in one direction; the law simply reaches further back than the Markov assumption permits, and reaches to the past itself rather than to any record of it.

And it is not a hidden memory. The difference is exactly the difference between Program B receiving an extra register updated by Program A, and Program B simply *consulting Program A itself*, reaching back to the entire preceding circuit, as a matter of law:



If the whole of Program A is available to Program B, no memory is needed: Program B can compute, from the definition of the preceding instructions and the initial condition, everything the hidden register would have told it, including whether the bit arriving now belongs to a degenerate or a uniform ensemble, even in a single case. The hidden Markov model *stores* the past in the present; the non-Markovian law *consults* the past directly. Ontologically, the first picture has bits plus an unmeasurable 2^n -dimensional object carried alongside them; the second has bits alone.

The rest of what follows turns this from a proposal into a derivation. The demonstration has two steps: first, split the quantum state into a statistical part and a phase part, and rewrite the dynamics as rules for those two parts; second, show that the phase part can be expanded away into the statistical history, terminating at a base case. What survives is a formulation of quantum dynamics in which every quantity is a probability or an operator, and in which the dependence on history is not a mystery but the visible, load-bearing structure of the law.

5.6 The polar decomposition

Complex numbers can be decomposed into two real numbers. One way is Cartesian, splitting into real and imaginary parts, $\mathbf{x} = \Re(\psi)$, $\mathbf{y} = \Im(\psi)$, and it gains us nothing: neither vector means anything on its own. The other way is polar, and it is exactly suited to our question. Every amplitude is its magnitude times a phase,

$$\psi = |\psi| \odot e^{i \arg(\psi)},$$

element-wise, so the two real vectors $|\psi|$ and $\arg(\psi)$ fix ψ completely. We will carry the square of the first rather than the first itself:

$$\mathbf{p} = |\psi|^2 \quad \mathbf{m} = \arg(\psi),$$

again element-wise. This is a variant of the polar decomposition rather than the textbook one, and it is still a decomposition in the strict sense: squaring costs nothing, since magnitudes are never negative and $\sqrt{\mathbf{p}}$ returns them, so the pair $(\mathbf{p}, \mathbf{m})$ fixes ψ exactly as $(|\psi|, \arg \psi)$ does. What the variant gains is the thing we came for. Computing $\mathbf{p}$ is the Born rule, so $\mathbf{p}$ is guaranteed to be a probability vector (the system's present statistical distribution, exactly the object Chapter 2 taught us to handle). The second vector, which we will call the *phase vector*, is a list of angles, one per configuration. The quantum state is reassembled as $\psi = \sqrt{\mathbf{p}} \odot e^{i\mathbf{m}}$ (element-wise square root, element-wise exponential); we saw this notation at work in Chapter 4's misleading-CMI example. Operators decompose the same way:

$$\Gamma = |U|^2 \quad \Theta = \arg(U),$$

the first factor being precisely the unistochastic matrix of Chapter 3. The plain magnitudes would have given neither: not a probability vector, and not a stochastic matrix.

It is possible to evolve $\mathbf{p}$ and $\mathbf{m}$ directly, without ever assembling a quantum state, and the derivation takes four steps. Substitute the polar forms into $\psi^{[t+1]} = U\psi^{[t]}$:

$$\psi_j^{[t+1]} = \sum_k U_{jk}^{[t+1 \leftarrow t]} \psi_k^{[t]} = \sum_k \sqrt{\Gamma_{jk}^{[t+1 \leftarrow t]}} p_k^{[t]} e^{i(m_k^{[t]} + \Theta_{jk}^{[t+1 \leftarrow t]})}.$$

The new amplitude is a sum of complex numbers. Write z_k for the k th of them, with magnitude $\sqrt{\Gamma_{jk} p_k}$ and angle $m_k + \Theta_{jk}$. The two rules are then the magnitude and the angle of $\sum_k z_k$, read off in turn.

Take the magnitude first, squared, since that is what $\mathbf{p}$ holds:

$$\left| \sum_k z_k \right|^2 = \sum_k |z_k|^2 + \sum_{k < l} 2 \Re(z_k \bar{z}_l),$$

where the cross terms have been collected in conjugate pairs using $z_k \bar{z}_l + z_l \bar{z}_k = 2 \Re(z_k \bar{z}_l)$, which is where the sum over *pairs* comes from and why it runs over $k < l$ rather than over all k and l . Each piece is now immediate. The first is $|z_k|^2 = \Gamma_{jk} p_k$. In the second, $z_k \bar{z}_l$ has magnitude $\sqrt{\Gamma_{jk} \Gamma_{jl} p_k p_l}$ and angle $(m_k - m_l) + (\Theta_{jk} - \Theta_{jl})$, so its real part is the first of those times the cosine of the second. Substituting both gives the discrete time-evolution rule for the probability vector:

$$p_j^{[t+1]} = \sum_k \Gamma_{jk}^{[t+1 \leftarrow t]} p_k^{[t]} + \sum_{k < l} 2 \sqrt{\Gamma_{jk}^{[t+1 \leftarrow t]} \Gamma_{jl}^{[t+1 \leftarrow t]}} p_k^{[t]} p_l^{[t]} \times \cos\left((m_k^{[t]} - m_l^{[t]}) + (\Theta_{jk}^{[t+1 \leftarrow t]} - \Theta_{jl}^{[t+1 \leftarrow t]})\right),$$

The angle comes off the same sum. Since $\Re(z_k) = \sqrt{\Gamma_{jk} p_k} \cos(m_k + \Theta_{jk})$ and $\Im(z_k)$ is the same with a sine in place of the cosine, the argument of $\sum_k z_k$ is the arctangent of the summed imaginary parts over the summed real ones, which is the rule for the phase vector:

$$m_j^{[t+1]} = \arctan 2 \left(\sum_k \sqrt{\Gamma_{jk}^{[t+1 \leftarrow t]} p_k^{[t]}} \sin(m_k^{[t]} + \Theta_{jk}^{[t+1 \leftarrow t]}), \right. \\ \left. \sum_k \sqrt{\Gamma_{jk}^{[t+1 \leftarrow t]} p_k^{[t]}} \cos(m_k^{[t]} + \Theta_{jk}^{[t+1 \leftarrow t]}) \right).$$

These two rules are exactly equivalent to unitary evolution, since nothing has been changed but coordinates, and the first one should be read slowly, term by term, because the entire interpretive content of quantum mechanics is sitting in it.

The first summation is $\sum_k \Gamma_{jk} p_k$: a Markov chain. Present distribution in, stochastic matrix on, next distribution out; this is Chapter 2 verbatim, with the unistochastic matrix $|U|^2$ playing the stochastic matrix. If the first summation were the whole rule, quantum mechanics would *be* classical probabilistic computing.

The second summation is everything else. Each of its terms couples a *pair* of distinct configurations $k < l$: their probabilities enter jointly under a square root, and the cosine of their relative phase (shifted by the operator's phases) multiplies the term, swinging it positive or negative. Negative terms are how probability *cancels*. This is the interference the complex numbers of Chapter 3 promised, now visible as an explicit correction to a Markov chain. And the term's structure explains at a glance the two facts we have been circling since Chapter 3. The first is why the correction depends on pairs of configurations. Every term in the second summation carries the factor $\sqrt{p_k p_l}$ for two *distinct* k and l , so a configuration contributes nothing on its own: the correction needs two configurations that both have probability to spare. When the distribution sits entirely on one configuration, every cross term is multiplied by zero and the whole correction vanishes, leaving a plain Markov chain. The second is why the correction is invisible to the Born statistics of the present moment: the cosine lives in the dynamics, not in $\mathbf{p}$ itself.

So the polar decomposition gives us a precise anatomy: *quantum dynamics is a Markov chain plus a phase-steered interference correction*. The question “what kind of statistical process is a quantum circuit?” reduces to the question of what steers the correction. The steering data is $\mathbf{m}$. If $\mathbf{m}$ were an irreducible physical ingredient, we would simply have rediscovered the hidden memory of Section 5.3 in polar coordinates, with all its problems intact. Our final move is to show that $\mathbf{m}$ is not irreducible.

5.7 The recursive expansion

Let us, as shorthand, write the two evolution rules as functions taking as inputs the relevant instruction and the two decomposed vectors:

$$\mathbf{p}^{[t+1]} = f\left(\Gamma^{[t+1 \leftarrow t]}, \Theta^{[t+1 \leftarrow t]}, \mathbf{p}^{[t]}, \mathbf{m}^{[t]}\right) \quad \mathbf{m}^{[t+1]} = g\left(\Gamma^{[t+1 \leftarrow t]}, \Theta^{[t+1 \leftarrow t]}, \mathbf{p}^{[t]}, \mathbf{m}^{[t]}\right).$$

Our goal is to represent the quantum circuit as a non-Markovian stochastic process: that is, to express its evolution solely in terms of probability vectors and the definitions of the operators, with no phase vector among the state variables. If we know the definition of a variable in terms of an equation, we can eliminate it from another equation through substitution. The phase vector at time t is itself the output of g applied one step earlier, so we substitute:

$$\mathbf{p}^{[t+1]} = f\left(\Gamma^{[t+1 \leftarrow t]}, \Theta^{[t+1 \leftarrow t]}, \mathbf{p}^{[t]}, g\left(\Gamma^{[t \leftarrow t-1]}, \Theta^{[t \leftarrow t-1]}, \mathbf{p}^{[t-1]}, \mathbf{m}^{[t-1]}\right)\right).$$

This, clearly, fails to get rid of the phase vector: it merely replaces a dependence on its value at the current time with a dependence on its value at the previous time. But nothing stops us from substituting again, pushing the residual $\mathbf{m}$ to time $t - 2$, and again, and again, each

substitution trading the phase's timestamp for one more layer of the system's statistical past:

$$\mathbf{p}^{[t+1]} = f\left(\Gamma, \Theta, \mathbf{p}^{[t]}, g\left(\Gamma, \Theta, \mathbf{p}^{[t-1]}, g\left(\Gamma, \Theta, \mathbf{p}^{[t-2]}, g(\dots)\right)\right)\right).$$

An expansion that never terminates would be a failure: an infinite regress in place of a definition. The expansion terminates if, and only if, there exist well-defined conditions under which the phase vector ceases to be relevant, so that the innermost g can be deleted and the recursion grounded. Do such conditions exist?

There is a clear one, and the mechanism behind it is already familiar. Consider the case in which the probability vector at some past time is a *degenerate distribution*. Look again at the second summation of the evolution rule for $\mathbf{p}$: under the square root, every term multiplies the probabilities of two *different* configurations, $p_k p_l$ with $k < l$, since the sum runs over $k < l$, never $k = l$. In a degenerate distribution, all elements are 0 except one, and so the product of two different elements is necessarily always 0. Every term of the second summation vanishes, the cosine never gets the chance to matter, and with it goes every dependence on $\mathbf{m}$:

$$p_j^{[t+1]} = \sum_k \Gamma_{jk}^{[t+1 \leftarrow t]} p_k^{[t]}.$$

The equation that remains is identical to that of a Markov chain, with the stochastic matrix given by the square magnitude of the unitary. (The same collapse happens inside g : from a degenerate distribution, the next phase vector is a function of the operator alone.) Hence, if the distribution at some moment is degenerate, then the step that follows it is Markovian, the phase upstream of that moment is irrelevant to everything downstream, and the recursive expansion stops there.

Call such a moment a *base case*. Every quantum circuit has one: the initial state. A quantum computer's memory begins in a definite configuration (all zeros, conventionally), which is a degenerate distribution. (Some presentations define circuits as beginning in a superposition; but every superposition is an eigenstate on some basis, so one can always prepend the unitary that manufactures it from a definite configuration and recover a degenerate start.) And this immediately explains an observation from Section 5.1 that had to wait: the *first* instruction in any quantum circuit has no history, and its behavior is exactly the Markov chain of its unistochastic matrix, which is why the first Hadamard behaved as $|\mathbf{H}|^2 = \Gamma_{\text{coin}}$, no questions asked. Instructions after the first inherit a history, and deviate from Markovianity precisely as the second summation directs.

Thus the expansion below is not an infinite regress but a finite recipe, unwinding back to the base case and stopping:

$$\mathbf{p}^{[t+1]} = f\left(\Gamma, \Theta, \mathbf{p}^{[t]}, g\left(\Gamma, \Theta, \mathbf{p}^{[t-1]}, g\left(\dots g\left(\Gamma^{[1 \leftarrow 0]}, \Theta^{[1 \leftarrow 0]}, \mathbf{p}^{[0]}\right)\right)\right)\right), \quad \mathbf{p}^{[0]} \text{ degenerate.}$$

Read what this equation says. The probability distribution at the next time interval is a function of: the definitions of the operators (given by the circuit), and the probability distributions at every prior time interval back to a base case. Nothing else. No quantum state. No phase. No hidden memory. The system's evolution has been expressed solely in terms of statistics and the laws: a genuinely, irreducibly non-Markovian stochastic process, exactly as Section 5.5 proposed. What was a proposal there has now been derived.

One refinement is needed before moving on. A base case need not be a *known* configuration; it need only be a definite one. Suppose we prepare the initial bit at random and measure it, so that in each individual run it has a definite (hence degenerate) value, while over the ensemble of runs the values follow some distribution. Every run grounds its own expansion in its own base case; the ensemble of expansions is governed by the statistical distribution of base cases, which may serve, itself, as a non-degenerate “base distribution” for the process as a whole. Degeneracy is required run by run, where the physics happens, not of our ensemble-level description of it. (In the next chapter we will find that measurement-like interactions *inside* a circuit create fresh base cases mid-computation, which is where this refinement starts pulling real weight.)

5.8 The phase as a sufficient statistic

Let us step back and consider what has been established.

The phase vector has been shown to be expandable, without residue, into a function of the operators and the system’s past probability distributions. So ask the question Chapter 2 trained us to ask: if the phase is a function of the history that suffices (together with the present distribution) to predict the next step, *what kind of thing is it?* It is a sufficient statistic. That is the definition. The phase vector is a compressed summary of the system’s statistical history, maintaining exactly what the second summation needs (the pairwise interference data) and discarding the rest. If all prior probabilities going back to a base case were provided, together with the definitions of the current and all prior instructions, that would be just as sufficient for computing the next distribution as being provided the probability and the phase. Both routes compute the same future; the phase is the walker’s position, the history is the sequence of steps.

Every property that made the quantum state embarrassing as *ontology* in Section 5.4 is, for a sufficient statistic, exactly what one expects. Unmeasurable in an individual run? Of course: a position is not one of the steps that produced it; a summary of a history is not a component of the present. Nowhere to live, 2^n entries among n bits? Of course: it is not stored in the hardware; it is computed *about* the hardware’s past, and functions of a history do not occupy space. Updated “instantaneously” by distant events, growing exponentially, invisible to every probe: each scandal, reread as a fact about bookkeeping, is a banality. The wavefunction formalism survives untouched as what it always was operationally, namely the most convenient way anyone has found to carry the relevant history forward step by step, and we will go on using it for calculation without apology. What does not survive is the obligation to house it in the world.

We opened Part II asking what kind of statistical process a quantum circuit is. The answer we have established: *a memoryless, non-Markovian stochastic process whose per-step law consults the system’s statistical history back to a base case, and whose famous “state” is not a constituent of the system but a summary of it in two parts: the present distribution $\mathbf{p}$ over configurations, and a phase vector $\mathbf{m}$ that is the sufficient statistic of the history behind it.* The bits always have values; the laws consult the history; the phase is a summary of that history, carried for convenience.

What we have *not* yet done is make the picture fully operational. Two gaps remain. First, degenerate distributions cannot be the whole story about base cases: real circuits entangle

their bits with environments and measuring devices mid-flight, and we claimed such events create base cases too. That claim needs support, and the tools of Chapter 4 supply it. Second, our expansion computes the statistics at the ends of circuits and at base cases, but a statistical process worthy of the name should say what happens at *every* time interval: there should be a definite, unique single-step transition matrix at each step, not merely a multi-step summary. Both gaps are closed in the next chapter, and closing the second one is where we part company with others who have walked the road this far.

Code: this chapter in Octave

```
% --- the two polar rules of Section 5.6, verbatim ---

1; % script marker so Octave allows function definitions below

function [p2, m2] = polar_step(U, p, m)
    G = abs(U).^2; Th = arg(U); n = length(p);
    p2 = zeros(n,1); m2 = zeros(n,1);
    for j = 1:n
        p2(j) = G(j,:) * p; % Markov chain term
        for k = 1:n
            for l = k+1:n % interference term
                p2(j) = p2(j) + 2*sqrt(G(j,k)*G(j,l)*p(k)*p(l)) ...
                    * cos((m(k)-m(l)) + (Th(j,k)-Th(j,l)));
            end
        end
        w = sqrt(G(j,:))' .* p;
        m2(j) = atan2(w' * sin(m + Th(j,:))', w' * cos(m + Th(j,:)));
    end
end

H = (1/sqrt(2))*[1 1; 1 -1]; S = [1 0; 0 1i];
T = [1 0; 0 (1+1i)/sqrt(2)];

% --- the Hadamard pair, no wavefunction anywhere ---
p = [1; 0]; m = [0; 0]; % base case: degenerate
[p, m] = polar_step(H, p, m); disp(p') % -> 0.5 0.5
[p, m] = polar_step(H, p, m); disp(p') % -> 1 0

% --- longer circuit with nontrivial phases, checked against psi ---
p = [1; 0]; m = [0; 0]; psi = [1; 0];
for U = {H, T, H, S, H}
    [p, m] = polar_step(U{1}, p, m);
    psi = U{1} * psi;
end
disp([p, abs(psi).^2]) % identical columns

% --- two-bit circuit ---
I2 = eye(2); CX = [1 0 0 0; 0 1 0 0; 0 0 0 1; 0 0 1 0];
p = [1;0;0;0]; m = zeros(4,1); psi = [1;0;0;0];
for U = {kron(H,I2), CX, kron(T,S), CX, kron(H,H)}
    [p, m] = polar_step(U{1}, p, m);
    psi = U{1} * psi;
end
disp([p, abs(psi).^2]) % identical columns
```

Chapter 6

Unistochastic Dynamics

Chapter 5 established what quantum dynamics *is*: a non-Markovian stochastic process, whose law consults the system's statistical history back to a base case, and whose phase is the sufficient statistic of that history. It ended with two questions open. First, degenerate initial states cannot be the only base cases, because real circuits create fresh ones mid-flight, and we need to say exactly when. Second, a statistical process worthy of the name should say what happens at *every* time interval rather than just at base cases and at the ends of circuits. We answer both, and answering the second is where we make its most distinctive commitment.

6.1 Complex numbers as bookkeeping

The recursive expansion of Chapter 5 was written in the polar variables $\mathbf{p}$ and $\mathbf{m}$, with the functions f and g standing in for two bulky trigonometric rules. If we allow ourselves to reintroduce complex numbers (now demoted, permanently, to notation), the same dynamics compresses into two short lines. Reassemble the polar data into a quantum state, evolve, and decompose again:

$$\mathbf{p}^{[t+1]} = \left| U^{[t+1 \leftarrow t]} \left(\sqrt{\mathbf{p}^{[t]}} \odot e^{i\mathbf{m}^{[t]}} \right) \right|^2 \quad \mathbf{m}^{[t+1]} = \arg \left(U^{[t+1 \leftarrow t]} \left(\sqrt{\mathbf{p}^{[t]}} \odot e^{i\mathbf{m}^{[t]}} \right) \right),$$

where $\odot$ denotes element-wise multiplication and the square root is element-wise. These are the same f and g as before (the reader may expand the square magnitude and recover Chapter 5's two summations), but the complex exponential now does the trigonometry for us. The same recursive substitution method applies verbatim: replace $\mathbf{m}^{[t]}$ with its definition one step earlier, and the update rule becomes a function of operators and probability distributions stretching back into the past,

$$\mathbf{p}^{[t+1]} = \left| U^{[t+1 \leftarrow t]} \left(\sqrt{\mathbf{p}^{[t]}} \odot e^{i \arg \left(U^{[t \leftarrow t-1]} \left(\sqrt{\mathbf{p}^{[t-1]}} \odot e^{i \arg(\dots)} \right) \right)} \right) \right|^2,$$

terminating, as before, at a base case. Nothing new has been claimed; the notation is simply more compact. We stress the point once, because it applies to everything here: from here on, whenever a quantum state or a complex amplitude appears, it appears as compression (a sufficient statistic being carried forward because carrying the full history is tedious), and every formula that uses one could be mechanically unwound into probabilities and operators alone.

6.2 The multi-step shortcut

The nested expansion is correct but clumsy. There is a drastic simplification available, and Chapter 3 already supplied it. Since matrix multiplication is associative, the entire circuit between two times can always be composed into a single multi-step unitary $U^{[t \leftarrow 0]}$, and the quantum state at any time is $\psi^{[t]} = U^{[t \leftarrow 0]} \psi^{[0]}$: one instruction, applied once.

Now let time 0 be a base case: a degenerate distribution, in each run some definite configuration b , so that $\sqrt{\mathbf{p}^{[0]}} = \mathbf{e}_b$ and the phase is irrelevant. Then

$$\mathbf{p}^{[t]} = \left| U^{[t \leftarrow 0]} \mathbf{e}_b \right|^2.$$

But multiplying a matrix by $\mathbf{e}_b$ merely selects its b -th column, and the square magnitude of a column of U is the corresponding column of $|U|^2$. So the whole history collapses into one Markov-like step under the unistochastic matrix of the composed circuit:

$$\mathbf{p}^{[t]} = \left| U^{[t \leftarrow 0]} \right|^2 \mathbf{p}^{[0]}.$$

If the base case is statistical (a distribution over definite configurations, in the sense of Chapter 5), each run selects its own column and the ensemble takes the weighted mixture, which is exactly what the right-hand side computes; the equation holds unchanged.

This discrete time-evolution rule is called *unistochastic dynamics*, and it fully captures the statistics of quantum information entirely in terms of operators and probability vectors, with no quantum state introduced at all. The probability distribution at every time interval is obtained merely by changing the value of t , which recomposes the circuit up to the moment of interest and applies its unistochastic matrix to the base case.

The fine print matters, and both of its clauses are already established. The rule yields correct predictions *only if the distribution on the right-hand side is a valid base case*: feed it a mid-history distribution (the uniform output of a first Hadamard, say) and it will predict the wrong thing, because the interference terms that a non-base-case history carries have been discarded. And the rule is *not composable*: Chapter 3 proved $|\prod U_k|^2 \neq \prod |U_k|^2$, so the unistochastic matrix of a circuit is not the product of the unistochastic matrices of its steps. Unistochastic dynamics is a base-case-to-anywhere law, not a step-to-step law. Both limitations point at the same missing ingredient: more base cases.

6.3 Base cases that appear mid-circuit

The initial state of a circuit is a base case by fiat. When does one appear in the middle of a circuit, with nobody putting it there?

The name invites the wrong picture. Nothing about the world starts over mid-circuit. Decoherence destroys no history: the phases are still there, lodged in correlations with whatever the bit interacted with. What changes is that those correlations stop bearing on any prediction about the bit by itself. A base case of this kind belongs to a description that has stopped following part of the world, and it exists only relative to that description.

Consider two bits, A and B. We are curious whether bit A can be conditioned on at a given time interval as the beginning of its own chain of historical dependence: that is, whether

the recursion for A's marginal statistics, unwound from some later time, may stop here rather than continuing back to the circuit's start. Two conditions must be satisfied:

First, A must interact with B such that the quantum mutual information between them becomes maximal. Recall from Chapter 4 what maximal QMI does to the marginal: tracing out B leaves A's reduced density matrix at minimal purity: fully decohered, off-diagonals gone, a plain classical distribution with no phase relations left *internal to A*. All of A's phase information has migrated into the joint correlations with B.

Second, A and B must not interact again for the duration under consideration. The phase information that migrated into the joint system is not destroyed. Chapter 4's misleading-CMI example showed it can be brought back by operations that touch both bits. If B returns to interact with A, the history buried in their correlations resurfaces, and A's statistics again depend on things that happened before the would-be base case. But if B never comes back, that buried history is permanently out of reach for every operator that will ever again touch A.

Given both conditions, we can drop B from consideration entirely by marginalizing over it. Bits dropped this way are referred to as the *environment*, as they are not part of the system actually being tracked. And the claim is: *the moment of last interaction is a valid base case for the marginal statistics of A*. We can treat a bit as a fresh starting point of history exactly when the QMI between it and its environment is maximal,

$$I(\text{bit} : \text{environment}) = S(\rho^{\text{bit}}) + S(\rho^{\text{environment}}) - S(\rho^{\text{bit, environment}}) \quad \text{maximal,}$$

and the environment is gone for good.

Note the word *marginal*. Base cases only occur on marginals. If we continue to compute the full joint distribution of A together with B, no new base case appears, because the joint state retains every phase relation and its history runs unbroken back to the circuit's start. It is only the reduced description of A alone, after the environment departs, that gets to restart its clock. David Albert makes the same point against the name *division event*: the division looks real so long as attention stays on the measured system, and the appearance goes as soon as the recording system is counted in, leaving the indivisibility of the whole “absolute and inviolable.”¹ (The reader may notice the family resemblance to Chapter 5's statistical base cases: a decohered marginal *is* a statistical distribution over definite configurations, which is exactly what run-by-run degeneracy requires.)

Widening the description is what makes a base case disappear, which invites a worry: widen far enough and none are left, so nothing grounds any of it. The widening has an end, though. Enlarge the description until nothing is outside it and one base case survives, the configuration the universe began in. That one is relative to nothing, since there is no larger system to describe it from, and what it was does not matter. It need only have been definite, which is all Section 5.7 ever asked of a base case. Everything after it, apparent base cases included, is the world running on from there.

One may charge that the beginning of the universe is a meaningless concept. Then so is the whole universe taken at once, and the worry goes with it. The objection only runs if there is a description with nothing outside it; deny that there is one and the widening never gets far enough to leave us without a base case.

¹D. Z. Albert, “Quantum mechanics without a wave-function,” *Synthese* 208 (2026), article 141, §4.

That first configuration does carry empirical consequences, since every later probability is computed from it. Albert takes this as a measure of how radical such a law is: a world that came into being five minutes ago, in exactly the state the world was then actually in, would differ observably from one that began with the Big Bang.² We grant the consequence and are untroubled by it. Nobody is in a position to know the configuration the universe began in, so no prediction follows from it that anyone can compute, and the picture gains no observable feature on that account which any other reading of quantum mechanics lacks.

Back to the condition itself. Checking whether a base case has been created requires computing QMI mid-circuit, which requires a density matrix at an arbitrary time interval. Chapter 4 gives it in one line: starting from a base case with distribution $\mathbf{p}^{[0]}$, place that distribution on the diagonal and conjugate by the composed circuit,

$$\rho^{[t]} = U^{[t \leftarrow 0]} \text{diag}(\mathbf{p}^{[0]}) \left(U^{[t \leftarrow 0]} \right)^\dagger,$$

then trace out whatever subsystem is under consideration and compute the three entropies.

In the circuits here the QMI condition holds exactly, since a CX gate manufactures perfect correlation. In the wild it rarely does. An environment records a system imperfectly, so the buried history is suppressed rather than put wholly out of reach, and base cases are *approximate*. The approximation becomes superb when the environment is macroscopic, for the same reason decoherence does.

Their ubiquity explains something presupposed since Chapter 2: *why the everyday world looks Markovian at all*. A system in constant contact with an eavesdropping environment has its history carried off at every bump and collision, and a process whose base cases appear faster than its interference terms can matter is a Markov chain. Section 5.6 shows why that suffices.

A decohered marginal is a statistical distribution over definite configurations, so run by run it sits on one of them. The interference correction there is a sum running over pairs of distinct configurations, and each piece of that sum carries a factor $\sqrt{p_k p_l}$ built from the pair's two probabilities. Put all the probability on one configuration and every one of those factors is zero. What remains is the Markov chain the correction was added to. That is why classical statistical mechanics never needed a phase vector. Quantum behavior, on this view, is what any system's statistics look like in the rare quiet between recordings.

6.4 A worked example: the interrupted Hadamard pair

Here is the whole machinery in the smallest circuit that exercises it. Take the Hadamard pair of Chapter 5 (which, uninterrupted, returns the bit deterministically to 0) and insert a CX between the two Hadamard gates, recording the first bit onto a second:

$$\begin{array}{ccccccc} \mathbf{e}_0 & \text{H} & \overline{\text{CX}} & \text{H} & \perp \\ \mathbf{e}_0 & & \text{CX} & & \perp \end{array}$$

²The case is Bertrand Russell's five-minute hypothesis. Albert, "Quantum mechanics without a wave-function," §3.

Compute the quantum state at the end (the reader has every tool required):

$$(\mathbf{H} \otimes \mathbf{I}_2) \text{CX} (\mathbf{H} \otimes \mathbf{I}_2) (\mathbf{e}_0 \otimes \mathbf{e}_0) = \frac{1}{2} \begin{pmatrix} 1 \\ 1 \\ 1 \\ -1 \end{pmatrix},$$

whose Born distribution is uniform over all four configurations. Marginalize down to the first bit alone:

$$(\mathbf{I}_2 \otimes \begin{pmatrix} 1 & 1 \end{pmatrix}) \left| \frac{1}{2} \begin{pmatrix} 1 \\ 1 \\ 1 \\ -1 \end{pmatrix} \right|^2 = \begin{pmatrix} 0.5 \\ 0.5 \end{pmatrix}.$$

The first bit ends *uniform*, whereas the bare Hadamard pair ended *degenerate*. The inserted CX changed the statistical behavior of the second Hadamard gate, despite acting between the two and commuting with neither.

The base-case machinery explains this exactly. The CX establishes maximal QMI between the two bits (it manufactured a Bell state, and Chapter 4 put its QMI at the maximum of 2), and the second bit is never touched again. Both conditions hold, so the moment after the CX is a valid base case for the first bit's marginal statistics. From that base case, the remaining circuit is a single Hadamard, and unistochastic dynamics finishes in one line:

$$\mathbf{p}_A^{[3]} = |\mathbf{H}|^2 \mathbf{p}_A^{[2]} = \begin{pmatrix} 0.5 & 0.5 \\ 0.5 & 0.5 \end{pmatrix} \begin{pmatrix} 0.5 \\ 0.5 \end{pmatrix} = \begin{pmatrix} 0.5 \\ 0.5 \end{pmatrix}.$$

The second Hadamard behaves as a plain coin because, for the first bit, *history has been restarted*: the interference terms that would have steered it back to 0 now live in joint correlations with an environment that never returns. This is Chapter 4's decoherence, re-described dynamically: what decoherence *is*, in our language, is the creation of a base case, which puts everything before the recording beyond the reach of the subsystem on its own. And note what the recording did to the gate: one and the same Hadamard, third slot of a circuit, has now exhibited its second distinct stochastic behavior, with the difference controlled by whether a *different* bit got entangled in between. The nomology consults the history; change the history, change the law's output.

6.5 What unistochastic dynamics still lacks

We can now compute correct statistics at any time interval, though only by returning to a base case and recomposing. The evolution rule of unistochastic dynamics does not allow us to divide the system's stochastic evolution into *single-step transition matrices*: we have a law that leaps from base cases to anywhere, and no law that walks.

Why should we want single-step matrices at all? A probability can be a statement about what somebody does not know, so long as one run is under consideration. Over an ensemble it becomes something else. A machine that dispenses a blue marble a quarter of the time and a red one three quarters of the time fills, after some thousands of marbles, two buckets on the floor, and the red bucket holds about three marbles for every one in the blue. That

proportion is a physical fact about the buckets, whatever anyone did or did not know before each marble fell.

A transition matrix is a record of the same kind. Imagine a demon of the sort Laplace described, able to see the configuration before a gate and after without disturbing either. It needs no theory at all. It writes down the value it saw going in and the value it saw coming out, and drops each run into a bucket chosen by the value that went in. The proportions within one bucket are as physical as the proportions in the marble buckets, and dividing a bucket by its size gives one column of the matrix, read off the world the way a census is read off a population. Such an entity could assign a single-step transition matrix to every step of every circuit, simply by watching.

That is to say: even if the evolution law is non-Markovian, if we believe the system possesses an underlying real ontology (bits with definite values at all times, which is precisely what we believe), then its trajectory statistics *must be capable of being fit to a Markov chain*. The single-step transition matrices may be such that computing them requires considering the historical state of the system; the same instruction may warrant different matrices in different contexts; but in each individual step of each individual run, the bits went from some definite configuration to some definite configuration, and the conditional frequencies of those transitions physically exist. They may exist but be unknowable in the same way that, if the evolution of the bits is fundamentally random, their values at a given time interval are believed to exist in the real world but cannot be known without measuring. Existence is not the question; the question is whether we can *compute* them.

The demon is a fiction, and the tally is not. If the bits carry definite values then at every step each bit went from one value to another, those transitions happened in some proportion, and the matrix recording that proportion is as real as the bits themselves, whether or not anyone is placed to count it.

Two constraints pin the answer down almost completely. We know, from the multi-step law, the correct probability distribution $\mathbf{p}^{[t]}$ at every time interval, so whatever the single-step matrix $\Gamma^{[t+1 \leftarrow t]}$ is, it must transport each distribution to its successor:

$$\Gamma^{[t+1 \leftarrow t]} \mathbf{p}^{[t]} = \mathbf{p}^{[t+1]}.$$

And we know one case exactly: an instruction that occurs immediately after a base case has no relevant history, and Chapter 5 proved its transition matrix is precisely the unistochastic matrix $|U^{[t+1 \leftarrow t]}|^2$, with no freedom at all.

A first attempt at a general rule might be piecewise: use the unistochastic matrix whenever it happens to transport correctly, and otherwise fall back on some matrix that forces the right output, such as the rank-one matrix that sends every input to the required distribution:

$$\Gamma^{[t+1 \leftarrow t]} \approx \begin{cases} |U^{[t+1 \leftarrow t]}|^2 & \text{if } |U^{[t+1 \leftarrow t]}|^2 \mathbf{p}^{[t]} = \mathbf{p}^{[t+1]} \\ \mathbf{p}^{[t+1]} \mathbf{1}_n^\top & \text{otherwise.} \end{cases}$$

The definition is discontinuous: an infinitesimal deviation between $|U|^2 \mathbf{p}^{[t]}$ and $\mathbf{p}^{[t+1]}$, no larger than rounding error, flips the assigned matrix between two entirely different ones. In a computer simulation this is untenable, and as physics it is worse: we should expect that a slight change in circumstances yields a slight change in the transition probabilities. We want the two branches to be endpoints of one continuous rule.

6.6 The unique single-step transition matrix

Two requirements pin the matrix down. We want to start with the unistochastic matrix $|U|^2$ as our default (it is exactly right after base cases, and it is the natural reading of “what this gate does”), and, when it fails to transport the current probabilities to the correct successors, *modify it as little as possible* so that it does. “As little as possible” can be made precise using the Kullback–Leibler divergence, the standard measure of statistical distance between two distributions. A stochastic matrix becomes a distribution once weighted by the one it acts on: the chance of finding the system at j now and at i next is $p_j^{[t]} \Gamma_{ij}$ under the fitted law, and $p_j^{[t]} |U|_{ij}^2$ under the gate’s own. The divergence between those two distributions over pairs is what we minimize:

$$\min_{\Gamma^{[t+1 \leftarrow t]}} \sum_{i,j} p_j^{[t]} \Gamma_{ij} \log \frac{\Gamma_{ij}}{|U^{[t+1 \leftarrow t]}|_{ij}^2} \quad \text{subject to} \quad \Gamma^{[t+1 \leftarrow t]} \mathbf{p}^{[t]} = \mathbf{p}^{[t+1]},$$

with Γ ranging over column-stochastic matrices. This problem has no closed-form solution, but it has something just as good: a unique solution expressible as a convergent limit, computed by an algorithm with a long pedigree in statistics known as *iterative proportional fitting*. The solution takes the form of the default matrix rescaled by two diagonal corrections,

$$\Gamma^{[t+1 \leftarrow t]} = \lim_{k \rightarrow \infty} \text{diag}(\mathbf{u}^{(k)}) |U^{[t+1 \leftarrow t]}|^2 \text{diag}(\mathbf{v}^{(k)}),$$

where the two scaling vectors are computed by alternating between two requirements. The first vector rescales the rows so that the matrix correctly transports the computed probabilities; the second rescales the columns so that they always sum to one, keeping the matrix a valid stochastic matrix:

$$\begin{aligned} \mathbf{u}^{(k)} &= \mathbf{p}^{[t+1]} \oslash \left(|U^{[t+1 \leftarrow t]}|^2 \left(\mathbf{v}^{(k)} \odot \mathbf{p}^{[t]} \right) \right) \\ \mathbf{v}^{(k)} &= \begin{cases} \mathbf{1}_n & \text{if } k = 0 \\ \mathbf{1}_n \oslash \left(\left(|U^{[t+1 \leftarrow t]}|^2 \right)^\top \mathbf{u}^{(k-1)} \right) & \text{otherwise,} \end{cases} \end{aligned}$$

$\oslash$ denoting element-wise division. The two processes fight against each other because fixing the rows disturbs the columns and vice versa, so they must be repeated indefinitely, converging toward the unique solution, or in practice repeated to any desired accuracy.

The construction has exactly the properties the piecewise attempt lacked. It is continuous: nearby transport requirements yield nearby matrices. It agrees with the exact result where the exact result is known: if $|U|^2$ already transports correctly (as it always does immediately after a base case), the scaling vectors converge to all ones and the limit *is* $|U|^2$, untouched. And it never invents transitions the gate forbids: entries of $|U|^2$ that are zero remain zero in the limit, since rescaling rows and columns cannot create probability where the unitary put none.

Watch it work on the Hadamard. For the first gate of the pair, the default already transports $(1, 0) \mapsto (0.5, 0.5)$, and the limit returns $|H|^2 = \Gamma_{\text{coin}}$ exactly. For the second gate, the default fails: Γ_{coin} maps $(0.5, 0.5)$ to $(0.5, 0.5)$, not to the required $(1, 0)$. The iteration

converges to

$$\Gamma^{[2\leftarrow 1]} = \begin{pmatrix} 1 & 1 \\ 0 & 0 \end{pmatrix} = \Gamma_{\text{zero}} :$$

the two matrices Chapter 5 read off the statistics, now reached by a single uniform rule that also covers the circuits where the statistics leave the matrix open. And the rule keeps producing new answers where hand-assignment would run out of patience. Insert a T gate between the two Hadamards (circuit H, T, H), and the Born marginals run $(1, 0) \rightarrow (0.5, 0.5) \rightarrow (0.5, 0.5) \rightarrow (0.854, 0.146)$. The T step gets the identity (its unistochastic matrix, which transports correctly), and the final Hadamard gets

$$\Gamma^{[3\leftarrow 2]} = \begin{pmatrix} 0.854 & 0.854 \\ 0.146 & 0.146 \end{pmatrix},$$

a *third* distinct stochastic behavior for the same gate: neither coin nor reset, but a biased reset whose bias is steered by the phase the T gate wrote into the history. Same instruction, three contexts, three transition matrices: the same instruction can be assigned different transition matrices if the instructions prior to it change. This fact does not get rid of the historical dependence or the non-Markovian nature of the stochastic dynamics, since computing each matrix required the two Born distributions it bridges, and computing those required the whole history back to a base case. What the construction removes is not the history but the *mystery*: at every time interval, there is now a definite, unique, computable answer to the question “with what probabilities do the bits transition right here?”

With per-step matrices in hand, simulating the quantum circuit as a stochastic process requires no new machinery at all: it is exactly the stochastic-jump simulation of Section 2.7. Assign the bits a definite value at the beginning of the circuit, sampled from the base-case distribution. At each time interval, look up the column of $\Gamma^{[t+1\leftarrow t]}$ selected by the bits’ current definite configuration, draw one uniform γ , and hop. Even the economics of Section 2.8 survives the trip: since rescaling rows and columns cannot create probability where the unitary put none, each fitted matrix inherits the zeros of $|U|^2$, so a gate that touches one bit yields a transition matrix that never moves any other bit, and the column being sampled is almost entirely zeros. The only thing the quantum case changes is where the matrices come from. In Chapter 2 the program handed them to us; here, each one is fitted by the procedure above, from the two Born distributions it bridges.

The simulation is a working model of our entire ontology: a definite configuration, a column selected by it, a γ , and a hop. In an individual execution the bits have definite values at each time interval, because there is always a fact about where the system is, mid-circuit, gate by gate. Over an ensemble of executions, the statistical distribution of the bits at each time interval precisely matches the Born rule, not merely at the end of the circuit but at every intermediate moment. (The chapter’s code runs exactly this experiment: two hundred thousand trajectories through the H, T, H circuit, matching all four Born distributions to within the sampling error of that many runs.) Nothing in the simulation is a hidden variable in the pejorative sense: the bits are the observables themselves, the things every measurement reveals. Nothing is a wavefunction: no complex vector is carried from step to step. The concession, made openly, is that the transition matrices are context-dependent, since fitting them required the circuit’s history, which is simply non-Markovianity showing up exactly where Chapter 5 said it lives: in the law, not in the system.

A terminological subtlety and a caveat remain. The fitted simulation is, formally, a time-inhomogeneous Markov chain, and time-inhomogeneity alone does not violate the Markov property; the non-Markovianity is the kind Chapter 5 defined, a property of *laws*. Markovianity of law demands that each instruction's matrix be fixed by the instruction itself, and no such assignment reproduces quantum statistics: the Hadamard has just needed three different matrices in three contexts, each computed from the history. The chain does not contradict this; it *embodies* it, being what remains after the non-Markovian law has already been consulted: history in, matrices out, and only then does the memoryless sampling begin. The caveat is that the chain asserts slightly more than quantum mechanics itself does. Quantum theory fixes the distribution at each time, but is silent about *joint* statistics across times (the probability of whole trajectories), since checking those would require measurements, which create base cases and change the circuit. The chain fills that silence in the most economical way, each step conditionally independent of the path so far, given the present. That filling-in is part of the convention we adopt with open eyes, and it is the precise point where the next section's disagreement lives.

The law itself is less of a convention than the filling-in is. Section 6.9 gives the reason for the quantity minimized above, and that reason carries the rescaled form along with it, so the zeros look after themselves. Section 6.10 settles which matrix the rescaling starts from. A reader who grants those two has nothing further to choose, and the law of this section is what they leave.

6.7 Divisibility: where we part company

Everything in Parts I and II so far follows a line of work that others have developed.

The correspondence between unitary quantum dynamics and non-Markovian stochastic processes has been developed with particular force by the physicist Jacob Barandes, whose “stochastic–quantum correspondence”³ establishes in full generality what our circuits have exhibited concretely: every quantum system can be recast as a stochastic process over definite configurations in an ordinary configuration space, with the wavefunction demoted, in his own image, to the fate of the luminiferous aether. On the largest questions, we and he are allies: the system really occupies a definite configuration at every moment and traces a trajectory through its configuration space; superpositions and coherences are artifacts of catching an indivisible process between the special moments he calls *division events*, which are the moments we have been calling base cases, created by exactly the environment-correlation condition of Section 6.3. We draw on this line of work extensively, and the terminological mapping should be made explicit: his indivisible processes are our base-case-to-anywhere unistochastic dynamics; his division events are our base cases; his derivation of Markovian ubiquity from frequently-created division events is similar to what appeared in Section 6.3.

The disagreement lies in what happens *between* base cases, and locating it precisely requires separating two claims that are easily blurred. The first is a theorem, and Barandes is right about it: the multi-step transition matrix generically cannot be *factored*. There is no stochastic matrix $\tilde{\Gamma}^{[t \leftarrow t']}$ satisfying the composition law $\Gamma^{[t \leftarrow 0]} = \tilde{\Gamma}^{[t \leftarrow t']}\Gamma^{[t' \leftarrow 0]}$ as a matrix identity (valid for every column at once, that is, for every counterfactual initial configuration), and his proof is direct: solving for the would-be factor forces negative entries. Nothing here

³J. A. Barandes, “The Stochastic-Quantum Correspondence,” arXiv:2302.10778 (2023).

contradicts that theorem, and note carefully that our fitted matrices do not pretend to: they satisfy the strictly weaker, vector-level requirement of transporting the *realized* distribution $\mathbf{p}^{[t']}$ to $\mathbf{p}^{[t]}$ (the distribution the process actually has, given its actual base case), which is always solvable. The matrix-level factorization fails; the realized process still hops.

The second claim is a policy, and here the paths diverge. Quantum mechanics fixes what Barandes calls the first-order data (the transition behavior from division events) and is silent about the rest: the step-by-step transition probabilities in between, and with them the probabilities of whole trajectories. His response is deliberate abstention. He works with the entire *equivalence class* of stochastic processes sharing the first-order data, leaves the higher-order conditional probabilities unspecified (noting they could, in principle, even vary contingently from one set of runs to another), and expressly leaves to future work whether any theoretical principle picks out one member of the class. The interval between division events is left, on purpose, unmodeled.

Our position is that the question he leaves open has an answer, and we take it up, at least tentatively and for the sake of the analysis. The first candidate anyone would reach for is one we do not adopt, and the four ways it goes wrong are what sent us looking for the law of Section 6.6.

That rejected candidate is *resampling*: give every column of $\Gamma^{[t+1 \leftarrow t]}$ the Born distribution $\mathbf{p}^{[t+1]}$, so that each moment is drawn afresh from the distribution belonging to it. The matrix is $\mathbf{p}^{[t+1]} \mathbf{1}_n^T$, which Section 6.5 has already written down as the fallback branch of its piecewise rule. Resampling is that branch taken at every step, in place of $|U|^2$ even where $|U|^2$ would have transported correctly on its own. It transports correctly at every step by construction, and it reproduces every single-time statistic quantum mechanics predicts, which is all quantum mechanics checks. That is the whole of its appeal, and the four defects that follow are why we set it aside.

Transition matrices can be measured, up to a point. Read a bit, let one gate act, read it again, and keep the runs sorted by what the first reading gave, so that the outputs fall into a separate bucket for each input value; the frequencies in one bucket are one column of the matrix. What comes back is always $|U|^2$, since the first reading makes a base case and $|U|^2$ is exact immediately after one. Resampling returns the same thing there, because a base case leaves the distribution degenerate on the value just read, and redrawing from the next moment's Born distribution then reproduces that very column. The two laws part company only across intervals where no reading was made, which is what Barandes said: no experiment settles the choice. What can still be said is that there is something to be right or wrong about.

The interval has an answer, by the tally of Section 6.5, so we model it rather than leave it open. The model is tentative, adopted for the sake of the analysis, since nothing said here singles out which answer is the right one.

The selection is also unique. The KL-minimal deformation of the unistochastic default exists at every interval, is unique, is continuous, reduces to the exactly-known answer wherever the exact answer is known, and reproduces the Born statistics at every moment. Given that a principled, well-behaved, everywhere-defined selection exists, we see every reason to model the interval and none to leave it unmodeled, since the alternative is a physics that must change the subject whenever asked what a system did between two special moments.

The KL-minimality principle is a *selection rule*: it is the assumption that nature's per-step behavior deviates from the gate's unistochastic default no more than the Born marginals force

it to. One could adopt a different selection rule and obtain a different chain reproducing the same single-time statistics; the intermediate transition probabilities are, in that specific sense, underdetermined by the empirical content of quantum mechanics. They are not *unconstrained*: transport, stochasticity, continuity, and agreement at base cases pin them tightly, and any candidate violating those is simply wrong. Concerns that no experiment settles narrow the field further, among them the hopping described below, since nothing in the physics asks for a particle that crosses between the arms while nothing acts on it.

Underdetermination of this kind will become a recurring character in Part III, where we will meet it again in the choice of ontic basis and the choice of a preferred slicing (also called a preferred frame, or a preferred foliation), and will argue that the correct response is neither denial nor paralysis: one acknowledges that a fact of the matter exists, adopts the most principled available convention as a working model of it, and keeps clear which is which. Where the indivisibilist leaves the underdetermined fact unmodeled, we model it: if we take the limit seriously, and we do, then there is a matter of fact about every time interval, whether or not we can prove we have guessed it exactly.

The three sections below fix the law of Section 6.6, each settling a different part of it. The first takes the four defects, one failure met in four circuits, what resampling does to the zeros of $|U|^2$, and establishes that a law has to keep those zeros. The second takes the minimizing: why a law should minimize anything at all, why that quantity, and why its answer comes out as a rescaling, which is what carries the zeros through. The last has the fifth defect, what goes wrong in a law that keeps the zeros faithfully and rescales some other matrix, and it establishes that the matrix has to be $|U|^2$ itself.

6.8 Divisibility and realism

The four defects of resampling are one failure seen four times. A gate's unistochastic matrix holds zeros wherever the gate forbids a transition, and resampling writes over them: every column becomes the next moment's Born distribution, which is positive wherever that distribution is, so every transition the gate forbids acquires a probability. Each defect below is that one substitution, met in a circuit where it shows.

Keeping the zeros is the whole of the repair. A zero says that a transition does not happen, and that is the entire content of each defect below. Where the gate is deterministic, as in the first two circuits, the zeros settle the law outright: they leave nowhere to put the probability except on the ones, and a column summing to one puts all of it there, so the matrix is forced to remain itself. In the third circuit they settle what the defect is about, that the distant bit does not move, and leave the fitting of Section 6.6 to set the entries. Where a gate has no zeros at all, as after a Hadamard, the fitting does all of the work.

The first defect of resampling shows up after a single gate. Send a bit through a Hadamard and then leave it alone:

$$\mathbf{e}_0 \quad \text{H} \quad \text{I} \quad \text{I} \quad \text{I}$$

Every gate after the first is the identity, the distribution sits at (0.5, 0.5) and stays there, and resampling has the bit changing value half the time at every step, without end. Read the bit as which arm of the interferometer of Section 8.2 the photon occupies, and the picture is a photon crossing between the arms and back while nothing acts on it. Nature may work

that way, and no experiment we know of would say. What can be said is that nothing in the physics asks for it: the churn is put there by the law, at every idle moment of every system.

The zeros remove this defect by themselves. A deterministic unistochastic matrix carries a single one in each column and zeros everywhere else, and keeping those zeros leaves nowhere to put the probability except on the one, where a column summing to one puts all of it. Such a matrix is therefore forced to remain itself. The identity's is zero off the diagonal, so the fitted law is the identity, and the bit keeps its value.

The second defect is the same one, in a circuit where the deterministic matrix is not the identity. A permutation never interferes, and its unistochastic matrix carries nothing but zeros and ones, which is to say it reports that the bit went somewhere definite. Follow the Hadamard with an X:

$$\mathbf{e}_0 \quad \text{H} \quad \text{X}$$

The distribution stands at (0.5, 0.5) before the X and at (0.5, 0.5) after, so nothing in the statistics moves at all. What happens underneath is not nothing. Here $|\text{X}|^2$ is the swap $\begin{pmatrix} 0 & 1 \\ 1 & 0 \end{pmatrix}$, whose zeros sit on the diagonal this time, and by the same reckoning the fitted law is that swap again: the bit flips, on every run, which is what an X does. Resampling hands both columns the distribution (0.5, 0.5), and the bit comes out of the gate on a coin toss.

Set that beside the idle bit above. Both circuits hold the marginals at (0.5, 0.5) from one step to the next. In the first the bit never moves, in the second it moves every time, and resampling returns the same coin for both. A law that cannot tell *always* from *never* is getting the statistics right by getting every run wrong. What tells the two apart is where the zeros fall, off the diagonal in one case and on it in the other, and resampling has kept neither.

The third defect shows up where Chapter 11 will care about it most. Quantum mechanics forces non-local laws, and there is a narrow way for a law to be non-local and a broad one. Under the law of Section 6.6, a gate acting on one bit of an entangled pair never changes the value of the distant bit, although what the near bit does depends on that distant value. Call that *pseudo-locality*: a dependence at a distance without a disturbance at a distance.

Put both bits in a statistical state then turn one of them:

$$\begin{array}{ccc} \mathbf{e}_0 & \text{H} & \text{R}_Y(\pi/3) \\ \mathbf{e}_0 & \text{H} & \text{I} \end{array}$$

The gate leaves the second bit alone, and its unistochastic matrix says so in its zeros, which sit at every entry that would change that bit. The fitted matrix keeps them exactly where they were:

$$|\text{R}_Y(\pi/3) \otimes \text{I}|^2 = \begin{pmatrix} 0.75 & 0 & 0.25 & 0 \\ 0 & 0.75 & 0 & 0.25 \\ 0.25 & 0 & 0.75 & 0 \\ 0 & 0.25 & 0 & 0.75 \end{pmatrix} \quad \Gamma = \begin{pmatrix} 0.119 & 0 & 0.015 & 0 \\ 0 & 0.119 & 0 & 0.015 \\ 0.881 & 0 & 0.985 & 0 \\ 0 & 0.881 & 0 & 0.985 \end{pmatrix}$$

The entries move, the zeros do not, because the fitting rescales rows and columns and a rescaled zero is zero. The second bit therefore never changes value, while the first changes on 45 per cent of runs. Under resampling every column is the same distribution (0.033, 0.033, 0.467, 0.467), no entry is zero anywhere, and both bits change on half the runs: a dependence becomes a disturbance, and the non-locality spreads from the laws into the values.

The fourth defect needs a circuit in which the first three compound. A second bit records the first, and every gate after that acts on the first bit alone:

$$\begin{array}{ccccccc} \mathbf{e}_0 & \text{H} & \overline{\text{C}}\text{X} & \text{H} & \text{H} & & \\ \mathbf{e}_0 & & \text{C}\overline{\text{X}} & & & & \end{array}$$

Number the gates in order, so that gate 1 is the Hadamard, gate 2 the recording, and gates 3 and 4 the Hadamards that follow; let moment 0 be the configuration before any gate acts and moment t the configuration just after gate t . Write x_t for the first bit's value at moment t and y_t for the second bit's. The recording sets $y_2 = x_2$, and the end of the circuit is moment 4. The question below is whether $x_4 = x_2$.

The recording makes a base case on the first bit's marginal, and the two Hadamards that follow undo one another, since $\text{H}^2 = \text{I}$. The state at the end is the state the recording left. Under the law of Section 6.6 the configuration comes back as well: the Hadamards leave the recording bit alone, by the zeros of the third defect, and with that bit held fast the pair has nowhere to return to but the configuration it set out from. A pair at 00 is at 00 again and a pair at 11 at 11, both with certainty. Under resampling the pair ends at 00 or 11 on even chances, and the recording bit differs from what it recorded on half the runs. Asked about the measured bit alone, the question is whether the value it has just after the recording is the value it has at the end. Under the law of Section 6.6 it is, on every run. Under resampling it is on half of them.

Read that second bit as the memory of an observer inside the laboratory, holding what the first bit was found to be, and let a second observer describe the laboratory from outside. Both are entitled to an account of the two Hadamards, and resampling gives them different ones.

The inside observer measured and found some value, which we call S ; it is 0 on about half the runs and 1 on the rest, and nothing below depends on which. By Section 6.3 the first bit's marginal restarts at $\mathbf{e}_S$, so on that observer's description the bit runs $\mathbf{e}_S$, then $(0.5, 0.5)$ after gate 3, then $\mathbf{e}_S$ again after gate 4, since $\text{H}^2 = \text{I}$. Resampling fills every column of a step with the distribution that step arrives at, so both columns of the last step are $\mathbf{e}_S$. The bit is back at S with certainty, whatever it did in between, and the argument is the same for either value S took. The inside observer's prediction is therefore

$$x_4 = x_2 \quad \text{on every run,}$$

which is a statement about the whole ensemble and not about the runs carrying one particular outcome.

The outside observer describes the two bits together, and on that description no base case has appeared at all: Section 6.3 gives them only on marginals, and the joint keeps every phase relation. The pair runs $(0.5, 0, 0, 0.5)$ over the configurations 00, 01, 10, 11, then $(0.25, 0.25, 0.25, 0.25)$, then $(0.5, 0, 0, 0.5)$. Resampling fills every column of the last step with $(0.5, 0, 0, 0.5)$, so the pair arrives at 00 or at 11 on even chances wherever it was beforehand, which gives

$$\Pr(x_4 \neq x_2) = \frac{1}{2}.$$

The law of Section 6.6 gives $\Pr(x_4 \neq x_2) = 0$ instead, by the zeros that hold the second bit fast.

The two statements are about one ensemble of runs, and each is the denial of the other. The inside observer has x_4 and x_2 equal in every run there is; the outside observer, using resampling, has them unequal in half of them. Leaving S unspecified is what makes this plain. Fix S at 0 and the disagreement can be mistaken for a claim about one run, or about the sub-ensemble with that outcome, which an appeal to ignorance might cover. Left general, the inside observer's claim is an identity holding run by run, and a proportion of one half contradicts it outright.

One law, two observers, two answers to the same question about the same bit: equal always, and equal half the time. Under a law that keeps the zeros the accounts agree instead, and agree for any sequence of gates on the first bit, since the zeros of the third defect hold the second bit fast and the inside observer's description is then exactly the part of the outside observer's that the record picks out.

The reply from outside is that the inside observer is misreporting: their memory was redrawn along with everything else, so the value they now name is the value the bit now has. That is true, and it settles nothing, because nothing available from the inside shows anything wrong. They measured, they reasoned correctly from a rule the framework supplies, and their prediction was confirmed by every check they could make. What was theirs to check, they checked. A law on which an observer's account of their own measurement and an outside account of that same measurement come apart, with nothing available to either of them that could decide between the two, has made what happened depend on who is describing it. No thought experiment of this kind ends in a measurement anyone can perform, and that is not the complaint. The complaint is that objective fact has gone.

And objective fact was the point of the exercise. A statistical reading of quantum theory is taken up in order to have systems that are in some definite configuration whatever anyone knows or describes, which is the whole of what it offers over an account built on the state vector alone. A transition law that hands back a description depending on who holds it gives that up while keeping every part of the machinery, which leaves nothing to recommend the program over the readings it set out to improve on.

Stacking records does not help. A second record written when the first was is redrawn along with it, so the two agree whenever they are compared and the check passes every time. Section 8.12 takes up a related objection of Albert's, that a world filled in this way holds records which agree at every instant and track nothing.

The demon also brings the two accounts of the laboratory into contact. Take the circuit in its general form: a bit is randomized with He_0 , a second bit records it, and the first bit is then acted on by gates composing to the identity, the two Hadamards above being one case. The inside observer's account has the first bit back at $\mathbf{e}_S$ at the end with certainty, and the outside observer's resampling account has it there on half the runs. Those are not two views of one fact. Sort the demon's buckets by the value the record held at the moment it was written, and the bucket for S holds a definite proportion of runs that ended at $\mathbf{e}_S$. That proportion is one number. Either the bit came back every time or it came back half the time.

Over many runs of the experiment there has to be at least one on which the outside account has the bit at $\mathbf{e}_S$ when it leaves moment 2 and at the other value when it arrives at moment 4, since otherwise the proportion would not come out as it does. Take that run. On it the outside account has $x_4 \neq x_2$, while the inside observer, reading their record and their bit at the end and finding the two in agreement, has $x_4 = x_2$. Both have the same x_4 , since

the record and the bit always agree at the end. So on that one run the two accounts disagree about x_2 :

$$x_2(\text{inside}) \neq x_2(\text{outside}).$$

This is not an empirical contradiction. The inside observer reports $x_4 = x_2$ on every run, and the outside account has $x_4 \neq x_2$ on half of them without being able to say which: a distribution over the pair names no run in particular. Nothing empirical distinguishes the runs on which the two come apart from the runs on which they agree. One has to exist regardless. If there is an observer-independent reality, that is a contradiction: the bit had one value at moment 2, and the two accounts assign it two. Resampling therefore leaves the configuration of the system observer-dependent.

The composed transition matrix for the interval says it in one line. Restricted to the two configurations the recording leaves, that matrix is the identity under the law of Section 6.6, while under resampling every one of its columns is the Born distribution $(0.5, 0, 0, 0.5)$. One of those is a deterministic transition and the other is as far from deterministic as a transition on that support can be, and which of them a stretch of the world performed is not settled by who describes it. The inside observer cannot locate the error because the runs have been relabeled: the set whose record read S when it was written is no longer the set whose record reads S at the end, the record having moved with the bit, and only the demon still holds the first set. The prediction was about that set, and of that set it is false.

Albert raises the unmodeled interval as an objection to the whole program, and independently of anything above.⁴ A theory that fixes the probability of each configuration and says nothing about the path between them is, he argues, incomplete in a way that matters. Fill the gap by drawing each moment independently, which is resampling, and the records a world contains agree with one another at every instant while indicating nothing whatever about the past. Nothing then has reliable dispositions, and he doubts such a world contains observers or measurements at all. He offers two completions of his own and calls the better of them a preposterous way to build a world.

The construction of Section 6.6 is a completion, and the fourth defect above has already shown the records holding under it. A gate that leaves the recording bit alone has $|U|^2$ equal to zero on every entry that would change it, and zeros survive the fitting, so the probability that a record changes is not merely small but exactly zero, against one half under resampling. What keeps the photon on its arm and what keeps a record readable are the same property of the law.

Albert argues this through a particular experiment, a particle passing two slits with a pointer recording which one it took. The double slit has not been introduced yet; it arrives in Chapter 8, and Section 8.12 runs that experiment under this law. The record holds, the particle keeps to a path, and the interference the recording ought to destroy is destroyed.

6.9 Divisibility and minimality

The quantity minimized in Section 6.6 was displayed and never argued for. Its reason is Schrödinger's, and it brings more with it than the quantity: minimizing it comes out as a rescaling of rows and columns, so the zeros of Section 6.8 survive without anyone having to

⁴D. Z. Albert, "Quantum mechanics without a wave-function," *Synthese* 208 (2026), article 141.

require them to, and the one thing left open is which matrix to scale from, which Section 6.10 settles.

The reason was given in 1931, for diffusing particles rather than circuits, and Schrödinger set it up in a form we have already met. “An auxiliary observer has observed the position of the particle at time t without, however, reporting us the result,” he writes, and asks what can be inferred about what that assistant saw. The assistant is the demon of Section 6.5, watching a moment nobody else is shown. Schrödinger then let a large number of particles start with one distribution and observed them arrive with another that the diffusion law on its own would not have produced, remarked in an aside that such an observation “may be more or less surprising,” and asked how they most probably got there. Counting the ways gave him a pair of integral equations that a rescaling of rows and columns solves.⁵ His question carries over to a circuit unaltered, and so does the surprise, which here is interference. A circuit running on $|U|^2$ alone is the Markov chain of Section 5.6 with the correction struck out. Wherever that correction is not zero, the distribution the chain hands the next moment is not the one the Born rule gives. The rest of this section is his argument run on a gate.

Keep two things apart before starting, because the argument runs on the difference between them. A *law* is what governs a step. A *tally* is what the demon of Section 6.5 writes down after watching many runs of it: so many went in at 0 and came out 0, so many went in at 0 and came out 1, and so on. A tally is a record of frequencies, and it is not yet anybody’s law. Dividing it by its column totals gives a matrix of the right shape, but the matrix so obtained is an observed frequency, which may sit anywhere its law allows and not only where that law is likeliest to put it. One cannot ask how probable a law is. One can ask how probable a tally is, given a law, and that is the only probability in what follows.

So suppose, for the length of the argument, that the gate’s own matrix $|U|^2$ is the law. Each run then draws its outgoing value from the column of $|U|^2$ belonging to the value it carried in, and every tally has a probability: the chance of each run doing what the tally says it did, multiplied by the number of ways the runs could have been assigned to produce that tally. Almost all of that probability sits on tallies close to $|U|^2$ itself, which is the law of large numbers. The rest is spread over tallies that depart from it, and those are the ones of interest, since the tally we are after departs from it.

It departs because of interference, and the departure is always within reach. Interference can only cancel. A configuration’s outgoing Born probability is the squared magnitude of a sum, with one term for each path leading into it. Where the gate allows no path, there is nothing in the sum, and the probability is zero. No arrangement of phases can put weight on an outcome the gate forbids. Cancellation the other way happens constantly: the second Hadamard of Chapter 5 drives an outcome of probability one half down to none at all. So the Born distribution at the next moment is always a distribution the $|U|^2$ law could have thrown up, improbably, by itself. It is the zeros of Section 6.8 again, now as a limit on which distributions can turn up, and it is why the fitting of Section 6.6 never has an empty problem to solve.

Now ask which tally it most probably was. Restrict attention to the tallies whose two

⁵E. Schrödinger, “Über die Umkehrung der Naturgesetze,” *Sitzungsberichte der Preussischen Akademie der Wissenschaften, Physikalisch-mathematische Klasse* 8 N9 (1931), 144–153, translated by R. Chetrite, P. Muratore-Ginanneschi and K. Schwieger as “On the Reversal of the Laws of Nature,” *European Physical Journal H* 46 (2021), article 28. The counting is his main calculation; its large-deviation form, which is the quantity of Section 6.6 exactly, is set out by the translators in their §1.2.

margins are the observed distributions, count each one's ways as above, and for a large number of runs one of them beats every other by a factor growing exponentially with the number of runs. Its exponent is exactly the quantity displayed in Section 6.6. So the phrase “deviates no more than the Born marginals force” has something plain behind it: among the tallies that show the observed distribution at both ends of a step, this is the one the gate's own matrix renders least improbable, and the others are rarer by a factor that grows without limit.

The counting says what shape the answer has as well, and this is where the rescaling comes from. A tally is a table of counts with one cell for each pair of values, and both of its margins are settled before anything is chosen: the column of runs that went in at j holds $Np_j^{[t]}$ of them, because that is how many went in at j , and the row of runs that came out at i holds $Np_i^{[t+1]}$, because that is how many came out at i . The question is which tally with those margins is the most likely one.

So ask how one tally can differ from another without either margin differing. Cell (i, j) holds the runs that went in at j and came out at i , so to shift a run from that cell to (i', j) is to suppose one of them came out at i' instead. One such supposition on its own leaves a gap, since the row for i is now short and the row for i' long. The smallest change that keeps every total where it was takes two runs and exchanges them across a rectangle: one run that came out at i now comes out at i' , and one run that went in at j' makes the opposite swap, coming out at i where it had come out at i' . Both columns hold the counts they held, both rows hold theirs, and four cells have changed.

Then ask when no exchange of that kind leaves a more probable tally than the one it started from. Two things pull against each other. The two runs now make different transitions, so the gate's own probabilities favor the exchange whenever the pair of transitions it moves to is the likelier pair. Against that, the counts have shifted, and with them the number of ways the runs can be distributed to produce the tally: a run taken out of a crowded cell and put into a sparse one leaves more ways than before, so that factor favors spreading the runs out. At the most likely tally the two pulls balance in every rectangle, and writing that balance down gives

$$\frac{\Gamma_{ij} \Gamma_{i'j'}}{\Gamma_{i'j} \Gamma_{ij'}} = \frac{|U|_{ij}^2 |U|_{i'j'}^2}{|U|_{i'j}^2 |U|_{ij}^2}$$

(the run-counts give the same ratios, each column being the matrix's column scaled by the same $p_j^{[t]}$). Each side is a ratio of ratios: on the left, how much likelier the step is to send j to i than to i' , divided by how much likelier it is to send j' to i than to i' . The condition says the tally and the gate agree on that comparison for every pair of inputs and every pair of outputs. They need not agree on the transition probabilities themselves, which the margins have already pushed around; they agree on how the four in any rectangle stand relative to one another. The margins say nothing whatever about those relative standings, which is why they are what the gate is left to fix.

That condition is exactly what a rescaling satisfies. Take any positive number for each row and each column, multiply row i by the first and column j by the second, and every one of these ratios comes out as it went in, each multiplier appearing once above the line and once below. So a rescaling of the gate's matrix meets the condition, and what remains is to see that nothing else does. Write M_{ij} for the ratio of Γ_{ij} to $|U|_{ij}^2$, so that the condition reads $M_{ij} M_{i'j'} = M_{i'j} M_{ij'}$. Fix any row i_0 and any column j_0 , take $i' = i_0$ and $j' = j_0$, and it

becomes $M_{ij}M_{i_0j_0} = M_{ij_0}M_{i_0j}$, so that

$$M_{ij} = \frac{M_{ij_0}}{M_{i_0j_0}} \cdot M_{i_0j},$$

a factor belonging to the row times a factor belonging to the column. Call the first u_i and the second v_j , both fixed once the reference row and column are chosen, so that $M_{ij} = u_i v_j$. Since M_{ij} was the ratio of Γ_{ij} to $|U|_{ij}^2$, multiplying back through gives

$$\Gamma_{ij} = u_i |U|_{ij}^2 v_j \quad \text{for every } i \text{ and } j,$$

and a matrix whose every entry reads that way is written

$$\Gamma = \text{diag}(\mathbf{u}) |U|^2 \text{diag}(\mathbf{v})$$

which is the form Section 6.6 fits, and the one the next section reaches by a different route. So the rescaling is the shape the most likely tally has, and the counting is what shows it. The zeros then look after themselves. A rescaled zero is zero, which is everything the four defects of Section 6.8 asked for, and nothing had to be added to the minimization to get it.

The last step is ours and not the counting's. The counting ranks tallies under a law we have supposed for the sake of it; adopting the winner as the law itself is a further move. What recommends the move is an accounting principle: credit the gate with what the gate does, and credit the history with no more than the marginals force. The matrix $|U|^2$ is the gate's own contribution, and the tally that departs from it least is the one leaving the least for the history to have done.

Schrödinger defended a step of the same kind, and in the same spirit:⁶

If we were asked our opinion about the previous history that the observed strongly abnormal distribution could have probably had, then we would have to reply . . . that the anomaly was probably caused by a precise time reversal of a normal diffusion process. Admittedly, this statement about the likely previous history would be only a probabilistic judgment, nevertheless it should, in my opinion, be granted the same degree of “almost certainty” as the corresponding statement about the likely future evolution.

A retrodiction of the likeliest history, on that accounting, stands exactly where a prediction of the likeliest future stands, and we take the interval the same way.

He never solved the problem he had posed. The 1931 paper has the answer in it already and wants only a way to compute the two factors, and that he did not have. He wrote that discussing his pair of equations “would be certainly interesting but probably not simple because it is nonlinear,” took their solvability for granted “because of the reasonable question which in an unambiguous and sharp manner leads to these equations,” and said outright that constructing the factors from the observed distributions was not what concerned him.

Twenty years later he still had no way to do it, and wave mechanics (his own theory) had failed to deliver the thing he wanted all along: *a description of physical reality with no gaps*

⁶Schrödinger, “Über die Umkehrung der Naturgesetze,” §6, in the translation cited above; the auxiliary observer is posed in his §1. The abnormal distribution he speaks of is one the diffusion law would rarely produce, which is the place interference occupies here.

in it. In the lectures where he set out what it had achieved, he is exact about what it supplies and what it withholds:⁷

We do give a complete description, continuous in space and time without leaving any gaps, conforming to the classical ideal... But we do not claim that this ‘Something’ is the observed or observable facts; and still less do we claim that we thus describe what nature (matter, radiation, etc.) really is. In fact we use this picture (the so-called wave picture) in full knowledge that it is neither... The gaps, eliminated from the wave picture, have withdrawn to the connection between the wave picture and the observable facts.

The wave evolves continuously under its equation, so that evolution has no gaps in it. A quantum state resembles nothing anybody observes, so the gaps reappear where the evolution has to meet what is seen. His own construction of 1931 is addressed to exactly those gaps, since what it computes is what most probably happened between two observations. He had written down what would fill them and could not carry it out.

The method was developed independently outside of physics. In 1937 a Dutch telephone engineer, J. Kruithof, had to forecast the traffic between exchanges from the traffic they already carried and the new totals leaving and arriving, and he solved it by giving every row a factor and every column a factor and repeating until the totals came right. Deming and Stephan put the same procedure to census tables three years later and gave it the name it still has.⁸

What the method leaves open is which matrix to scale from, since it is done against that matrix rather than deriving it, and the next section closes that question on other grounds.

6.10 Divisibility and counterfactuals

The four defects of Section 6.8 belong to resampling. A fifth belongs to a second rival, which keeps the fitting of Section 6.6 and changes only the matrix it begins from, and what goes wrong there takes a circuit of a different shape to see. The demon’s buckets test more when the incoming distribution is not degenerate, and the rotation of Section 6.9 will serve again. Preparing such an input takes a gate like anything else, so let someone outside the experiment take a bit with an initial value of 0 and act on it with X^n , choosing $n = 0$ or $n = 1$ on a coin flip, and hand what comes out to a single rotation:

$$\mathbf{e}_0 \quad X^n \quad R_Y(\pi/3)$$

Two things hold here that did not hold for a bit read and then sent through a gate. What reaches the rotation is uniform rather than degenerate, so both columns of the matrix are

⁷Schrödinger, *Science and Humanism*, pp. 144–145. The remarks on the integral equations in the paragraph above are from §3 of the 1931 paper.

⁸J. Kruithof, “Telefoonverkeersrekening,” *De Ingenieur* 52 (1937), E15–E25; W. E. Deming and F. F. Stephan, “On a Least Squares Adjustment of a Sampled Frequency Table When the Expected Marginal Totals are Known,” *Annals of Mathematical Statistics* 11 (1940), 427–444. The matrix-scaling literature that grew from the same root is entered most easily through R. Sinkhorn and P. Knopp, “Concerning nonnegative matrices and doubly stochastic matrices,” *Pacific Journal of Mathematics* 21 (1967), 343–348, whose concern is the case where the matrix being scaled has zeros in it, as ours does.

exercised at once. And what reaches it is a statistical distribution over the two eigenstates, whose density matrix is the diagonal $\frac{1}{2}\mathbf{I}$ of Section 4.3, so there is nothing off the diagonal to interfere: the outgoing distribution is (0.5, 0.5) again, and $|\mathbf{R}_Y(\pi/3)|^2$ transports it correctly on its own.

With $\cos(\pi/6) = \sqrt{3}/2$ and $\sin(\pi/6) = 1/2$, the rotation and its unistochastic matrix are

$$\mathbf{R}_Y(\pi/3) = \begin{pmatrix} \sqrt{3}/2 & -1/2 \\ 1/2 & \sqrt{3}/2 \end{pmatrix}, \quad |\mathbf{R}_Y(\pi/3)|^2 = \begin{pmatrix} 3/4 & 1/4 \\ 1/4 & 3/4 \end{pmatrix}.$$

Whoever flipped the coin knows which n each run got, so the runs can be sorted by the value that went in without anyone intervening. The bucket that began at 0 comes out (3/4, 1/4) and the bucket that began at 1 comes out (1/4, 3/4): both columns of $|\mathbf{R}_Y(\pi/3)|^2$, and counting the runs is all it took to get them.

Chapter 5 had already derived that matrix, before any question of fitting arose. The polar decomposition of Section 5.6 wrote the evolution of the probabilities as a Markov chain carrying $|U|^2$ as its matrix, plus a correction that is not linear in $\mathbf{p}$. Every term of the correction carries $\sqrt{p_k p_l}$ for two distinct configurations, so a distribution sitting on a single value multiplies all of it by zero and the Markov chain is the whole of the rule. Each branch of the coin flip arrives that way, degenerate at 0 on the runs with $n = 0$ and degenerate at 1 on the runs with $n = 1$. So across that step the exact rule is a Markov chain whose matrix is $|U|^2$.

The buckets therefore confirm a matrix the book already had in hand. A law that begins from a different one contradicts that derivation, and the rest of this section follows one that does: its statistics come out right at every moment, and its counterfactuals cannot all be true at once.

With the tally in hand, the two laws can be taken in turn. Section 6.6 starts from $|\mathbf{R}_Y(\pi/3)|^2$, which already transports, so the scaling vectors come back all ones and the answer is that matrix untouched. It agrees with the buckets.

Scott Aaronson's *Schrödinger theory* fits the same way, by rescaling rows and columns until the marginals come right, and starts from $|U|$ in place of $|U|^2$, so here the scaling begins from $|\mathbf{R}_Y(\pi/3)|$. It takes its name from the construction of Section 6.9, older than this use of it; scaling $|U|$ is Aaronson's own step.⁹

Here the arithmetic is short. The entries are $\sqrt{3}/2$ and $1/2$, each column sums to $(\sqrt{3}+1)/2$, the marginals wanted are already uniform, and the fitting has nothing left to do but divide:

$$\frac{2}{\sqrt{3}+1} \begin{pmatrix} \sqrt{3}/2 & 1/2 \\ 1/2 & \sqrt{3}/2 \end{pmatrix} = \begin{pmatrix} 0.634 & 0.366 \\ 0.366 & 0.634 \end{pmatrix}.$$

That matrix carries (0.5, 0.5) to (0.5, 0.5) exactly, so every statistic the experimenter can compute comes out right. The buckets hold 3/4 where it says 0.634.

The complaint here is milder than the fourth defect's. Nothing has become observer-dependent: the tally settles the matter and one of the two answers is wrong. What the example shows is that a law beginning from $|U|$ recovers the transitions only for someone who already knows how every run started. Anyone holding the ensemble description alone gets a

⁹S. Aaronson, "Quantum Computing and Hidden Variables," *Physical Review A* 71 (2005), 032325, §V B; the scaling is his equations (2) and (3).

different matrix for the same step, so what such a law says about the rotation depends on how much its user knew.

There is a deeper issue still, and the demon of Section 6.5 is what draws it out. The demon does not see gates. It sees a value go in and a value come out, and it sorts the runs by the value that went in, so the runs of a step fall into one bucket for each value the step receives. Every column of the matrix those buckets exhibit is a tally: the column for an input of 0 is read off the bucket that went in at 0, and the column for an input of 1 off the bucket that went in at 1. Hand a law one bucket on its own, though, and it hands back a whole matrix. One column of it is exercised, every run in that bucket having gone in at the same value, and the rest are counterfactuals: claims about what the step would have done with a value that bucket never delivered, and so claims about the buckets the law was never shown.

Under the law of Section 6.6 those claims hold up. Fitted to the runs that went in at 0, it returns $|R_Y(\pi/3)|^2$ completely, and the column it offers for an input of 1 is $(1/4, 3/4)$, which is what the other bucket holds. Fitted to the runs that went in at 1 it returns the same matrix again, and the column it offers for an input of 0 is $(3/4, 1/4)$, which is what the first bucket holds. One matrix covers the step, and every claim in it can be checked against the bucket that has that input.

Scaling $|R_Y(\pi/3)|$ gives two matrices instead, the first fitted to the runs that went in at 0 and the second to the runs that went in at 1:

$$\begin{pmatrix} 3/4 & 1/2 \\ 1/4 & 1/2 \end{pmatrix}, \quad \begin{pmatrix} 1/2 & 1/4 \\ 1/2 & 3/4 \end{pmatrix}.$$

Each matrix agrees with its own bucket and contradicts the other's. The agreement is forced: every run in a bucket went in at the same value, so the transport requirement by itself fixes that bucket's column, whatever matrix the law began from. The contradiction is the trouble. The first matrix says that a run going in at 1 would come out $(1/2, 1/2)$, and the bucket of runs that do go in at 1 holds $(1/4, 3/4)$; those runs fall under the second matrix, whose spare column makes the same claim in reverse about an input of 0. Neither claim is ever met, since arriving at the input it describes replaces the matrix that made it.

Two of the four columns are therefore unreachable. The claim about an input of 1 holds only while the input is 0, and the claim about an input of 0 only while the input is 1, so each is retracted by the very circumstance that would test it. A claim no circumstance can meet does no work, and striking both leaves one column from each matrix, the column its own bucket exercises. Those two columns are $(3/4, 1/4)$ and $(1/4, 3/4)$, which is $|R_Y(\pi/3)|^2$. Beginning from $|U|$ and discarding what can never be reached returns the matrix one would have begun from anyway.

Those columns are the ones the scaling leaves undefined: a column whose marginal is zero sums to zero, and dividing to get transition probabilities gives $0/0$. Aaronson settles it with a limit, mixing a vanishing amount of the maximally mixed state into the input, and that limit is where the two matrices above come from. A later repair of the theory writes $|U|^2$ into those columns instead, which removes the disagreement.¹⁰ That repair is a case split, one rule for

¹⁰Aaronson, §II, for the limit. The repair is G. Malavolta, H. Nieuwboer, A. Ramachandran and M. Walter, “Robustness of Hidden-Variable Theories and Matrix Scaling,” arXiv:2609.14519 (2026), §1.3, which also refutes Aaronson’s conjecture that his theory is robust. The modification was made for that reason rather than this one, and its authors give the choice as one reasonable option among others.

the columns carrying probability and another for the columns without. Scaling $|U|^2$ needs no such split: where the input is definite it returns $|U|^2$ by the rule it applies everywhere else.

Put it in terms of what the demon has to do. It manipulates nothing, so to learn a step's matrix it has to go looking: find the instances of that step scattered through the world, sort them by the value that went in, and tally each group on its own. For a one-bit gate that is two searches and two tallies, the runs that arrived at 0 and the runs that arrived at 1, each divided by how many there were. Two columns, one matrix. Chapter 5 says what the matrix will be wherever nothing interferes, and the demon's tallies say the same thing: $|U|^2$.

That matrix describes the step only if the law was the same for both groups. Under $|U|^2$ it was, and the demon's two columns are two columns of one law. Under $|U|$ the groups were governed by the two matrices above, and the demon still writes $|U|^2$ down, a column from each. The two laws disagree about both columns: where the first reads $(3/4, 1/4)$ and $(1/2, 1/2)$, the second reads $(1/2, 1/2)$ and $(1/4, 3/4)$. From each the demon takes the column that law's own bucket exercised, so every column it writes down comes from a law the other column's law contradicts. The matrix is spliced from two laws and is the law of neither. What the demon never gets to tally is the pair of columns where those laws disagreed, and those stay out of reach for the demon as much as for anyone.

The hiding is doing necessary work. The transport requirement fixes each bucket's own column whatever matrix a law begins from, so every way a law can depart from $|U|^2$ has to sit in the columns its bucket does not exercise. That is why scaling $|U|$ predicts everything $|U|^2$ predicts: its whole deviation is parked where nothing can be counted, and it conspires to have the demon tally $|U|^2$ after all. A deviation anywhere else would turn up in a bucket and be refuted on sight. Beginning from $|U|^2$ leaves nothing parked anywhere. Each bucket's spare columns are the columns the other buckets hold, and what the demon reads off the world is what the world did.

This is not the history dependence of Chapter 5 wearing a different coat. A reader may object that the second Hadamard of a pair already behaves differently from the first, so a gate's matrix depends on what came before it there too. It does. The difference is that the second Hadamard receives 0 on half the runs and 1 on the other half, so both of its columns are tallied from runs that happened, and the matrix fitted to either half alone is the matrix fitted to all of them. Nothing is spliced and nothing lies out of reach. That dependence is on a history the runs share, which the demon watches like anything else, and not on which value a given run carried in.

The agreement is no feature of this one circuit. A step that receives a definite value is a base case, so the correction of Section 5.6 vanishes term by term and the outgoing distribution is exactly the column of $|U|^2$ belonging to that value. The row multipliers in the fit come back all ones, and the fitted matrix is $|U|^2$ completely, whichever value the step received. Every instance of a step with a definite input therefore writes down the same matrix, and no column of it can be contradicted by another instance.

That settles a question the Hadamards raise. Send a bit at 0 through three of them. The first two return it to 0, so the bit reaching the third carries the value 0 on every run and the third's other column is never tallied in that circuit at all. Start the same circuit at 1 and the first two return it to 1, so the bit reaching the third carries a 1 every time, and the column the first arrangement never reached is the one this arrangement tallies. Both fit the third Hadamard to the same matrix.

So every column is within the demon's reach, and it needs no intervention to get there. One that had only ever watched circuits started at 0, and wanted to know how the third Hadamard treats an incoming 1, goes looking through the world for instances of the same arrangement started at 1 and tallies those instead. The column it could not fill from the first set of runs is one the second set fills, and by the paragraph above the two agree. A counterfactual the demon cannot reach would have to be a column no arrangement anywhere exercises, and nothing depends on what a law says about one of those.

So we add a requirement to the ones already imposed: a law may assign a step no counterfactual that cannot be reached. The runs that arrived at one value make up one bucket, gathered from wherever in the world they happened, and the buckets between them exhibit one matrix, a column apiece. A law handed any single bucket must return that whole matrix, so that each of its counterfactual columns says what the bucket holding that value shows. Those buckets are there to be read.

Kruithof asked something of the same shape of his own method, and his reason for it carries over. Forecasting traffic from one year to a later one, he required the answer to come out the same whether he went straight there or worked through some intermediate year on the way, since otherwise his forecast would depend upon a date he had picked for no reason. The requirement above has that motive pointed at a different arbitrary choice, which is the bucket a law happens to be handed. His own condition the fitting meets exactly: fitting twice is itself a rescaling, and the rescaling that meets given marginals is unique, so the detour through the intermediate year changes nothing.

That requirement settles what a law of this form may begin from. The scaling form was argued for in Section 6.9; the question at hand is only which matrix the scaling starts from. Write A for the matrix a law rescales, Γ for the matrix the buckets exhibit, and $\mathbf{u}$ and $\mathbf{v}$ for the row and column multipliers of Section 6.6.

Fitting A to the bucket whose input is degenerate at j leaves only column j to answer for, and the transport requirement forces that column to come out as the tally:

$$u_i A_{ij} v_j = \Gamma_{ij}.$$

That pins each u_i to Γ_{ij}/A_{ij} apart from the single shared factor v_j , which can be left where it stands. Every remaining column is divided by its own sum in a moment, and a factor common to all of them cancels there.

Those remaining columns carry no transport constraint, only the requirement that each sum to one, so column k comes out as $u_i A_{ik} v_k$ with $v_k = 1/\sum_m u_m A_{mk}$. The requirement is that each of them be the tally as well, which is to say that $\Gamma_{ik} = u_i A_{ik} v_k$ for every i and k , or

$$\Gamma = \text{diag}(\mathbf{u}) A \text{diag}(\mathbf{v}).$$

So a law meets the requirement only when the matrix it rescales is the matrix the buckets exhibit, give or take a rescaling of that matrix by rows and columns. Such a rescaling changes nothing any fit can see, since rescaling A and then fitting returns what fitting A returns, so the requirement fixes the starting matrix outright.

Where nothing interferes, that matrix is $|U|^2$, a definite value j going in coming out distributed as $|U|_{ij}^2$. Scaling $|U|$ meets the requirement only for gates whose moduli factor into a row term times a column term, which the Hadamard and the permutations do and

a rotation through a general angle does not. The defect stays out of sight until a gate like $R_Y(\pi/3)$ is put in front of it.

Aaronson records the failure himself, as a violation of what he calls decomposition invariance. His axiom asks for more than the requirement above. It ranges over every way of writing a statistical state as a mixture of pure ones, including ways that sort no runs by any value they carry, and in that form it cannot be held alongside the zero-preservation of the third defect, which is a theorem of his own.¹¹

¹¹Aaronson, §II for the axiom and §III, Theorem 2 for the impossibility. The requirement above ranges only over the sorting the demon performs, which is why the theorem leaves it alone.

Code: this chapter in Octave

```
% --- iterative proportional fitting, as displayed in Section 6.6 ---

1; % script marker so Octave allows function definitions below

function G = ipf(A, pt, pt1, iters)
    n = length(pt); v = ones(n,1);
    for k = 1:iters
        u = pt1 ./ (A * (v .* pt)); % fix the rows (transport)
        v = 1 ./ (A' * u); % fix the columns (sum to 1)
    end
    G = diag(u) * A * diag(v);
end

% --- fit the H, T, H circuit to a Markov chain ---
H = (1/sqrt(2))*[1 1; 1 -1]; T = [1 0; 0 exp(1i*pi/4)];
circuit = {H, T, H};
psi = [1; 0]; p = {abs(psi).^2};
for k = 1:3
    psi = circuit{k} * psi; p{k+1} = abs(psi).^2; % Born marginals
end
for k = 1:3
    G{k} = ipf(abs(circuit{k}).^2, p{k}, p{k+1}, 500)
end
% G{1} = COIN, G{2} = identity, G{3} = [0.854 0.854; 0.146 0.146]:
% three different stochastic behaviors for the same Hadamard gate.

% --- simulate: definite bit values at every time interval ---
N = 200000; counts = zeros(4,2);
for run = 1:N
    x = 1; counts(1,x) = counts(1,x) + 1; % definite value, always
    for k = 1:3
        x = 1 + (rand() > G{k}(1,x)); % stochastic transition
        counts(k+1,x) = counts(k+1,x) + 1;
    end
end
counts / N % matches [p{1} p{2} p{3} p{4}]' at every step

% --- the interrupted Hadamard pair (H, CX, H): base-case minting ---
e0 = [1;0]; I2 = eye(2);
CX = [1 0 0 0; 0 1 0 0; 0 0 0 1; 0 0 1 0];
psi = kron(H,I2) * CX * kron(H,I2) * kron(e0,e0);
marginalA = kron(I2,[1 1]) * abs(psi).^2 % -> [0.5; 0.5], not [1; 0]

% --- Section 6.8: resampling versus a law that keeps the zeros ---
% IPF restricted to the support, since these distributions contain zeros.
```

```

function G = ipf_support(A, p0, p1, iters)
    n = length(p0); G = zeros(n);
    c = find(p0 > 1e-14); r = find(p1 > 1e-14);
    As = A(r,c); v = ones(numel(c),1);
    for k = 1:iters
        u = p1(r) ./ (As * (v .* p0(c))); v = 1 ./ (As' * u);
    end
    G(r,c) = diag(u) * As * diag(v);
    G(:,c) = G(:,c) ./ sum(G(:,c),1);
end

% the idle bit: after H, every gate is the identity and p stays at [0.5; 0.5]
pid = [0.5; 0.5];
Gid = ipf_support(eye(2), pid, pid, 400);
hop_z = Gid(2,1)*pid(1) + Gid(1,2)*pid(2) % -> 0: the bit keeps its value
hop_rs = 2*pid(1)*pid(2) % -> 0.5 under resampling

% --- Section 6.8: a deterministic gate, and the zeros that carry locality ---
% H then X: a permutation needs no repair, so the fit returns |X|^2 itself.
X = [0 1; 1 0];
pH = abs(H*e0).^2; pX = abs(X*H*e0).^2; % [0.5;0.5] -> [0.5;0.5], unmoved
ipf_support(abs(X).^2, pH, pX, 2000) % -> the swap: the bit flips, always
pX * ones(1,2) % resampling: a coin, every run
% same marginals as the idle bit above, opposite truth: never move vs always move

% H on both, then a turn on the first bit only. The default is zero on every
% entry that would move the second bit, and rescaling cannot revive a zero.
th = pi/3; RY = [cos(th/2) -sin(th/2); sin(th/2) cos(th/2)];
psi2 = kron(H,H) * kron(e0,e0); U2 = kron(RY,I2);
q1 = abs(psi2).^2; q2 = abs(U2*psi2).^2;
abs(U2).^2 % zero wherever bit 2 would change
Gz = ipf_support(abs(U2).^2, q1, q2, 4000) % same zeros, different entries
b2 = [0;1;0;1]; b1 = [0;0;1;1];
sum(sum((b2 != b2') .* (Gz .* q1')) % -> 0: bit 2 never changes
sum(sum((b1 != b1') .* (Gz .* q1')) % -> 0.448: bit 1 does
sum(sum((b2 != b2') .* ((q2*ones(1,4)) .* q1')) % -> 0.5 under resampling

% --- Section 6.8: H, CX, H, H -- the record that should survive, and does not ---
g = {kron(H,I2), CX, kron(H,I2), kron(H,I2)};
ps = kron(e0,e0); Pm = {abs(ps).^2};
for k = 1:4; ps = g{k}*ps; Pm{k+1} = abs(ps).^2; end
Pm{3}, Pm{5} % state at t=4 is the state at t=2
G3 = ipf_support(abs(g{3}).^2, Pm{3}, Pm{4}, 4000);
G4 = ipf_support(abs(g{4}).^2, Pm{4}, Pm{5}, 4000);
G4 * G3 % -> identity on 00 and 11: it returns
(Pm{5}*ones(1,4)) * (Pm{4}*ones(1,4)) % resampling: 00 and 11, half each

```

```

mem = [0;1;0;1];
sum(sum((mem != mem') .* ((G4*G3) .* Pm{3}'))))          % -> 0: memory intact
sum(sum((mem != mem') .* (((Pm{5}*ones(1,4))*(Pm{4}*ones(1,4))) .* Pm{3}')))) % -> 0.5

% the reader's own account of the same circuit: they measured, so their bit
% restarts from the recorded value, and  $H^2 = I$  brings it back with certainty
f2 = [1;0]; f3 = abs(H*f2).^2; f4 = abs(H*H*f2).^2;
Fk = ipf_support(abs(H).^2,f3,f4,3000) * ipf_support(abs(H).^2,f2,f3,3000);
Fr = (f4*ones(1,2)) * (f3*ones(1,2));
Fk(1,1), Fr(1,1)          % -> 1 and 1: the reader predicts certainty either way
Lk = G4*G3; Lr = (Pm{5}*ones(1,4))*(Pm{4}*ones(1,4));
Lk(1,1)+Lk(2,1), Lr(1,1)+Lr(2,1) % -> 1 with the zeros kept, 0.5 without: the
                                % laboratory account and the reader's part ways

% --- Section 6.10: which matrix the scaling begins from ---
th = pi/3; RY = [cos(th/2) -sin(th/2); sin(th/2) cos(th/2)];
B2 = abs(RY).^2          % -> [0.75 0.25; 0.25 0.75]
% Sorting by the value that went in gives each bucket a degenerate input, so
% only its own column is exercised. Fitting a base A to one bucket fixes the
% row multipliers from that column and normalizes the rest.
function G = fit_bucket(A, j, q)
    u = q ./ A(:,j);          % the exercised column must come out as the tally
    G = diag(u) * A; G = G ./ sum(G,1);
end

b0 = B2(:,1); b1 = B2(:,2); % the two buckets, counted and not assumed

% begin from  $|U|^2$ : both buckets return the same matrix, and it is  $|U|^2$ 
fit_bucket(B2,1,b0), fit_bucket(B2,2,b1)

% begin from  $|U|$ : each bucket gets its own column right and the other wrong,
% and the two spare columns are both [0.5; 0.5], which neither bucket holds
B1 = abs(RY);
fit_bucket(B1,1,b0), fit_bucket(B1,2,b1)

% the same two bases on the unsorted ensemble: uniform in, uniform out
pu = [0.5; 0.5];
ipf(B2, pu, pu, 500)          % ->  $|U|^2$  again, so sorting changes nothing
ipf(B1, pu, pu, 500)          % -> [0.634 0.366; 0.366 0.634]

% why no earlier circuit showed this: for a Hadamard  $|U|$  is a scalar multiple
% of  $|U|^2$ , so the two bases cannot be told apart at all
A2 = abs(H).^2; A1 = abs(H);
max(max(abs(fit_bucket(A1,1,A2(:,1)) - fit_bucket(A2,1,A2(:,1))))) % -> 0
max(max(abs(ipf(A1,pu,pu,500) - ipf(A2,pu,pu,500)))) % -> 0

```

Chapter 7

Quantum Algorithms

As discussed previously, the phase vector is a sufficient statistic on the system's past statistical state. In a hidden Markov model, this sufficient statistic would need to be included in the present, as part of an additional memory carried by the system. That additional memory is additional information the system gets to operate on. If that is true for the hidden Markov model, it is equally true for the non-Markovian dynamics that replaces it: the historical dependence of the system's dynamics allows systems to operate on more information than their present configuration holds, with the extra information stored, as it were, in the past.

Once we understand this, we can begin to answer several questions that have been accumulating since Part I. Why is classical mutual information misleading for quantum systems? Where does the extra correlation live? How can a single bit be used to transmit two bits of information? And what is it that a quantum algorithm has to work with, given that measurement only ever reveals ordinary bits? We answer all three, and the answers are one answer: *the law consults the history*.


7.1 A picture of the phase: the hypergraph

A way to visualize this additional resource will help. The phase vector has one entry per configuration, and for n bits that is 2^n entries (far too many to attach to the bits one apiece). But there is a natural home for them. Each configuration of the memory picks out a *subset* of the bits: the bits whose value is 1 in that configuration. And the subsets of n items correspond exactly to the possible *edges* on a network of n nodes, provided we count generously: the empty subset (an “empty edge,” drawn floating beside the network), the single-bit subsets (self-loops on each node), the pairs (ordinary edges), and the larger subsets (*hyperedges*, joining three or more nodes at once, which is the reason the structure is called a *hypergraph*). Both collections number 2^n , so every phase can be assigned to its own edge.


Concretely: number the entries of the phase vector from 0, and write each index in binary with n digits, B_1 leading. Each digit says whether its node takes part in the edge, 1 for yes and 0 for no. Entry 0 is all zeros, which involves no node at all, so it goes on the empty edge. An entry whose binary form holds a single 1 involves one node, so it goes on that node's self-loop. An entry of all ones involves every node at once, so it goes on the edge that joins them all, which is a hyperedge when $n > 2$.

For two bits, the four phases land as follows: m_{00} (entry 0) on the empty edge, m_{10} (entry

2) on B_1 's self-loop, m_{01} (entry 1) on B_2 's self-loop, and m_{11} (entry 3) on the edge connecting them:



For three bits there are eight phases and eight edges. Entry 1 is m_{001} , a single 1 in the last place, so it sits on B_3 's self-loop; entry 7 is m_{111} , all three nodes at once, the genuine hyperedge, drawn meeting at a point in the center:




A loose analogy may be made to a neural network. The values in the phase vector all lie on connections between the bits, akin to the weights on a neural network's connections. A recurrent neural network additionally has self-loops, where a node connects back to itself with its own weight. A neural network also has a bias, an overall background parameter belonging to no connection in particular; here, that bias is the empty edge. The analogy is loose, but it is instructive: like a neural network's weights, the phases are where the *interactions* between the parts are parameterized. As we are about to see in detail, gates do their most interesting work not on the nodes but on the edges.


The picture comes with the same warning label as everything in Part II. The bit values written inside the nodes are the system's configuration (its ontic state, on our reading). The phases written on the edges are a *sufficient statistic of the collection's history*. The hypergraph is a diagram of one kind of stuff annotated with a compressed record of where its statistics have been. The phases are drawn *about* the bits, not located *in* them, which is exactly why no phase can be assigned to any bit alone and why, as the next sections show, phases slide along the edges in ways no local property could.

Take the simplest case, a single bit B . It has two configurations and so two edges: the empty edge carries m_0 and the node's self-loop carries m_1 . Before any gate acts the bit simply has a value, and both phases are zero.

Begin with B at 0 and apply one Hadamard gate. The distribution becomes uniform, and both phases stay at zero:



Begin with B at 1 instead and apply the same gate. The distribution is uniform again and the empty edge is still zero, but the self-loop now carries π :



Nothing in the distribution tells these two apart. Measure the bit in either case and 0 and 1 come back equally often. The whole of the difference sits on the self-loop, and what it records is where the bit started.

Now apply the Hadamard a second time. The bit that began at 0 returns to 0, and the bit that began at 1 returns to 1. The two inputs to that second gate agree in every respect but one: both distributions are uniform, both empty edges read 0, and they differ only in whether the self-loop carries 0 or π . The gate reads the configuration and the phases and nothing else, so that single difference is what decides the outcome: a self-loop of 0 sends the bit back to 0, and a self-loop of π sends it back to 1.

This is what a sufficient statistic does. A bit may have been through any number of operations, and almost none of that past still bears on what happens next. The phase keeps the part that does, compressed into one number per configuration, and that is enough to finish the job.

Something further follows from the run that started at 0. That bit never acquires a phase at all: every entry of $\mathbf{m}$ is zero before the first gate, after it, and after the second. The interference is present all the same, since the second gate takes a uniform distribution back to a degenerate one, which no stochastic matrix on the present distribution can do. The phase-free rule of Section 5.2, $\mathbf{p}^{[t+1]} = \left| U \sqrt{\mathbf{p}^{[t]}} \right|^2$, reproduces that circuit exactly, and a system evolving under it holds no phase at any step. Whatever accounts for the interference here, the phase cannot be it, because there is none to appeal to.

Section 5.7 already showed why. Expanding the evolution recursively eliminates the phase vector by substitution, leaving the distributions at every prior time together with the definitions of the operators and nothing else. The phase is a compression of that history rather than a further ingredient of the world, which is what it means to call it a sufficient statistic. What accounts for the interference is non-Markovianity: the second gate's behavior depends on how the present distribution came about. Two bits can share a distribution, share a phase of zero, and still be handed different futures, because their histories differ.

7.2 Correlation hidden in the phase, revisited

Chapter 4 exhibited a quantum state whose measured statistics showed no correlation at all (classical mutual information zero) while its quantum mutual information was maximal. We can now say where the missing correlation was: on the edges. Consider the circuit


$$\begin{array}{cccccc}
 \mathbf{e}_0 & & \mathbf{H} & & \overline{\mathbf{C}\mathbf{X}} & & \perp \\
 \mathbf{e}_0 & & \mathbf{R}_X(\pi) & & \overline{\mathbf{C}\mathbf{X}} & & \mathbf{H} & & \perp
 \end{array}$$

which produces the quantum state

$$\psi = \frac{1}{2} \begin{pmatrix} -i \\ i \\ -i \\ -i \end{pmatrix},$$

with the uniform Born distribution (0.25, 0.25, 0.25, 0.25): two uncorrelated fair coins, as far as any value-measurement is concerned.

Now apply one more Hadamard to the first bit, and watch what the interference term of Chapter 5 does. Under $H \otimes I_2$, each configuration's amplitude flows to two output configurations (those agreeing with it on the second bit), multiplied by $+\frac{1}{\sqrt{2}}$, except for the flow from configurations with $B_1 = 1$ back to themselves, which carries $-\frac{1}{\sqrt{2}}$. We can represent the effect on the statistical distribution using an *amplitude tree*: each parent configuration at the top, with its amplitude; each child configuration at the bottom, receiving contributions along the branches (solid branches carry $+\frac{1}{\sqrt{2}}$, dashed branches $-\frac{1}{\sqrt{2}}$):



At the child 01, the contributions $+\frac{i}{2} \cdot \frac{1}{\sqrt{2}}$ and $-\frac{i}{2} \cdot \frac{1}{\sqrt{2}}$ meet and cancel; likewise at 10. What survives is

$$(H \otimes I_2) \psi = \frac{1}{\sqrt{2}} \begin{pmatrix} -i \\ 0 \\ 0 \\ i \end{pmatrix}, \quad |(H \otimes I_2) \psi|^2 = \begin{pmatrix} 0.5 \\ 0 \\ 0 \\ 0.5 \end{pmatrix},$$

the perfectly correlated distribution. One local gate, touching one bit, converted invisible phase structure into visible value correlation, because the differing signs on 01 and 10, sitting on the hypergraph's edges where value-measurements cannot reach, were the correlation all along, waiting for an interference term to route them back into the probabilities. This is what the QMI was counting in Chapter 4, and it is the first demonstration of the chapter's slogan: information the present values do not carry, the history does. And gates can retrieve it.

7.3 Superdense coding: two bits from one

Now we put a full bit of history-stored information to work.

Alice and Bob share a Bell pair, prepared as in Chapter 3: Alice holds B_1 , Bob holds B_2 . Alice wants to send Bob a two-bit message (00, 01, 10, or 11), but she is only allowed to send him a single bit: her half of the pair. She encodes her message by applying, *to her bit*

alone, one of four operations: nothing, X, Z, or X followed by Z. Then she mails the bit. Bob, holding both bits, decodes by applying a CX (Alice's bit as control) followed by a Hadamard on Alice's bit, and measures both bits. The four complete circuits:


$$\begin{array}{ccccc}
 & & 00 & & 01 \\
 \mathbf{e}_0 & \text{H} & \overline{\text{CX}} & \text{I} & \overline{\text{CX}} & \text{H} & \perp \\
 \mathbf{e}_0 & & \text{CX} & & \text{CX} & & \perp \\
 & & 10 & & 11 \\
 \mathbf{e}_0 & \text{H} & \overline{\text{CX}} & \text{Z} & \overline{\text{CX}} & \text{H} & \perp \\
 \mathbf{e}_0 & & \text{CX} & & \text{CX} & & \perp
 \end{array}$$

The reader can verify (the chapter's code does) that the four circuits end, deterministically, in the configurations 00, 01, 10, and 11 respectively: Bob reads off both bits of the message without error.

To see how one transmitted bit carried two, freeze each circuit at the moment the bit is in the mail (after Alice's encoding, before Bob's decoding) and look at the transmitted pair's quantum state:

$$\psi_{00} = \frac{1}{\sqrt{2}} \begin{pmatrix} 1 \\ 0 \\ 0 \\ 1 \end{pmatrix} \quad \psi_{01} = \frac{1}{\sqrt{2}} \begin{pmatrix} 0 \\ 1 \\ 1 \\ 0 \end{pmatrix} \quad \psi_{10} = \frac{1}{\sqrt{2}} \begin{pmatrix} 1 \\ 0 \\ 0 \\ -1 \end{pmatrix} \quad \psi_{11} = \frac{1}{\sqrt{2}} \begin{pmatrix} 0 \\ 1 \\ -1 \\ 0 \end{pmatrix}.$$

The Born distributions split the four messages into two groups. For 00 and 10, the pair is perfectly correlated; for 01 and 11, perfectly anti-correlated. So one bit of the message is sitting in the *values*: Bob learns it by checking whether his two bits agree. The second bit of the message is nowhere in the values: within each group, the two distributions are identical. It is in the *phase*, and on the hypergraph it is visible at a glance:



The presence or absence of a π on the network distinguishes the rows of the message just as agreement or disagreement of the values distinguishes the columns. Bob's decoding circuit is nothing but the Chapter 4 trick performed twice: local operations that route edge information back into the values, where measurement can read it.

The bit Alice mailed behaves as if it carried two bits of information, despite being, by every direct test, a single bit. One might try to explain this by giving the bit a hidden

internal memory and treating the phase as a physical object riding along with it. But we know better by now: the phase is a sufficient statistic over the pair's *joint history*, and Alice's four encodings are four *distinguishable preparations of that history*. Her Z gate did not stuff an extra bit inside the mailed bit; it made the pair's past different in a way the law consults when Bob's gates act. The arithmetic is exactly Chapter 4's QMI count: a Bell pair holds two bits of mutual information, one in the values (the CMI) and one in the phases, and superdense coding simply retrieves both.

7.4 Phase kickback

We have seen phases move from edges into values. Phases also move along the network in a subtler way (from one bit's neighborhood to another's), and the mechanism gets its own section, because it is the engine inside most quantum algorithms.


Recall two facts. Applying the Hadamard gate twice in a row to a bit with an initial value of 0 deterministically returns it to 0. And the CX gate is, by its truth table, *passive on the control*: the control bit's value is never altered; the gate only ever writes onto the target.

Now set up the following two-bit circuit. We apply the Hadamard twice to B_1 , but in between the two Hadamards, we apply CX with B_1 as the control. The target, B_2 , is prepared with the value 1 and given a Hadamard of its own first:

$$\begin{array}{cccccc} \mathbf{e}_0 & \text{H} & \overline{\text{CX}} & \text{H} & \perp & \\ \mathbf{e}_1 & \text{H} & \text{CX} & & \perp & \end{array}$$

A strange thing happens: B_1 ends deterministically with a value of 1, not 0. The gate that is passive on the control has changed the control, and only the control: B_2 , the bit the CX supposedly writes onto, ends in a uniform distribution, exactly as it began after its Hadamard. The reader should sit with how backwards this is. The one bit the gate is allowed to change is statistically untouched; the one bit the gate cannot change is redirected completely.

The hypergraph shows precisely what happened. After the first column of Hadamards, the quantum state is $\frac{1}{2}(1, -1, 1, -1)$: uniform values, with a phase of π on B_2 's self-loop and π on the connecting edge. When the CX gate is applied, the state becomes $\frac{1}{2}(1, -1, -1, 1)$: the values are untouched, but the π on the connecting edge has been *kicked back* onto B_1 's self-loop:



This is known as *phase kickback*. It does not alter the value statistics of B_1 in the moment (both snapshots have uniform Born distributions), but it changes B_1 's *history*, and therefore its behavior at a future logic gate. When the second Hadamard is applied to B_1 , the interference term consults that history and steers the bit to 1 rather than 0.

Line this up against the two other versions of the same sandwich we have already computed. Hadamard, nothing, Hadamard: the bit returns to 0 (Chapter 5). Hadamard, recording CX

onto a fresh $\mathbf{e}_0$ target, Hadamard: the bit comes out uniform, because the recording created a base case and carried the history off (Chapter 6). Hadamard, kickback CX onto an $\mathbf{H}\mathbf{e}_1$ target, Hadamard: the bit comes out 1: the interaction did not carry the history off but *edited* it. One gate in the middle, three preparations of its target, three completely different fates for a control bit whose values the gate never touches. There is no cleaner exhibit of our central claim: what a quantum gate does is a function of the history it is handed, and the phases are the record in which that history is kept.

And notice what kickback says about ownership. Before the CX, the π lay on the connecting edge: a relation between the bits, belonging to neither. After, it lies on B_1 's self-loop. No probe of B_1 alone, at any single moment, can detect the difference; yet the difference commands B_1 's future. Properties like this (unlocatable in any part, undetectable at any instant, decisive over time) are exactly what one expects of a statistic of a joint history, and exactly what one should refuse to expect of a piece of local physical furniture.

7.5 When classical simulation is hard

Evolving a classical probability distribution takes just as much work as evolving a quantum one. Both are vectors with an entry for every configuration, both are advanced by a matrix built from a local gate, and both take 2^n pieces of arithmetic to push one step. Nothing in that comparison favors either side.

The difference is that a classical stochastic process can be simulated without evolving the distribution at all. Start the run at a definite configuration. At each gate, the configuration you are in picks out a column of the gate's matrix; draw a random number, read off which row it lands on in that column's cumulative total, and move there. Repeat for every gate and you have one run. Repeat the whole thing over many shots and the frequencies estimate whatever the circuit outputs, with the full distribution never built. Section 2.7 did precisely this.

That shortcut is unavailable here, because a gate has no settled behavior of its own. Section 6.2 showed the same Hadamard governing its step by two different matrices depending on where in the circuit it sits.

	<i>step 1</i>	<i>step 3</i>
classical, COIN, X, COIN	$\begin{pmatrix} .5 & .5 \\ .5 & .5 \end{pmatrix}$	$\begin{pmatrix} .5 & .5 \\ .5 & .5 \end{pmatrix}$
this one, H, T, H	$\begin{pmatrix} .5 & .5 \\ .5 & .5 \end{pmatrix}$	$\begin{pmatrix} .854 & .854 \\ .146 & .146 \end{pmatrix}$

In a classical circuit the gates never need composing, each being indifferent to the others. Here they are not. The phases a gate carries act on the phases already present, so what a gate does depends on how it sits in relation to everything before it, so the step law belongs to the circuit rather than to the gate, and recovering it means knowing what every earlier gate did. Nothing short of the whole vector will supply that. The sampling shortcut is gone, and what replaces it is the same work a classical machine would do only if it insisted on carrying the distribution: one pass over 2^n entries for every gate. Section 7.6 shows why no cheaper route exists.

Composing is not always obligatory. Where the product has a short description, the short-cut returns in another form. Finding which circuits those are takes a change of bookkeeping: the same circuits, written down differently.

Section 3.4 gave the observables, the Pauli strings built from them, and the parity each returns. The move is to write a state as a list of those expected values rather than a list of amplitudes. Amplitudes are indexed by configurations, 00, 01, 10 and 11 for two bits; expected values are indexed by Pauli strings, one entry to a string.

For one bit that is four entries, in the order I, X, Y, Z. A bit known to be 0 on the computational basis is

$$(1 \ 0 \ 0 \ 1)^\top :$$

the identity is always 1, Z is +1 because the value is known, and X and Y are 0 because nothing is known of them. Lists for separate bits combine by the Kronecker product as state vectors do, so two bits give sixteen entries indexed II, IX, IY, IZ, XI, XX and on to ZZ: 4^n of them for n bits.

Gates act on such a list directly, through a real matrix called the operator's *Pauli transfer matrix*, with entries

$$R_{jk} = \frac{1}{2^n} \text{Tr}(P_j U P_k U^\dagger),$$

where P_j and P_k range over the Pauli strings. It is a truth table of the same sort the stochastic matrices were: column k is the input string, row j the output, and the entry says how much of one the gate turns into the other. For the Hadamard it is

	I	X	Y	Z
I	1	0	0	0
X	0	0	0	1
Y	0	0	-1	0
Z	0	1	0	0

with one entry to a column. Read down the X column and the single 1 sits in the Z row: H turns X into Z. Read the Z column and it turns Z into X. The Y column carries -1 , so Y comes back negated. In this bookkeeping the Hadamard is a swap and a sign, and nothing else.

Taken on its own the change is a step backwards. The list of expected values has 4^n entries where the amplitudes had 2^n , so writing a state this way costs the square of what it cost before, and evolving one gate costs accordingly. What it opens up is a question one could not ask in the other bookkeeping: what do the gates do to the entries, one entry at a time?

The answer is that some gates merely permute them (every entry of the table a whole number, 0 or ± 1 , one to a column), as H's table above does, and S and CX do the same. A gate of that kind is called a *Clifford* gate: conjugating any Pauli string by it returns another Pauli string rather than a combination of several. H, S and CX generate the whole Clifford group between them. T does not belong, and the reason is one entry in its table: $\text{TXT}^\dagger = (X + Y)/\sqrt{2}$, two terms where H and S give one.

That rescues the bookkeeping completely. If every gate only shuffles the entries, there is no need to carry 4^n of them: carry the strings whose value is certain, and the shuffling says where each goes. Enough certainties determine the state, and Section 3.4's commuting pair is already an example. The Bell pair of Section 3.5 is the only state of two bits with $\text{ZZ} = +1$ and $\text{XX} = +1$ together, so four amplitudes are replaced by two parities, and its third certainty $\text{YY} = -1$ follows from them.

A Clifford gate sends a Pauli string to a single Pauli string, so a state described by n strings is still described by n strings after the gate, and the description never grows. That is the Gottesman–Knill theorem,¹ and a circuit built from Clifford gates alone can be simulated efficiently however deep it runs and however many bits it uses, the work rising no faster than with the square of the number of bits.²

The chapter’s code implements it, and the bookkeeping is small enough to set out. Each Pauli string is held as two bits per position: whether an X stands there, and whether a Z does. Neither is the identity, the first alone is X, the second alone is Z, and both together make Y, since Y is an X and a Z in the same place up to a factor. One more bit to a string records its sign. A register prepared at 0 is left unchanged by Z on the first bit, by Z on the second, and so on down, so the simulation begins from those n strings and holds nothing else.

Each gate then rewrites the strings, and the rewriting is bit flipping. A Hadamard on bit a exchanges X and Z there, so it swaps the two bits in that column of every string. An S sends X to Y, so it switches the z bit on wherever the x bit already is. A CX copies in both directions at once, the x bit forward from control to target and the z bit back again. Each gate is a pass over n rows, whatever the circuit has done beforehand, and the table never grows.

Reading bit a comes in the two cases that make the sampling work. If some string carries an X or a Y at position a , the strings do not yet fix that bit, the reading is a fair coin, and the table is rewritten to record what came up. If every string carries only I or Z there, the value is already settled, and it is recovered by multiplying together the strings that bear on it and reading the sign. The table in the listing holds $2n$ strings rather than n , the extra ones being partners that turn that last multiplication into a few lines. For the Bell pair the two cases appear in order: the first bit read is a free coin, the second is settled, and the correlation falls out of the bookkeeping rather than having been put in.

What those shortcuts amount to is a description of the quantum state shorter than its list of amplitudes, and how short a description exists depends on the circuit:

<i>what is held</i>	<i>how large</i>	<i>when it suffices</i>
the amplitudes	2^n numbers	always
n Pauli strings	about n^2 bits	Clifford circuits
a sum of such descriptions	one term per branch	a few gates beyond

All three fix the same state exactly, as the Bell pair just did with two strings in place of four amplitudes. So 2^n is the size of the least thoughtful faithful description, and no measure of the state itself.

The gates that break the short description are the ones whose conjugation returns a sum. T sends X to $(X + Y)/\sqrt{2}$, two terms where H and S return one, so each such gate can double the number of pieces the description needs. Bravyi and Gosset turn that into a simulation whose running time stays manageable in the number of bits and in the number of Clifford gates, and explodes only with the count of the others, at roughly $2^{0.23t}$ for t of them.³ The

¹D. Gottesman, “The Heisenberg Representation of Quantum Computers,” arXiv:quant-ph/9807006 (1998).

²S. Aaronson and D. Gottesman, “Improved Simulation of Stabilizer Circuits,” *Physical Review A* 70 (2004), 052328.

³S. Bravyi and D. Gosset, “Improved Classical Simulation of Quantum Circuits Dominated by Clifford Gates,” *Physical Review Letters* 116 (2016), 250501.

dial is neither the number of bits nor the depth nor how entangled the circuit becomes, all of which Clifford circuits run up freely. It is how many gates push the state past a short description.

Ensemble statistics reach the same place from another direction. For any stochastic process the past carries a definite quantity of information bearing on the future, and no model can work with less; the memory an optimal *classical* model must carry exceeds it for almost every process. Gu, Wiesner, Rieper and Vedral construct quantum models that break the classical bound,⁴ and Korzekwa and Lostaglio prove advantages in time as well as in memory for simulating stochastic processes, describing the coherence that does the work as an internal memory of the system during its evolution.⁵

Each result above rules out a family of methods, and a family ruled out is not a problem proved hard. Whether some further method sidesteps the composing step for every circuit is the open question of the field, and nothing here settles it. What can be said is that the only machine anyone has built that skips the step is a quantum computer, which runs its own circuit and lets one sample it.

7.6 Reading the answer out

Composing the gates and squaring, $|U_d \cdots U_1|^2 \mathbf{p}$, is not how a circuit is typically computed. The state vector is pushed through one gate at a time instead, $\psi \mapsto U_1 \psi \mapsto U_2 U_1 \psi$ and onward, and squared at the end. Both give the same answer, and the second is cheaper by a factor of 2^n : composing touches 4^n entries for each gate, pushing the vector touches 2^n .

The two answer different questions. Composing works out how the gates relate for every configuration the register might have started in: column j of the composed matrix is the vector computation run from configuration j , and there are 2^n columns. The vector route answers for the one configuration the register actually started in. Neither route avoids any work the other does. Ask for the whole map and the vector route has to be run 2^n times, which is precisely what composing costs. But if we care only about the output for a single input configuration, and by convention a quantum circuit begins with every bit initialized to 0, then there is time to be saved.

What the state vector is, then, is the answer to the narrower question, which is why it can be smaller than the operator. That is the shape Chapter 5's sufficient statistic had as well: enough to carry the one history the system is actually in, and silent about the histories it is not.

The narrower answer is a vector of amplitudes, and with them the phases. Those phases are where algorithms are usually described as doing their work, and the description is a fair one: an algorithm arranges for phases to cancel here and reinforce there. Nothing in this book takes them for physical, but picturing them as a real property of the system is a convenience hard to do without, which is why the hypergraph of Section 7.1 was drawn at all. On it a phase sits on an edge rather than on a node, and an edge may join any number of them: the self-loops touch one bit, the ordinary edges two, the hyperedges more, and the empty edge

⁴M. Gu, K. Wiesner, E. Rieper, and V. Vedral, "Quantum mechanics can reduce the complexity of classical models," *Nature Communications* 3, 762 (2012).

⁵K. Korzekwa and M. Lostaglio, "Quantum Advantage in Simulating Stochastic Processes," *Physical Review X* 11 (2021), 021019.

none. So a phase is not a property that can always be localized to a bit, and most of them cannot be.

Read that way the arrangement is a hidden Markov model of the kind Section 5.3 described, with 2^n units of memory added to the n bits and held in the connections. It is the same memory Section 7.5 found a classical imitator must carry. And it explains why a gate here does more than a gate in a classical circuit does. There, adding a gate changes the values of the bits it touches. Here, adding a gate changes how every gate in the circuit stands in relation to every other, which is a fresh set of 2^n phases.

That is the resource, and the same fact is the obstacle: a hidden memory is hidden. Phases are not directly measurable. Measurement reveals bit values, and nothing but bit values, every time, and so an algorithm that lets its answer end up in the phases has computed nothing anyone can read. Algorithms that make use of the phases must therefore be cleverly constructed so that, after the phases have done the work of the calculation, the work is routed back into the probabilities (interference, arranged on purpose) so that the answer can be read off the distribution of measured values. Every technique here is a routing trick of exactly this kind: the Chapter 4 rotation, Bob's decoding circuit, the final Hadamard of the kickback circuit.

Since the readout is a distribution, many quantum programs must be executed many times over to collect their statistics; each execution is called a *shot*. The art of quantum algorithm design is almost entirely the art of managing this obstacle: arranging for wrong answers to cancel and right answers to accumulate, so that few shots (ideally one) suffice. It is also why the speedups are selective. A routing trick has to be found problem by problem, and for most problems none is known, which is why they form a short list rather than a general rule.

The kickback circuit is already a complete, if tiny, example of the art, and the next section works it through as a full algorithm. The single run *interacted with both possible control values at once* in the only sense we admit: the gate's behavior is a function of everything that comes before it in the circuit, and after a Hadamard what comes before spans both control values. That is what grows exponentially. One interaction consults a summary of the whole history; one classical evaluation consults a single configuration.

7.7 Deutsch's algorithm

Here is the smallest problem on which a quantum computer provably beats a classical one, and it is built from the kickback circuit of Section 7.4 and nothing else.⁶

There are exactly four functions from one bit to one bit. Two are *constant*, sending both inputs to the same output; two are *balanced*, returning each output for exactly half the inputs:

<i>Constant</i>	<i>Balanced</i>
$0 \mapsto 0, \quad 1 \mapsto 0$	$0 \mapsto 0, \quad 1 \mapsto 1$
$0 \mapsto 1, \quad 1 \mapsto 1$	$0 \mapsto 1, \quad 1 \mapsto 0$

Someone hands us a sealed box computing one of the four and will not say which, and we are allowed to run it *once*. The question is not which function the box holds, which one run cannot settle, but which *category* it belongs to.

⁶D. Deutsch, "Quantum theory, the Church-Turing principle and the universal quantum computer," *Proceedings of the Royal Society of London A* 400 (1985), 97–117.

Classically one run cannot settle even that. It reports f at a single input, and each category contains a function taking that value there, so whatever comes back, two candidates survive, one from each category. Settling the category classically takes two runs, one per input, after which comparing $f(0)$ with $f(1)$ decides it.

Before the box can be a quantum gate it has to be made reversible, and two of the four functions are not: a constant function sends both inputs to the same output, so its output does not determine its input. The standard repair gives the box a second bit to write on. It takes x together with an extra bit y , leaves x alone, and returns y with $f(x)$ added modulo 2:

$$U_f : \quad x, y \mapsto x, y \oplus f(x).$$


Each of the four functions becomes a permutation of the four two-bit configurations this way, and permutations are what Chapter 1 called reversible. As gates: the constant $f = 0$ does nothing at all; the constant $f = 1$ flips y unconditionally, an X on the second bit; the balanced $f(x) = x$ flips y exactly when x is 1, which is CX; and the balanced $f(x) = \neg x$ flips y exactly when x is 0, which is CX with an X on y in front of it. As circuits, with x on the upper row and y on the lower one throughout:

$f = 0$	$f = 1$	$f(x) = x$		$f(x) = \neg x$
I	X	$\overline{C}X$	$C\overline{X}$	X
		$C\overline{X}$	$\overline{C}X$	

Each box touches only y , and only ever conditionally on x , which is what $y \oplus f(x)$ says. Writing the configuration as xy , the four boxes are these permutations:

<i>Constant</i>		<i>Balanced</i>	
$f = 0$	$f = 1$	$f(x) = x$	$f(x) = \neg x$
$00 \rightarrow 00$	$00 \rightarrow 01$	$00 \rightarrow 00$	$00 \rightarrow 01$
$01 \rightarrow 01$	$01 \rightarrow 00$	$01 \rightarrow 01$	$01 \rightarrow 00$
$10 \rightarrow 10$	$10 \rightarrow 11$	$10 \rightarrow 11$	$10 \rightarrow 10$
$11 \rightarrow 11$	$11 \rightarrow 10$	$11 \rightarrow 10$	$11 \rightarrow 11$

Making the box reversible has not made the classical problem any easier, and the reason is sharper than “one query is not enough.” Every single input-and-output pair that one of the constant boxes can produce is also produced by one of the balanced boxes:



Whatever the box returns, the pair we observe sits on both sides of that picture, so the one thing a single classical query establishes is consistent with either category. The eight arrows are the whole classical obstruction, drawn out.


Now prepare x at 0 and y at 1, put a Hadamard on each, run the box, and put a Hadamard on each again:

$$\begin{array}{ccccccc} x & \mathbf{e}_0 & \text{H} & U_f & \text{H} & \perp \\ y & \mathbf{e}_1 & \text{H} & U_f & \text{H} & \end{array}$$

Measure x . It reads 0 for both constant boxes and 1 for both balanced boxes, on every run, with certainty. One query settles the category.

The mechanism is the kickback of Section 7.4. A second bit prepared at 1 and then Hadamarded is the one arrangement in which a controlled flip of y disturbs y 's values not at all, and instead deposits a π on the first bit's phases. The box's whole job is to flip y conditionally, so the box writes its answer into x 's phases rather than into anybody's values. A constant box treats both values of x alike, flipping y never or always, so the phase it deposits on the two x configurations is the same; a phase common to both steers nothing, the two Hadamards on x cancel exactly as they did in Chapter 5, and x returns to 0. A balanced box flips y for one value of x and not the other, so it deposits a π on one x configuration and not on the other, and that *difference* between the two phases is precisely what the second Hadamard converts back into a definite value: x comes out 1.

The hypergraph shows all four boxes at once. Each one leaves the values uniform and writes only on the edges, so the four differ only in where the π sits once U_f has acted:



By the numbering of Section 7.1, the two edges that involve B_1 (its self-loop and the connecting edge) carry the phases of the configurations where x is 1, and the two that do not (the empty edge and B_2 's self-loop) carry those where x is 0. Reading each picture as those two halves, the constant boxes give halves that match and the balanced boxes give halves that differ by π . That is the whole algorithm: the second Hadamard on B_1 responds to a difference between the two halves and to nothing else.

The four pictures are also two pictures. Each constant box carries a π exactly where the other carries none, on every edge, which is a global phase in the sense of Section 3.2. The rule of Section 5.6 consults the phases only through their differences, so a shift laid on all of them alike changes nothing it reads. The two constant boxes therefore leave x with histories that agree in the only respect the law consults: they are sufficiently equivalent, in the sense of

Section 5.8, and less than identical. The two balanced boxes stand in the same relation to each other. Fixing the shift by putting 0 on the empty edge leaves one picture to a category, and they are the pair already drawn in Section 7.4. A constant box leaves the hypergraph where the first Hadamards put it, before the kickback. A balanced box carries the π off the connecting edge and onto B_1 's self-loop, after the kickback. Deutsch's question is whether the kickback happened, and one query answers it because the kickback is a property of how the box treats both inputs.

Read in our terms, the bit x has one value in each run, as it always does. What the law consults is the circuit that produced that value, and after the first Hadamard that circuit spreads x evenly over both configurations. The *distribution* is a structural feature of that history, not a quantity carried alongside the bit. Being constant or balanced is a property of how f behaves across both inputs, which is to say a property of that same distribution and not of any single configuration; and the box's effect on x 's phases is a function of the history it acts on. That is how one interaction answers a question about two inputs. What one query settles is settled by a law whose form depends on a statistic that no single trial carries.

The same algorithm runs on more than one input bit. Hand over a box computing a function of n bits, promised to be constant or to be balanced in the sense of returning each output on exactly half of its 2^n inputs, and the circuit above carries over directly: a Hadamard on every input, one query, Hadamards again, and a reading that comes out all zeros when the box is constant and something else when it is balanced. Deutsch and Jozsa gave the extension in 1992, called the Deutsch–Jozsa algorithm.⁷ One query still settles the category, where a classical machine needs $2^{n-1} + 1$ of them in the worst case.

7.8 Why updating appears mandatory

One question left open in Part I can now be answered, and the answer leads into the next chapter.

In probabilistic computing, applying Bayes' theorem is *optional*. If a bit is measured and we decline to update the probability vector, nothing downstream goes wrong: the un-updated vector still describes the ensemble, our predictions are simply less sharp than they could be, and anyone who conditions on the measurement later gets the same answers as someone who conditioned immediately. Classical ignorance is free.

In quantum computing it appears not to be. Recall from Chapter 5 that the phase vector's evolution rule takes the probability vector as an input: the phases at the next step depend on the probabilities at this step. Consider what this means for anyone who treats the phase as a physical object: a piece of physical furniture whose law of motion has, as an input, a *description of someone's uncertainty*. A physical thing steered by a subjective thing is not a coherent picture of anything, and this incoherence is among the strongest reasons to refuse the phase realism that Part II has been dismantling. On our reading there is no puzzle: the phase is a statistic of the statistical history, so of course it is a function of statistics.

But there is a practical consequence, and it is a serious one. When a measurement occurs, it creates a base case (the measured system becomes maximally correlated with a record that does not return), and the system's consultable history is restarted from the updated,

⁷D. Deutsch and R. Jozsa, "Rapid solution of problems by quantum computation," *Proceedings of the Royal Society of London A* 439 (1992), 553–558.

conditioned distribution. The law itself now consults the post-measurement statistics. An observer who declines to update, who keeps evolving the old description as though the measurement had not happened, is no longer merely unsharp; they are feeding the wrong history into that rule, computing the wrong phases, and thereby predicting interference effects that will not occur, or missing ones that will. Their predictions *deviate from what is empirically observed*.

The obligation binds only someone who has already decided to describe the measured system *by itself*, and nothing forces that decision. A measurement is an interaction, and an interaction has two sides. Rather than apply the update of Section 3.7, zeroing the configurations the outcome has ruled out and renormalizing what survives, which is the step usually called the collapse of the wavefunction, one may enlarge the description to take in whatever the system became correlated with: the pointer, the record, the apparatus, the observer. The joint description then evolves under its own unitary in the ordinary way, and nothing is zeroed anywhere in it. Nothing has been lost, because nothing was discarded.

Chapter 4 exhibited both sides of this trade. The reduced dynamics of a system coupled to an environment is a sum over the environment's outcomes, and it is the *summing* that costs, not the coupling. A base case, in Chapter 6's sense, is created when a system becomes maximally correlated with a record that does not come back. Keep the record in the description and it has not gone anywhere; there is no base case for the enlarged system, its history is unbroken, and the law consults that unbroken history exactly as before.

So the rule is not that the update must be applied. It is that one may not both drop the correlate from the description *and* carry on with the distribution that preceded the measurement. Either keep the apparatus and evolve the whole of it, or discard the apparatus and apply the update. What produces wrong predictions is discarding the apparatus while continuing as though the old description still held. Which of the two to do is not settled by the physics but by how much of the interaction one is willing to carry along, and since any description that can actually be worked with stops somewhere, the update is what keeps it correct at the point where it stops.

The most famous illustration of this fact is an experiment in which a particle's statistics change depending on whether anything (or anyone) could, even in principle, have conditioned on which path it took: the Mach-Zehnder interferometer. It is the subject of Chapter 8, where the machinery of Part II finally meets the experiment that convinced a century of physicists that particles must be waves.

7.9 Quantum cryptography

Three protocols follow, because each depends upon something Part II established rather than upon any computational difficulty. Take them in order: encrypting a message, agreeing on a key, and proving who you are.

Start with encryption. The XOR operator is balanced, and it stays balanced when either input is held fixed. Whatever value the first input takes, the output is 0 for one setting of the second and 1 for the other. It is also reversible in the sense that any two of the three values fix the third. Write P for a plaintext bit, K for a key bit, and C for the ciphertext bit. Then

$$P_i \oplus K_i = C_i, \quad C_i \oplus K_i = P_i.$$

Because XOR stays balanced with one input fixed, a key drawn uniformly makes the ciphertext uniform whatever the plaintext was. The ciphertext carries no information about the message at all. This is the *one-time pad*, and no amount of computation breaks it, because there is nothing in the ciphertext to break.

The cost is severe. The key must be as long as the message and can never be reused, which is why the pad is rare in practice and why ciphers like ChaCha20 stretch a short key into a long keystream instead. And the guarantee holds only if the key is unpredictable. A pseudorandom generator is a deterministic machine whose output can be reconstructed by anyone who learns its seed, and classical environmental noise is deterministic too, however impractical it may be to track. A bit prepared with a Hadamard gate and then measured is different. The transition is stochastic at the level of the law itself, with no deeper machinery posited to produce it, so the theory leaves nothing underneath for an attacker to reconstruct. Section 8.10 was careful to leave open whether such machinery exists, and that caution applies here: what the pad needs is that quantum theory supplies none of it. One gate and one measurement give a key bit distributed uniformly, which is what the pad requires.

Then comes the harder problem, *key distribution*. The pad is a symmetric cipher, so Alice and Bob must already share the key. Sending it unencrypted defeats the purpose, and sending it encrypted is circular, since Bob has nothing to decrypt with. Classically this is solved by Diffie–Hellman, which prevents an eavesdropper from reading the key in transit. The quantum protocol cannot prevent that. What it can do is guarantee that reading the key leaves a mark.

The mechanism goes back to Section 5.1. Two Hadamard gates in succession cancel, and a measurement placed between them destroys the cancellation. A measurement placed between them establishes QMI between the bit and whatever holds the record of the outcome, and once that record is carried away and never brought back, the history the second gate would have consulted is gone. The gate has nothing left to undo, and its output is uniform instead of the value Alice started from. The protocol built on this is called *BB84*, after Bennett and Brassard, who introduced it in 1984.⁸

1. Alice generates a keystream of bits, each e_0 or e_1 at random.
2. For each bit she tosses a fair coin and applies a Hadamard gate or does not.
3. She sends the bits to Bob.
4. For each bit Bob tosses his own fair coin and applies a Hadamard gate or does not, then measures.
5. Once the whole stream has arrived, the two announce publicly which bits they applied the gate to.
6. They discard every bit where their choices disagreed, since only the agreeing ones are guaranteed to match.
7. They publicly compare a sample of what remains.

⁸C. H. Bennett and G. Brassard, “Quantum cryptography: Public key distribution and coin tossing,” in *Proceedings of IEEE International Conference on Computers, Systems and Signal Processing* (Bangalore, 1984), 175–179.

8. Disagreement in the sample means an eavesdropper. Agreement means the rest of the stream is safe to use as a pad.

Where their choices agreed, either both applied a Hadamard and the two cancel, or neither did, and Bob recovers exactly the bit Alice sent. Consider what an eavesdropper faces on one of those bits.

Suppose Alice sent it with no gate applied. The eavesdropper measures it, reads Alice's value, and sends an identical bit onward. Bob measures without a gate and gets what Alice sent. Nothing is disturbed, and she has the key bit.

Suppose instead that Alice applied a Hadamard. If the eavesdropper simply measures, her measurement lands between Alice's gate and Bob's, and Section 7.8 gives the consequence: the bit is left maximally correlated with a record that does not come back, and the history restarts from there. Bob's Hadamard has nothing to cancel, so his result is uniform rather than Alice's value. Half the time he records a value Alice never sent, and the disagreement surfaces when the two compare their samples.

The eavesdropper has an answer to this. A Hadamard of her own, applied before she measures, undoes Alice's, so she reads the bit in the form Alice prepared it. A second Hadamard afterward puts the bit back into the form Alice sent. Bob's gate then cancels hers, he recovers Alice's bit, and nothing gives her away.

The answer works only when Alice used a gate. If Alice did not, the sandwich causes the damage instead. The eavesdropper's first Hadamard puts the bit into a form Alice never prepared, her measurement fixes a value there, and her second Hadamard hands Bob a bit that measures uniform. Bob applies no gate, reads it, and disagrees with Alice half the time.

So the eavesdropper must sandwich her measurement in Hadamards exactly when Alice used one, and measure the bit directly when Alice did not, and she cannot know which, because the choices are not announced until the whole stream has arrived. Her choice is wrong on half the bits, and half of those reach Bob wrong, which is why each checked bit catches her with probability $1/4$. Over n checked bits that is $1 - (3/4)^n$: about 99% at sixteen and about 99.99% at thirty-two. The guarantee is again physical. No technique lets her read a bit in transit without leaving a record correlated with it.

One attack survives. If the eavesdropper sits between Alice and Bob and impersonates each to the other, she can run the protocol twice, hold two keys, and relay every message after decrypting and re-encrypting it. Diffie–Hellman has the same weakness, and classically it is answered with digital signatures: a piece of information only the claimed party could produce, held in a trusted public database.

The quantum version exploits a limitation rather than a capability. Alice hands Bob a bit prepared in one of $\mathbf{e}_0$, $\mathbf{e}_1$, or $\frac{1}{\sqrt{2}}(\mathbf{e}_0 + \mathbf{e}_1)$, without saying which. Bob measures and reads, say, 0. That is consistent with the first preparation and with the third, and nothing in the single bit tells him which he received. Having failed to learn the preparation, he cannot reproduce it; this is the no-cloning theorem.⁹ On our reading the reason is familiar. The preparation is a fact about the ensemble and the history, and the bit Bob holds carries a configuration, not a history.

Gottesman and Chuang turn that into a signature scheme.¹⁰ The public key is a string of such bits, published freely, safe to hand out because nobody who measures them can work

⁹W. K. Wootters and W. H. Zurek, "A single quantum cannot be cloned," *Nature* 299 (1982), 802–803.

¹⁰D. Gottesman and I. Chuang, "Quantum Digital Signatures," arXiv:quant-ph/0105032 (2001).

backward to the preparations. The private key is the list of preparations itself, known only to its owner. Asked to prove his identity, Bob prepares a fresh string by his instructions and sends it; Alice fetches the published string and checks that the two agree.

She can check without learning either, using the *swap test*. Its middle gate is a controlled swap, or Fredkin gate, written CSW. It exchanges the values of the two lower bits when the most significant bit, the control, reads 1, and leaves them alone when it reads 0:

A	B	C	A CSW BC
0	0	0	000
0	0	1	001
0	1	0	010
0	1	1	011
1	0	0	100
1	0	1	110
1	1	0	101
1	1	1	111

$$\text{CSW} = \begin{pmatrix} 1 & 0 & 0 & 0 & 0 & 0 & 0 & 0 \\ 0 & 1 & 0 & 0 & 0 & 0 & 0 & 0 \\ 0 & 0 & 1 & 0 & 0 & 0 & 0 & 0 \\ 0 & 0 & 0 & 1 & 0 & 0 & 0 & 0 \\ 0 & 0 & 0 & 0 & 1 & 0 & 0 & 0 \\ 0 & 0 & 0 & 0 & 0 & 1 & 0 & 0 \\ 0 & 0 & 0 & 0 & 0 & 1 & 0 & 0 \\ 0 & 0 & 0 & 0 & 0 & 0 & 0 & 1 \end{pmatrix}$$

The circuit is

control	$\mathbf{e}_0$	H	$\overline{\text{CSW}}$	H	$\perp$
first bit	ψ		CSW		
second bit	φ		CSW		

Measure the control. If the two states are identical it reads 0 with certainty. If they are orthogonal it reads 0 half the time. States that are neither give something in between: comparing $\mathbf{e}_0$ against $\frac{1}{\sqrt{2}}(\mathbf{e}_0 + \mathbf{e}_1)$ reads 0 three times in four.

A single test therefore proves little, since a mismatched pair can pass by luck. Three quarters is the worst case among the three preparations above, so sixteen tests bring the chance of a forged signature passing all of them to about 1% and thirty-two bring it to about 0.01%. Longer signatures make the failure rate as small as one likes.

The scheme has a real limit, and it is the mirror image of what makes it work. No-cloning forbids reconstructing a preparation from one copy; it forbids nothing about many. Given enough identically prepared bits, an attacker can rotate each one differently before measuring and reverse-engineer the preparation from the statistics. Fresh bits are needed because one bit cannot be measured twice to the same end. The first measurement establishes maximal QMI between the bit and the apparatus, which writes that interaction into the history the bit's later behavior answers to, washing out the preparation the attacker was trying to read. A classical public key may be handed to anyone who asks, any number of times. A quantum one weakens with every copy distributed, which is a strange property for something called public.

What the three protocols share is that each leans on a different part of the same dynamics, and that none of them requires a wave, a collapse, or a particle in two places.

The pad leans on the stochastic part alone. Which value a measured bit takes is not fixed by anything in the theory, and no known law of physics predicts it. An attacker has nothing to work from, however much is known about the machine that produced the key.

The signature scheme leans on the non-Markovian part. A bit's configuration carries no record of how it was prepared, and yet the preparation governs how that bit behaves in a later test. Two people holding identically prepared bits are therefore not equally placed: the one who knows the preparation can say how a swap test will come out, and the one who

holds only the bit cannot. The private key is that gap between configuration and history and nothing more.

Key distribution leans on what a measurement does in between the two Hadamards. An eavesdropper who measures the bit in transit correlates her apparatus with it, establishing maximal QMI between the two, and that correlation is what destroys the interference the two gates depended on. Why a correlation with a record carried off elsewhere should decide what a bit does next is the subject of Chapter 8.

7.10 Quantum teleportation

Alice holds a bit in a state ψ she does not know. She and Bob also hold one bit each of a shared pair, and that pair's preparation is part of the circuit below: it is the first two gate columns. In practice those two gates are applied while Alice and Bob are together, and they separate afterwards, each carrying one bit of the pair away with them. The circuit is

$$\begin{array}{llllll} \text{Alice, unknown} & \psi & & \overline{\text{CX}} & \text{H} & \perp \\ \text{Alice, half-pair} & \mathbf{e}_0 & \text{H} & \overline{\text{CX}} & \text{CX} & \perp \\ \text{Bob, half-pair} & \mathbf{e}_0 & & \text{CX} & & \end{array}$$

Write $\psi = \alpha \mathbf{e}_0 + \beta \mathbf{e}_1$. The first two gate columns put Alice's second bit and Bob's into $\frac{1}{\sqrt{2}}(\mathbf{e}_{00} + \mathbf{e}_{11})$, and Alice's CX and Hadamard then carry the three bits to

$$\frac{1}{2}(\mathbf{e}_{00} \otimes \psi + \mathbf{e}_{01} \otimes X\psi + \mathbf{e}_{10} \otimes Z\psi + \mathbf{e}_{11} \otimes XZ\psi),$$

the first factor in each term being the pair Alice is about to read and the second what Bob's bit is left holding. Every coefficient is $\frac{1}{2}$, so each of the four readings comes up with probability $\frac{1}{4}$, and the four are the only possibilities.

Alice reads her two bits and sends the pair to Bob. Write u' for what the unknown bit reads and a' for what her half-pair reads. The primes are there because neither value is guaranteed to be the one the bit held beforehand: Alice's CX takes her half-pair as its target and so can perturb it, and her Hadamard acts on the unknown bit and so can perturb that. Either may come through unchanged on any given run; what is gone is the guarantee.

There is a pattern across the four terms. The leading bit of Alice's pair is 1 in exactly the two terms carrying a Z, and the trailing bit is 1 in exactly the two carrying an X. So whichever pair she reads, the second factor beside it is the same expression in her two bits:

$$\text{Bob's bit holds } X^{a'}Z^{u'}\psi.$$

That is what makes the protocol work, and it is why Bob is told to apply $X^{a'}Z^{u'}$: not some operator chosen to fit, but the very one the circuit has already applied to his bit. Applying it a second time undoes it, since $X^2 = Z^2 = \text{I}$:

$$X^{a'}Z^{u'} X^{a'}Z^{u'}\psi = \pm \psi,$$

where the sign is -1 only when both bits are 1, because bringing the two Xs together across the intervening Z costs one use of $ZX = -XZ$. Written out case by case, with Bob's two gate

slots holding either the gate or the identity:

$$\begin{array}{llllll}
 u'a' = 00 : & \text{Bob} & \psi & \text{I} & \text{I} & \perp & X^0 Z^0 \psi = \psi \\
 u'a' = 01 : & \text{Bob} & X\psi & \text{I} & X & \perp & X^1 Z^0 X\psi = \psi \\
 u'a' = 10 : & \text{Bob} & Z\psi & Z & \text{I} & \perp & X^0 Z^1 Z\psi = \psi \\
 u'a' = 11 : & \text{Bob} & XZ\psi & Z & X & \perp & X^1 Z^1 XZ\psi = -\psi
 \end{array}$$

The sign in the last row is a global phase in the sense of Section 3.2, which changes nothing observable. So Bob's bit ends in ψ on all four readings, for every α and β , and Alice never learned what they were.

A standard caution attaches to circuits like this one. A stochastic process needs a law of the form $\mathbf{p}^{[t+1]} = \Gamma^{[t+1 \leftarrow t]} \mathbf{p}^{[t]}$, and quantum mechanics supplies no such thing. What it supplies is the composed matrix from the start of a circuit to its end, relating two moments and saying nothing whatever about the ones between, from which the usual conclusion is that one should not speak of the bits' values in between at all.

The conclusion holds only where the circuit interferes. Section 5.6's correction to the Markov term is a sum over *pairs* of configurations, and the term belonging to a pair $k < l$ carries the factor $\sqrt{\Gamma_{jk} \Gamma_{jl} p_k p_l}$, so it can be non-zero only when both configurations carry probability and both send a branch to the same output j . The Γ there is $|U|^2$, by Section 5.6, so the branches being asked about are the non-zero entries of $|U|^2$ and not of U : replace every gate in a circuit by its squared magnitudes, run the result as an ordinary Markov chain, and ask whether any two occupied configurations ever arrive at the same place. In the tree of Section 3.4 that is two paths meeting at one outcome. Where none ever do, every such term is zero, $\mathbf{p}^{[t+1]} = |U^{[t+1 \leftarrow t]}|^2 \mathbf{p}^{[t]}$ holds step by step, and the chain built from the squared magnitudes is guaranteed to give the circuit's own probabilities. One may then speak of the configurations in between, with no appeal to Section 6.6's construction and no selection rule to defend. That a circuit of this kind has well-defined trajectories at all is reason enough to take $|U|^2$ seriously where it works.

Teleportation is such a circuit. The overlap test can be run on the whole protocol in one pass, Bob's corrections included, once those are written as gates rather than as instructions he carries out by hand. They are conditional on what Alice read, so make them gates controlled by Alice's bits: a CX controlled by her half-pair, and a CZ controlled by her unknown bit. The *controlled Z* works exactly as CX does except that the gate it conditionally applies to the target is Z rather than X:

$$\text{CZ} = \begin{pmatrix} 1 & 0 & 0 & 0 \\ 0 & 1 & 0 & 0 \\ 0 & 0 & 1 & 0 \\ 0 & 0 & 0 & -1 \end{pmatrix},$$

which negates the amplitude of the one configuration with both bits at 1 and leaves the other three alone. The matrix is symmetric, so which bit is called the control makes no difference to what it does. With the corrections written that way there is no measurement left anywhere, and the protocol is a single six-gate circuit.

Its two CX gates and its CZ are permutations: each sends a configuration to exactly one place, so $|U|^2$ has a single non-zero entry in each row and no output can have two configurations arriving at it. That is read off the matrices and assumes nothing. Only the

two Hadamards could bring paths together, and whether they do is a question about which configurations are occupied when they act. The recording is what decides it for the second of them: the CX from the unknown bit onto Alice's half makes that bit's value a function of the other two, so afterwards no two occupied configurations differ in it alone, and a Hadamard on it splits each into fresh outputs instead of merging any pair. A gate that records empties the interference sum, which is Section 6.3's condition arriving in its simplest form.

Tracing every branch settles it. Replacing each gate by its squared magnitudes and following every path that carries weight gives eight trajectories:

start	H	CX	CX	H	CX	CZ
000	000	000	000	000	000	000
000	000	000	000	100	100	100
000	010	011	011	011	010	010
000	010	011	011	111	110	110
100	100	100	110	010	011	011
100	100	100	110	110	111	111
100	110	111	101	001	001	001
100	110	111	101	101	101	101

Each of the eight ends on a different configuration, and that is the check. Repeats in the earlier columns do not bear on it, since the table lists complete trajectories: a path that divides later is written out once per continuation, so its shared beginning appears more than once, which is why the first column holds only two values among eight rows. Two paths meeting would show as the opposite shape, two rows differing at one step and agreeing at the next, and the table has none. Had any pair met, they would have agreed from that point onward and finished together, so every endpoint beyond the meeting would appear twice. Eight distinct endpoints therefore rules out a meeting at every step and not merely at the last. Every interference term is therefore zero, the single-step law is $|U|^2$ at every gate, and the whole trajectory can be read back off its endpoint: no two runs that end differently went the same way, and no two that went differently end the same.

Reading the first column against the last, every run that began with the unknown bit at 0 ends with Bob's bit at 0, and every run that began at 1 ends with Bob's bit at 1. Not on average and not with high probability: on all eight, which is every run the circuit has. Marginalizing Alice's two bits away splits the eight into two groups of four, and the split is exactly by the value the unknown bit started with.

With every term of the second summation zero, Section 5.6's law reduces at each gate to its first summation alone, which is a Markov chain with $\Gamma = |U|^2$. The circuit can therefore be written out with those matrices standing in for the gates. For the Hadamard, $|H|^2$ is Chapter 2's COIN, a fair toss either way; for the CX, $|CX|^2 = CX$, its entries being 0 and 1 already. We can therefore express the preparation protocol as a Markov chain:

Alice, unknown	p		$\overline{CX}$	COIN	$\perp$
Alice, half-pair	e₀	COIN	$\overline{CX}$	$C\overline{X}$	$\perp$
Bob, half-pair	e₀		$C\overline{X}$		

Follow one run of it, writing u for the unknown bit's value, a for Alice's half-pair and b for Bob's. Section 1.5 showed CX to be the reversible XOR: the target becomes the XOR of the two inputs, and the control is left exactly as it was.

The first COIN and the first CX prepare the shared pair. The COIN gives Alice's half a value a , and the CX then has her half as its control and Bob's as its target, so his becomes

$$b = 0 \oplus a = a.$$

From that point the two halves are guaranteed to carry the same value, and that is what they take with them when Alice and Bob separate.

The second CX does the encoding. Its control is the unknown bit and its target is Alice's half, so hers becomes

$$a' = a \oplus u,$$

while u , being the control, is left exactly as it was. The last COIN then overwrites u with a fresh value u' , so the unknown bit's original value is effectively erased. Alice reads her two bits and sends them to Bob: a fresh coin u' , and $a' = a \oplus u$.

Bob's half of the protocol is a Markov chain as well, and for the same reason. His gates are $X^{a'}$ and $Z^{u'}$, and $|X|^2$ and $|Z|^2$ are both permutation matrices, so neither can branch and neither can create an overlap. His decoding can therefore be written out in the same way, and set beside the quantum one. He begins holding b in all four cases, and each circuit either flips it or does not, so what he computes is $b \oplus 1$ or $b \oplus 0$:

$u'a' = 00 :$	Bob	b	I	I	$\perp$	$b' = b \oplus 0$
$u'a' = 01 :$	Bob	b	I	X	$\perp$	$b' = b \oplus 1$
$u'a' = 10 :$	Bob	b	I	I	$\perp$	$b' = b \oplus 0$
$u'a' = 11 :$	Bob	b	I	X	$\perp$	$b' = b \oplus 1$

Compare the two tables and they differ in one place: where the quantum one has Z in the rows for $u'a' = 10$ and 11 , the classical one has the identity, since $|Z|^2 = I$ and that gate cannot move a configuration on any run. The consequence is visible in the rows. The first and third are now the same circuit and compute the same thing, and so are the second and fourth, so the classical table has four rows and two distinct circuits. Within each pair the only difference is the leading bit u' , which therefore makes no difference to what Bob does. Classically the second half of Alice's message is waste.

The trailing bit does all of it. The X slot is filled exactly when $a' = 1$, so the two cases are a single instruction,

$$b' = b \oplus a',$$

and the two substitutions are already in hand: Bob's half is $b = a$, and Alice sent $a' = a \oplus u$. So

$$b' = a \oplus (a \oplus u) = u.$$

The a cancels and his bit holds the unknown bit's original value, on every run of the experiment. That is Section 7.9's one-time pad, with the pair's preparation handing the key to both parties while they were still together: a is the key, u is the plaintext, and a' is the ciphertext. It is uniform because a is, so it discloses nothing on its own, and it yields u only in combination with a key Bob holds and nobody else does. The eight trajectories above are that pad run through every case it has.

Computing ψ through the protocol is easy enough, and doing so tells one nothing about what the bits are up to. What makes this protocol interesting is that we are licensed to

represent it as a Markov chain without specifying any transition law at all. That license follows from the polar decomposition together with the absence of overlap, and it costs no assumption. It holds for the one basis the circuit is written on, at least; Chapter 10 takes up what happens when definite values on more than one basis are demanded at once.

Represented that way, what the bits are doing is in plain view. The values arrive correctly because they are communicated as a one-time pad, and the protocol is a teleportation and not a copying because Alice's last step erases the unknown bit.

None of this replaces the quantum mechanical formalism. The polar decomposition has two evolution laws, one for the probabilities and one for the phases, and everything above concerns the first of them taken on its own. The classical circuit accounts for the bit values and for nothing else. Carrying the phases is still required to account for the protocol in full, and in particular to say why the Z correction is there at all, since at the level of configurations it does nothing at all.

The Z is there because the protocol scrambles the phase along with the configuration, and the pad repairs only the configuration. Of Bob's two corrections, X swaps the two probabilities and negates the relative phase, while Z leaves the probabilities alone and shifts the relative phase by π . On the reading $u/a' = 10$ his bit already carries the right probabilities and a phase wrong by exactly π , so Z is the whole of the repair. Leave it out and the protocol fails. Suppose the bit Alice sent was itself prepared by a Hadamard acting on a bit at 0, so that a second Hadamard ought to bring it back to 0 with certainty, exactly as Chapter 5's pair does. With both corrections in place Bob's bit does that. With the Z omitted, on the two readings that call for it, his bit reads 1 with certainty instead.

What the two corrections restore between them is the sufficient statistic of Section 5.8, and a sufficient statistic is no record of the history. It keeps what is needed to predict the next step and discards the rest, so it can be ambiguous about what actually happened. Bob's bit has been through a sequence of operations Alice's bit never saw, and once the corrections are applied the two histories agree in the only respect the law consults. They are sufficiently equivalent, which is everything required of them and less than identical.

That is also where the non-Markovian machinery comes back. The protocol has a Markovian account and its bits can be spoken of plainly, and nothing in that account mentions a phase at all. Saying why Bob's bit goes on to behave as Alice's would have, under whatever gate comes next, means carrying the phase as well, which is the work of Section 5.6's second law and of the expansion that turns it back into history.

So the interest here lies somewhere other than in dispensing with ψ . In a circuit of this kind one can say what the bits are doing, meaningfully and in detail, without adding a transition law as a further assumption. The formalism keeps everything it had; what it gains is an account of the configurations underneath it.

Code: this chapter in Octave

```

% --- the gates ---
e0 = [1;0]; e1 = [0;1]; I2 = eye(2);
H = (1/sqrt(2))*[1 1; 1 -1];
X = [0 1; 1 0]; Z = [1 0; 0 -1];
CX = [1 0 0 0; 0 1 0 0; 0 0 0 1; 0 0 1 0];

% --- Section 7.2: correlation hidden in the phase ---
Rx = @(t) [cos(t/2) -1i*sin(t/2); -1i*sin(t/2) cos(t/2)];
psi = kron(I2,H) * CX * kron(H, Rx(pi)) * kron(e0,e0)
abs(psi).^2 % uniform: CMI = 0
abs(kron(H,I2) * psi).^2 % -> [0.5;0;0;0.5]: correlated

% --- Section 7.3: superdense coding, end to end ---
bell = CX * kron(H,I2) * kron(e0,e0);
decode = kron(H,I2) * CX;
enc = {I2, X, Z, Z*X}; % messages 00, 01, 10, 11
for k = 1:4
    mid = kron(enc{k},I2) * bell; % the pair, in the mail
    [abs(mid).^2, abs(decode*mid).^2] % mid statistics | Bob reads k-1
end

% --- Sections 7.4-7.5: the three fillings of the H-?-H sandwich ---
mB1 = kron(I2,[1 1]); % marginal of B1
mB1 * abs( kron(H,I2) * kron(H,I2) * kron(e0,e0) ).^2 % -> [1;0]
mB1 * abs( kron(H,I2) * CX * kron(H,I2) * kron(e0,e0) ).^2 % -> [.5;.5]
mB1 * abs( kron(H,I2) * CX * kron(H,H) * kron(e0,e1) ).^2 % -> [0;1]

% --- Section 7.7: Deutsch's algorithm, all four oracles, one query ---
oracle = {kron(I2,I2), kron(I2,X), CX, kron(I2,X)*CX};
for k = 1:4 % const0, const1, f(x)=x, f(x)=~x
    mB1 * abs( kron(H,I2) * oracle{k} * kron(H,H) * kron(e0,e1) ).^2
end % -> [1;0],[1;0],[0;1],[0;1]

% --- Section 7.9: the pad, BB84, and the swap test ---
p = 1; % a plaintext bit, held fixed
mean(xor(p,[0 1])) % -> 0.5: the pad hides it
abs(H*e0).^2 % -> [.5;.5]: one key bit

n = 20000; err = [0 0]; % mismatches: no Eve, then Eve
for k = 1:n
    bit = randi([0 1]); g = randi([0 1]); % Alice's bit, both gates
    for eve = 0:1
        v = e0*(bit==0) + e1*(bit==1);
        if g, v = H*v; end
        if eve

```

```

    ge = randi([0 1]); % Eve guesses the gate
    if ge, v = H*v; end
    m = rand < abs(v(2))^2; % she measures and resends
    v = e0*(m==0) + e1*(m==1);
    if ge, v = H*v; end
end
if g, v = H*v; end
b = rand < abs(v(2))^2; % Bob measures
err(eve+1) = err(eve+1) + (b ~= bit);
end
err/n % -> [0, ~0.25]: she is caught

CSW = eye(8); CSW(:, [6 7]) = CSW(:, [7 6]); % 101 <-> 110
Hc = kron(kron(H,I2),I2); % H on the control bit
mc = kron(eye(2),ones(1,4)); % marginal of the control
st = @(u,v) mc * abs(Hc*CSW*Hc*kron(e0,kron(u,v))).^2;
[st(e0,e0), st(e0,e1), st(e0,H*e0)] % -> reads 0 w.p. 1, .5, .75

% --- Section 7.10: teleportation -- a circuit whose branches never meet ---
CZ13 = eye(8); CZ13(6,6) = -1; CZ13(8,8) = -1; % Bob's Z, controlled by psi
tel = {kron(I2,kron(H,I2)), kron(I2,CX), kron(CX,I2), kron(H,eye(4)), ...
    kron(I2,CX), CZ13}; % the last two are Bob's corrections
p1 = 0.4; psi = [sqrt(1-p1); sqrt(p1)*exp(2i)];
st = kron(psi,kron(e0,e0)); feeders = 0; markov = 0;
for k = 1:6 % Section 5.6's correction needs two
    G = abs(tel{k}).^2; % supported configurations meeting at
    s = abs(st).^2 > 1e-12; % one output; count how many ever do
    for j = 1:8; feeders = max(feeders, sum((G(j,:)') > 1e-12) & s)); end
    markov = max(markov, norm(G*abs(st).^2 - abs(tel{k}*st).^2));
    st = tel{k}*st;
end
feeders % -> 1: no two branches ever meet
markov % -> 0: |U|^2 carries p by itself
E = []; % every trajectory, and where it ends
for x = 0:1
    f = 1 + 4*x;
    for k = 1:6
        G = abs(tel{k}).^2; nf = [];
        for i = f; nf = [nf, find(G(:,i)') > 1e-12)]; end
        f = nf;
    end
    E = [E; [repmat(x,numel(f),1), f'-1]];
end
rows(E) % -> 8 complete trajectories
numel(unique(E(:,2))) % -> 8: each ends on its own configuration

```

```

all(mod(E(:,2),2) == E(:,1))          % -> 1: Bob's bit equals x on every run

% the same four gates classically, with COIN in place of H
COIN = [0.5 0.5; 0.5 0.5];
cls = {kron(I2,kron(COIN,I2)), kron(I2,CX), kron(CX,I2), kron(COIN,eye(4))};
pc = zeros(8,1); pc(1) = 1-p1; pc(5) = p1; q = kron(psi,kron(e0,e0));
for k = 1:4; pc = cls{k}*pc; q = tel{k}*q; end
norm(pc - abs(q).^2)                  % -> 0: same distribution as the quantum run
ok = true;                            % the pad: b = a, a' = a XOR u, b' = b XOR a'
for u = 0:1; for a = 0:1
    ok = ok && (mod(a + mod(a+u,2), 2) == u);
end; end
ok                                    % -> 1: Bob's bit is u on every run
pol = @(v) [abs(v(2))^2, mod(arg(v(2))-arg(v(1))+pi,2*pi)-pi];
[pol(psi); pol(X*psi); pol(Z*psi)]    % X swaps p and negates delta; Z shifts
                                      % delta by pi and leaves p alone
st = kron(H*e0,kron(e0,e0));          % the Z is not optional: prepare with H,
for k = 1:4; st = tel{k}*st; end      % teleport, correct, then one more H
r = st([7 8])/norm(st([7 8]));        % the reading u'a' = 11
[abs(H*(X*Z)*r).^2, abs(H*X*r).^2]    % -> [1;0] with Z, [0;1] without it

% --- Section 7.5: the Pauli transfer matrix of a gate ---
function P = paulis(n)
    B = {eye(2), [0 1;1 0], [0 -1i;1i 0], [1 0;0 -1]};
    if n == 1, P = B; return; end
    Q = paulis(n-1); P = {};
    for j = 1:4, for k = 1:length(Q), P{end+1} = kron(B{j},Q{k}); end, end
end
function R = ptm(U, n)
    P = paulis(n); d = 2^n; m = 4^n; R = zeros(m);
    for j = 1:m, for k = 1:m
        R(j,k) = real(trace(P{j} * U * P{k} * U')) / d;
    end, end
    R(abs(R) < 1e-9) = 0;
end
ptm(H,1)          % rows and columns I X Y Z: the table shown in the text

% One entry to a column is the Clifford property. T fails it.
Sg = [1 0; 0 1i]; Tg = [1 0; 0 exp(1i*pi/4)];
col1 = @(R) all(abs(sum(abs(R)) - 1) < 1e-9);
[col1(ptm(H,1)), col1(ptm(Sg,1)), col1(ptm(CX,2)), col1(ptm(Tg,1))]
% -> 1 1 1 0: H, S and CX permute the strings; T splits one into two
ptm(Tg,1)(:,2)' % T's X column: two entries of 0.707, not one of 1

% --- Section 7.5: a Clifford circuit simulated from its Pauli strings ---
% The tableau holds 2n Pauli strings as bit vectors: x(k,j) and z(k,j)

```

```

% say which of I, X, Y, Z row k puts on bit j, and r(k) is the sign.
% Rows n+1..2n are the strings the state leaves unchanged; rows 1..n
% are their partners, which make the readout cheap. No amplitudes.
function g = gph(x1,z1,x2,z2)
    if x1==0 && z1==0, g = 0;
    elseif x1==1 && z1==1, g = z2*(2*x2-1);
    elseif x1==1 && z1==0, g = z2*(1-2*x2);
    else g = x2*(2*z2-1); end
end
function [x,z,r] = rsum(x,z,r,h,i)      % row h <- row h times row i
    s = 2*r(h) + 2*r(i);
    for j = 1:columns(x), s += gph(x(i,j),z(i,j),x(h,j),z(h,j)); end
    r(h) = mod(mod(s,4)/2,2);
    x(h,:) = mod(x(h,:)+x(i,:),2);  z(h,:) = mod(z(h,:)+z(i,:),2);
end
function [x,z,r] = gh(x,z,r,a)          % Hadamard: swap X and Z on bit a
    r = mod(r + x(:,a).*z(:,a),2);
    t = x(:,a);  x(:,a) = z(:,a);  z(:,a) = t;
end
function [x,z,r] = gs(x,z,r,a)          % S: X -> Y on bit a
    r = mod(r + x(:,a).*z(:,a),2);  z(:,a) = mod(z(:,a)+x(:,a),2);
end
function [x,z,r] = gcx(x,z,r,a,b)       % CX, control a, target b
    r = mod(r + x(:,a).*z(:,b).*mod(x(:,b)+z(:,a)+1,2),2);
    x(:,b) = mod(x(:,b)+x(:,a),2);  z(:,a) = mod(z(:,a)+z(:,b),2);
end
function [out,x,z,r] = meas(x,z,r,a)     % read bit a on the ontic basis
    n = columns(x);  p = find(x(n+1:2*n,a)==1,1);
    if ~isempty(p)
        % undetermined: a fair coin
        p = p+n;
        for k = 1:2*n
            if k~=p && x(k,a)==1, [x,z,r] = rsum(x,z,r,k,p); end
        end
        x(p-n,:) = x(p,:); z(p-n,:) = z(p,:); r(p-n) = r(p);
        x(p,:) = 0; z(p,:) = 0; z(p,a) = 1;  r(p) = (rand<0.5);  out = r(p);
    else
        % already settled by the strings
        x(2*n+1,:) = 0; z(2*n+1,:) = 0; r(2*n+1) = 0;
        for k = 1:n
            if x(k,a)==1, [x,z,r] = rsum(x,z,r,2*n+1,k+n); end
        end
        out = r(2*n+1);
    end
end
end

% H on bit 1, CX, then S on bit 2: sampled from the tableau, and

```

```

% checked against the amplitudes the tableau never computes.
S = [1 0; 0 1i]; nb = 2; N = 20000;
pq = abs(kron(I2,S) * CX * kron(H,I2) * [1;0;0;0]).^2;
cnt = zeros(4,1);
for t = 1:N
    x = [eye(nb); zeros(nb); zeros(1,nb)];
    z = [zeros(nb); eye(nb); zeros(1,nb)]; r = zeros(2*nb+1,1);
    [x,z,r] = gh(x,z,r,1);
    [x,z,r] = gcx(x,z,r,1,2);
    [x,z,r] = gs(x,z,r,2);
    b = 0;
    for k = 1:nb, [o,x,z,r] = meas(x,z,r,k); b = b*2 + o; end
    cnt(b+1)++;
end
[cnt/N, pq]          % the tableau samples the right law
0.5*sum(abs(cnt/N - pq)) % ~0.005 at this many shots

```


Chapter 8

Quantum Interference

Every argument so far has run on circuits. We will now turn to the experiment that convinced a century of physicists that particles must be waves, that Richard Feynman called the only mystery, and that any statistical account of quantum mechanics must either explain or die on. We will explain it. The explanation will use nothing beyond the machinery already built: base cases, quantum mutual information, and one hidden assumption dragged into the light.

8.1 The double slit

Consider a laser aimed at a sheet with a single narrow slit in it. On the far side of the slit is a screen, so that light can be seen where it lands. If light were a stream of classical, Newtonian particles, one would expect the particles that pass through to land in a narrow band in line with the slit. That is not what happens: the light spreads out across the screen, fanning wider than the slit itself, as if it *diffracted* from the slit, a behavior characteristic of waves.




Now consider a slight modification: a sheet with *two* slits. Rather than opening both, we first conduct two separate experiments: one with the bottom slit blocked, one with the top slit blocked. Since only a single slit is open at a time, each outcome is effectively the single-slit diffraction experiment again, centered on whichever slit is open. From these two experiments we know the probability distribution of where particles land when they pass through the top slit, and the distribution of where they land when they pass through the bottom slit.

What should happen when both slits are open? The natural expectation is that the statistical distribution on the screen will simply be the *mixture* of the two single-slit distributions: each particle goes through one slit or the other, half and half, by the symmetry of the arrangement. Since we know what each case does, the total should be the equally-weighted

average of the cases. We will refer to this expectation as the *additivity assumption*:

$$P(x) = \frac{1}{2} P(x \mid \text{top}) + \frac{1}{2} P(x \mid \text{bottom}).$$

If the experiment is conducted in the real world, the additivity assumption is not fulfilled. What appears on the screen is not the sum of the two humps but a striped pattern: alternating bands of light and darkness, including dark bands at places where either slit alone would have delivered plenty of light:



$$P(x) \neq \frac{1}{2} P(x \mid \text{top}) + \frac{1}{2} P(x \mid \text{bottom}).$$

Keep in mind that these bands show up even if a single particle is fired at a time, with the pattern reconstructed over a long period by accumulating statistics. A lone particle has nothing to interact with, so interaction between particles cannot explain them.

This fact has led many physicists (Feynman and David Deutsch prominently among them) to argue that one cannot think of quantum theory as a statistical theory at all.¹ The reasoning: if the pattern appears one particle at a time, then, logically, the particle must itself spread out as a wave, pass through both slits, and interfere with itself, in order to produce the right statistics when it reaches the screen.

That conclusion is tempting, and the rest of the chapter is given to taking it apart. To do so cleanly we trade the double slit for a smaller, finite, circuit-sized version of it, an arrangement exhibiting the same interference effect with only two paths instead of a continuum of screen positions.

8.2 The Mach–Zehnder interferometer

The Mach–Zehnder interferometer is built from a component called a *beam splitter*: a semi-transparent mirror. Half the light hitting it passes through as if it were a window; half reflects as if it were a mirror. A single particle of light meeting a perfect beam splitter thus has an equal probability of being transmitted or reflected.


In circuit terms, one bit suffices to represent which of two paths the photon is on (horizontal or vertical, values 0 and 1), and the beam splitter is the unitary operator

$$B = \frac{1}{\sqrt{2}} \begin{pmatrix} 1 & i \\ i & 1 \end{pmatrix},$$

whose square magnitudes are an exact fifty–fifty: $|B|^2 = \text{COIN}$.

¹R. P. Feynman, R. B. Leighton, and M. Sands, *The Feynman Lectures on Physics*, vol. 3 (Reading, MA: Addison–Wesley, 1965), §1-1, where the double slit is introduced as containing the “only mystery” of quantum mechanics; D. Deutsch, *The Fabric of Reality* (London: Allen Lane, 1997), ch. 2.

To construct the complete interferometer, mirrors are placed on both paths so that, whichever path the photon takes after the first beam splitter, it is reflected back to a common location, where a second beam splitter is placed:



If a fifty-fifty splitter is all a beam splitter is, the second one should re-randomize whatever arrives and the photon should exit each port half the time. In practice, the photon *always* exits the vertical port. The circuit and the arithmetic:

$$\mathbf{e}_0 \quad \text{B} \quad \text{B} \quad \perp \quad |\text{B B } \mathbf{e}_0|^2 = \left| \frac{1}{2} \begin{pmatrix} 0 & 2i \\ 2i & 0 \end{pmatrix} \begin{pmatrix} 1 \\ 0 \end{pmatrix} \right|^2 = \begin{pmatrix} 0 \\ 1 \end{pmatrix}.$$

The reader will recognize the two-Hadamard puzzle of Chapter 5: same gate twice, coin becomes certainty. The Mach-Zehnder interferometer *is* the H-H sandwich; everything we have said about the sandwich is about to apply.

Now block one of the two intermediate paths with a barrier and collect statistics on where the photon lands. What is found: a quarter of the time the photon exits the vertical port, a quarter of the time the horizontal port. The two do not sum to one, because half the time the photon hits the barrier and is absorbed, and nothing is detected at all. If we post-condition on the cases where a particle *is* detected, the exit statistics are uniform. This holds whichever path the barrier is placed on:

$$P(x \mid \text{top}) = \begin{pmatrix} 0.5 \\ 0.5 \end{pmatrix} \quad P(x \mid \text{bottom}) = \begin{pmatrix} 0.5 \\ 0.5 \end{pmatrix},$$

where “top” and “bottom” name the surviving path. Average these with equal weights and the additivity assumption predicts a uniform distribution when both paths are open. Yet, as we have already seen, with both paths open the outcome is a *degenerate* distribution (the photon always exits vertically):

$$\begin{pmatrix} 0 \\ 1 \end{pmatrix} \neq \frac{1}{2} \begin{pmatrix} 0.5 \\ 0.5 \end{pmatrix} + \frac{1}{2} \begin{pmatrix} 0.5 \\ 0.5 \end{pmatrix}.$$

The additivity assumption fails here exactly as it fails at the double slit, and one may be inclined to run the same argument: the particle must have traversed both paths as a wave and interfered with itself at the second beam splitter. Indeed, since this system is a single bit, one *could* even map its two amplitudes onto a wave on the two physical paths; the trick works, as it always does, for one particle. The next section carries that out in full, because it goes further than one might expect, and the section after it says where it stops.

8.3 The wave picture drawn out

The state has two entries and the interferometer has two arms, so the obvious thing to try is a wave on each arm. Let the wave on arm j be an ordinary electromagnetic wave,

$$E_j(x, t) = |\psi_j| \cos(kx - \omega t + \arg \psi_j),$$

with $|\psi_j|^2$ the intensity and $\arg \psi_j$ the phase. Nothing quantum has been assumed here. These are the waves of classical optics, and the question is how far they carry.

They carry further than one might expect. Start with all the light on the horizontal arm and apply the first beam splitter:


$$B\mathbf{e}_0 = \frac{1}{\sqrt{2}} \begin{pmatrix} 1 \\ i \end{pmatrix}.$$

Each arm now has amplitude $1/\sqrt{2}$ and therefore half the intensity, which is what a half-silvered mirror does to a beam. The phases differ. The transmitted wave keeps $\arg \psi = 0$; the reflected wave has $\arg \psi = \pi/2$, a quarter cycle behind. That quarter cycle is the i in the off-diagonal entries of B , and reflection really does impose it.

At the second splitter each of those two waves splits again, so four waves leave it, two into each port. Transmission carries the phase through and reflection adds another quarter cycle:

arriving from	by	leaving by	amplitude, phase
horizontal arm	transmission	horizontal port	$\frac{1}{2}$, 0
vertical arm	reflection	horizontal port	$\frac{1}{2}$, π
horizontal arm	reflection	vertical port	$\frac{1}{2}$, $\pi/2$
vertical arm	transmission	vertical port	$\frac{1}{2}$, $\pi/2$

The two waves leaving by the horizontal port are half a cycle apart. The two leaving by the vertical port are in step. Drawn out below, with the sum of each pair under the rule:



Half a cycle apart, the two cancel exactly and the horizontal port stays dark. In step, they add to amplitude 1 and intensity 1, and all the light leaves by the vertical port. That is $|B B \mathbf{e}_0|^2 = (0, 1)$, arrived at without writing down a quantum state.

The barrier case comes out too. Block the vertical arm and only one wave reaches the second splitter, with amplitude $1/\sqrt{2}$. It splits into two of amplitude $1/2$, one per port, and with nothing to cancel against, each port receives intensity $1/4$. Those are exactly the quarter-and-quarter statistics of the previous section, with the remaining half absorbed at the barrier.

So the wave picture is not a loose analogy here. It is quantitative, it uses only classical optics, and it gets every number right. This is why the conclusion that the photon is a wave has been so hard to dislodge.

8.4 Why the wave cannot be in the world

Bell stated the view being examined here as clearly as anyone. Setting out de Broglie's answer to the question of wave or particle, he wrote:

Is it not clear from the smallness of the scintillation on the screen that we have to do with a particle? And is it not clear, from the diffraction and interference patterns, that the motion of the particle is directed by a wave? De Broglie showed in detail how the motion of a particle, passing through just one of two holes in the screen, could be influenced by waves propagating through both holes.²

The phrase to fix on is Bell's last one: *waves propagating through both holes*. That is what this section denies. Quantum mechanics itself is not in dispute; its states, its evolution and its predictions all stand. What is denied is the propagating. No wave passes through the two holes, or down the two arms, or through the arrangement at all, because no wave in space can carry what the quantum state carries.

The mapping onto a wave in space works for one particle and no more, and the trouble starts earlier still, before a second particle is anywhere in view.

Notice what Section 8.3 chose. Its circuit has a deterministic answer, every photon leaving by the vertical port, and the wave account managed it. Now take the second beam splitter away. Each photon leaves by whichever arm it is on, the statistics are an even split, and the wave account says what one would expect of a wave: half the intensity travels down each arm, and both halves arrive.

A detector reports nothing of the kind. Half an intensity on an arm is half a chance of finding the photon on that arm. Find the photon on one arm and leave the wave on the other where it is: the account still gives half a chance of finding a photon there as well, and a detector placed there never fires. Leave the wave at half intensity on the arm where the photon was found, and the account gives half a chance of its being gone when that arm is checked again; it is always still there.

So the wave account has to be supplemented in both places at once. Something must remove the half-intensity wave from the arm where nothing was found and raise the other to full strength, at the moment of detection and at no other. That supplement is the *collapse postulate*, and every difficulty belonging to it belongs here: when the contraction happens, what counts as a detector, and how a wave spread across two arms comes to be wholly within one of them with nothing traveling between. This is the measurement problem, and the wave picture runs into it as soon as the outcome is not already certain. The supplement is not available in any case. Section 9.5 shows that a collapse which really happens is a different theory, departing from quantum mechanics in ways detectable in principle, so it cannot belong to any reading of quantum mechanics as it stands.

The second difficulty is the one just named, and with the picture drawn out in full it can be made exact.

The interferometer has two places a photon can be, the two arms between the splitters, and that is what let the wave account fit it. Call them $x = 0$ and $x = 1$. Other arrangements offer more places, so write M for however many a given one provides. A wave carries one value

²J. S. Bell, "Six possible worlds of quantum mechanics" (1986), in *Speakable and Unsayable in Quantum Mechanics*, second edition (Cambridge: Cambridge University Press, 2004), p. 191.

at each place, so it supplies M numbers. The state of N particles spread over those places has M^N amplitudes, one for every way of putting all of them somewhere at once. Places and configurations are different things, and configurations outnumber places. A single wave is enough only at $N = 1$, the case Section 8.3 worked.

How many waves the picture needs is not left open. A particle can be put into superposition over every place at once, and enough beam splitters will arrange it however many arms the interferometer is given, so each place has to carry an amplitude for that particle. Where the particle certainly does not go, that amplitude is zero, which is a value like any other. So the picture requires one amplitude per particle per place, NM of them, and this is forced rather than chosen: having waves in space is having exactly that.

Now count what sits at a place. Each particle's wave has a value at every place, since splitting into two beams is precisely that: a value on this arm and a value on that one. So every particle leaves one number at each place, and N particles leave N .

Count what a place has to account for. The state carries one amplitude for every way of placing all N particles at once, M^N in all, and every one of those has to be recoverable from the numbers sitting at the places. Shared out over the M places, each is answerable for M^{N-1} .

The question is therefore whether N numbers can do the work of M^{N-1} . For one particle that is one against one. For two particles in two places, two against two. For three particles in two places, three numbers against four.

The buckets below are the two places of the interferometer. Each square in a bucket is one of the values that place has to account for, shaded when a wave supplies it and dashed when nothing does.

one particle			1 needed, 1 supplied
	$x = 0$	$x = 1$	
2 particles			2 needed, 2 supplied
	$x = 0$	$x = 1$	
3 particles			4 needed, 3 supplied, one short
	$x = 0$	$x = 1$	

For four particles it is eight against four, for six it is thirty-two against six, and the shortfall grows without limit after that, since one side doubles with every particle added while the other goes up by one.

Read the other way round, the shortfall is a fact about ψ : it carries more numbers than any collection of waves in space can hold.

Two arms is the most forgiving case there is. The counts agree only for a single particle, whatever M may be, and for two particles in exactly two places. Three places and two particles already fall short, six waves against nine; four places and two, eight against sixteen. The interferometer is where the luck stretches furthest, which is part of why the wave reading held its ground so long.

There are three ways to try to close the gap.

One is to stop tying the waves to particles and simply put as many values at each place as the state needs, M^{N-1} rather than N . That is enough by construction, so ask instead what tells them apart at the same point. With a wave per particle the answer was easy: this one belongs to that particle. Here there is no such answer. The k th value would have to mean the amplitude at this place given that the other particles are arranged in configuration k , so its label is a configuration of the rest rather than anything in space, and how many labels there are depends on how many particles happen to exist elsewhere. Add a particle on the far side of the room and the number of values at every point doubles. A field indexed that way has the structure of the state and not of space, and what it describes is configuration space over again, reached by a longer road.

Another is to say that ψ is one field rather than one per particle. That fares worse. A single field carries one value per place, M numbers against M^N ; at two particles in two places it supplies two where four are wanted, and the gap widens with every place added. It fails at every M the moment N exceeds one.

The last is to cut down which configurations count. Perhaps no two particles may occupy one place, or perhaps the particles cannot be told apart, so that swapping them gives no new configuration. Either way the state is smaller than M^N , and it is fair to ask whether the smaller number fits. Sometimes it does. Two particles that may not share, among three places, have exactly six configurations and exactly six waves.

The fit does not survive, and the reason is visible in the count itself. Placing N particles in M places with no sharing leaves

$$\frac{M!}{(M-N)!} = M(M-1)(M-2)\cdots(M-N+1)$$

configurations: a product of N factors, each near M . Treating the particles as interchangeable divides that by a fixed number and letting them share multiplies it back up, but the shape does not change. It is always N factors of roughly M , growing like M^N , while the waves grow like M alone. So for any N above one the configurations overtake the waves once there are places enough: two particles that cannot share and cannot be told apart match at five places and fall behind at six; allow them to share and they match at three and fall behind at four. A further embarrassment waits in the restricted counts. Fifteen configurations over six places is two and a half apiece, and a field cannot carry half a value at a point.

The shape of this failure has been met once already. Section 3.4 drew a single bit as a point on a sphere, a picture that fits because a bit's state and a point on a sphere both carry two degrees of freedom, and noted that a sphere for each bit runs out for the same reason a wave for each particle runs out here: what is drawn grows one per bit while what must be described doubles. Spheres, waves and the marginals of Section 2.2 all fail at the point where the parts stop being independent, and they fail because a joint description that does not factorize is not recoverable from descriptions of the parts.

Three photons in an interferometer fits on a bench, and a wave in space has already run out before reaching it.

One route remains for anyone who wants the wave to be real. Stop trying to fit the wave into the world, and let the world sit inside the wave, with the M^N amplitudes describing a space of their own in which the three-dimensional world is a pattern rather than the container. That is the interpretation called Many Worlds, and it does not escape the difficulty either: Section 9.1 finds that the worlds themselves are not settled by the state. What has already

been surrendered, before any of that, is the wave as a thing in space, which was the whole attraction of the picture.

A reader who knows Feynman's formulation may object that the argument so far is about waves, and that a sum over paths dispenses with waves altogether.³ It does, and it does not help. The sum is built by slicing a circuit and inserting every intermediate configuration between one gate and the next, which turns a single amplitude into a sum of products of matrix elements, one product per path. The amplitude trees of Chapters 3 and 7 are that construction drawn out, and Feynman's integral is its continuum limit. Look at what the nodes of the Chapter 7 tree carry: 00, 01, 10 and 11, the joint configurations of both bits. A path there is a history of the pair, and for N particles a path is a history of all N positions at once. The sum runs over configuration space exactly as the state does, so the counting above applies to paths without a word changed.

Reading those paths as real alternative histories does not place them in space either. The histories a run does not follow have to be somewhere, and the only room left is across many worlds, which is the same trade under another name. Quantum field theory, replacing paths of positions with paths of field configurations, repeats the problem for field configurations.

The wave picture fails earlier still, and for a reason the counting cannot show. Go back to two particles in two places, where the waves were just enough and the count raised no objection at all. Those four field amplitudes are not free among themselves. Two of them belong to the first photon and two to the second, and the picture offers nothing further. Where the first photon turns up depends on its own wave and on that alone; where the second turns up depends on its own.

One might expect the two waves to work on each other, so that the pair delivers what a joint description delivers. Section 2.2 closed that route. Three joint distributions over two bits were set out there with identical marginals, and recombining the marginals returned the uncorrelated joint every time. Nothing in the separate descriptions recorded the correlation. That is why the probability vector has to be the size it is, and amplitudes are no different: a wave apiece is a description apiece, and a joint that does not factorize cannot be assembled from them. So the product is all the picture can deliver, and the chance of two independent things both happening is the product of their separate chances:

$$P(ab) = |\psi_A(a)|^2 |\psi_B(b)|^2.$$

Take the pair in the state $\frac{1}{\sqrt{2}}(\mathbf{e}_{00} + \mathbf{e}_{11})$, whose statistics are $P(00) = P(11) = 1/2$ and $P(01) = P(10) = 0$: the two photons are always on the same arm. Each one separately is on either arm half the time, so $|\psi_A(0)|^2 = |\psi_B(1)|^2 = 1/2$ whatever the assignment, and the factorized formula returns

$$P(01) = \frac{1}{2} \cdot \frac{1}{2} = \frac{1}{4},$$

where the true answer is zero. No choice of the four numbers repairs this, because the marginals fix them and the marginals are uniform.

Put this back where the counting failed. Three photons on two arms have eight configurations, so the joint carries eight amplitudes, while the wave picture gives each photon its own wave and carries six. The two numbers are close enough to invite the thought that one could

³R. P. Feynman, "Space-Time Approach to Non-Relativistic Quantum Mechanics," *Reviews of Modern Physics* 20 (1948), 367–387.

be worked into the other. Going down always works: the joint says where each photon is likely to turn up, and many different joints agree there, the way correlated and anti-correlated bits gave the same marginals. The six is a summary with things left out. Going up is one fixed recipe, the Kronecker product of the three waves, and it reaches exactly the joints that factorize. The pair above was already outside that set, and a third photon does not make room for it.

One further thing goes wrong, and it goes wrong in the single-particle case the counting concedes. Take the interferometer of Section 8.3, where the map fits, and choose any two phases θ_0 and θ_1 , one per configuration. Write $D = \text{diag}(e^{i\theta_0}, e^{i\theta_1})$ and replace every gate U in the circuit by $DUD^\dagger$. Preparation is untouched, since $D\mathbf{e}_i$ is $\mathbf{e}_i$ again up to a phase. Measurement is untouched too, since conjugating the projector $\mathbf{e}_j\mathbf{e}_j^\dagger$ of Section 3.8 by a diagonal matrix returns it unchanged: the two phases $e^{i\theta_j}$ and $e^{-i\theta_j}$ meet and cancel. Every probability is untouched as well, because the conjugation multiplies the entry in row j and column i by $e^{i\theta_j}e^{-i\theta_i}$ and leaves its magnitude alone. Nothing observable has moved.

What has moved is the offset between the arms. Conjugating the beam splitter by $\text{diag}(1, e^{i\varphi})$ carries the reflected wave from phase $\pi/2$ to phase $\pi/2 + \varphi$, and φ was ours to pick. The quarter cycle that Section 8.3 read off the i in B can be set to any value whatever by a redefinition no experiment can detect, the later gates being conjugated too and absorbing it. What survives is what happens around a closed circuit, one that leaves a configuration and comes back to it.

The global phase of Section 3.2 was unobservable too, and it was never claimed to be physical. A global phase multiplies both entries at once, so the relative phase between the arms stays where it was, and that is what interference depends on. Galilean relativity works the same way: a change of frame alters every absolute velocity, but absolute velocity was never physical, and relative velocity, which is, comes through unchanged. The redefinition above moves the relative phase itself, so the part of the phase the wave picture needs to be real has no fixed value either.

A physical object has definite properties, and the offset between two waves would be one of them. It has no definite value: two accounts assign it differently and no experiment tells them apart. Either the wave is not a physical object, or it is one and the account of it is incomplete, naming a quantity the account cannot fix.

The same structure appears for a charged particle in an electromagnetic field, where the loop integral of the potential is measurable and the potential at a point is not, and it has been raised as an objection to wave function realism on those grounds.⁴ The finite version needs no electromagnetism at all.

Three failures, then, and they do not sit at the same depth. The first says the wave account is incomplete, since a collapse postulate has to be added before it can describe an outcome that is not certain. The second says it runs out of room almost at once, and cannot be carried on without giving up three-dimensional space, which was the whole of its appeal. The third says that even in the one circuit where it reproduces every number, the offset it puts between the arms is a choice of notation.

What holds the account together is a coincidence of counting. For one particle the state carries exactly one amplitude per place, which is what a wave in space supplies, so the optics

⁴M. Sanchioni, “Exploring the Interplay Between Wave Function Realism and Gauge Symmetry Interpretations in Quantum Mechanics,” *Foundations of Physics* 54 (2024), article 59.

can be read straight off the state. That is why it fits the double slit and the interferometer alike: both are run one particle at a time. Add a second and the alignment is gone. An account that holds in the smallest case and fails in the general one cannot be the general account, and what has to be refused is the step from a convenience of the simplest experiments to a picture of quantum theory as a whole.

8.5 Interaction-free measurement

Recall the two facts established in Section 8.2: with no barrier, the photon is always detected at the vertical port; with a barrier present, the photon is sometimes detected at the horizontal port. Put them together and notice what follows. A detection at the horizontal port *only ever happens when a barrier is present*. So if we detect the photon at the horizontal port, we can be certain a barrier is present between the beam splitters.

But if we detected the photon at all, it was not absorbed, so it did not touch the barrier. It took the open path, end to end. Yet its arrival at the horizontal port certifies the barrier's presence. We have detected the presence of an object using a particle that demonstrably never interacted with it. This “interaction-free measurement” was first pointed out by the physicists Avshalom Elitzur and Lev Vaidman,⁵ and it is a reproducible laboratory fact.

Carlo Rovelli describes meeting it in person. Visiting Anton Zeilinger's laboratory in Innsbruck, he was shown a table carrying a laser, mirrors and two detectors, arranged exactly as above. With both paths open, every photon arrived at one detector and none at the other. Putting his hand into one of the paths sent half of them to the detector that had been dark. He asks how blocking one path can cause photons traveling along the other to change where they land, and answers: nobody knows.⁶

He does not name the effect, but that is what he was looking at. A photon reaching the formerly dark detector cannot have touched his hand, since a photon that touches it is absorbed and never arrives at all. Its arrival certifies that the hand is in one of the paths, and it certifies this without having been anywhere near it. Whether anybody knows why is the business of the next three sections.

What this proves is that the presence or absence of a barrier on a path the photon *does not traverse* can influence the photon's statistical behavior. (Assuming the photon was on one path or the other at all, which Section 10.6 returns to.) That is rather counterintuitive, since we expect a particle to be influenced only by what it directly and locally interacts with. Yet the experiment leaves no wiggle room: the photon's statistics at the second beam splitter have a non-local dependence upon the presence or absence of an object at a distance, on a path the photon did not take.

8.6 The hidden assumption

Once we take this into account, we can see exactly why the additivity assumption is fallacious, and the reason has nothing to do with waves.

⁵A. C. Elitzur and L. Vaidman, “Quantum mechanical interaction-free measurements,” *Foundations of Physics* 23 (1993), 987–997.

⁶C. Rovelli, *Helgoland: Making Sense of the Quantum Revolution*, trans. E. Segre and S. Carnell (New York: Riverhead Books, 2021), ch. 1.

Let us make the hidden conditions explicit. We will write BT for “a barrier is present on the top path” and BB for “a barrier is present on the bottom path.” The single-slit-style experiments that taught us $P(x \mid \text{top})$ and $P(x \mid \text{bottom})$ were conducted *with a barrier present* on the other path, since that is how one forces the photon onto a known path. So what the additivity assumption actually asserts is

$$P(x \mid \neg BB, \neg BT) = \frac{1}{2} P(x \mid \text{top}, BB) + \frac{1}{2} P(x \mid \text{bottom}, BT) :$$

the statistics with *no* barriers anywhere are predicted by summing statistics gathered with a barrier present. Dropping the barrier conditions from the right-hand side is legitimate only if the photon’s behavior does not depend on a barrier on the path it did not take:

$$P(x \mid \text{top}, BB) = P(x \mid \text{top}, \neg BB) \quad P(x \mid \text{bottom}, BT) = P(x \mid \text{bottom}, \neg BT).$$

And that is precisely what interaction-free measurement refutes. The conditions do not drop.

Deutsch’s version of the experiment has the same structure with four slits in place of two: opening a second pair darkens a point the first pair had lit, from which he concludes that “something must be coming through that second pair of slits” to interfere with the light from the first.⁷ That is the additivity assumption again, in a four-slit form, and it carries the same hidden condition: that a slit can matter to the photon only by sending something along it.

This diagnosis is not ours alone. The philosopher Tim Maudlin, examining Feynman’s claim that the double slit is impossible to explain classically, separates the reasoning into two propositions: that each particle passes through exactly one slit (Feynman’s Proposition A, in Maudlin’s telling), and that a particle’s later behavior does not depend on whether the other slit is open (Maudlin’s own Proposition B). Only the second does the anti-statistical work, and, as he observes, “Proposition B is not a proposition of classical physics, classical probability theory or classical logic.”⁸ It is an assumption about the *laws*: his Proposition B is precisely what we have been calling the additivity assumption, and what the experiments refute is Proposition B, not Proposition A.

The point can be made constructive: fill in the correct values, and additivity is restored. With a barrier present, a detected photon took the open path and exits uniformly, which gives the first column below. With no barrier present, the photon is always detected at the vertical port; since over many runs some of those photons traversed the top path and some the bottom (we mean that literally: in each run the photon has a path), the vertical-port certainty fills in the second column:

$$\begin{aligned} P(x \mid \text{top}, BB) &= \begin{pmatrix} 0.5 \\ 0.5 \end{pmatrix} & P(x \mid \text{top}, \neg BB) &= \begin{pmatrix} 0 \\ 1 \end{pmatrix} \\ P(x \mid \text{bottom}, BT) &= \begin{pmatrix} 0.5 \\ 0.5 \end{pmatrix} & P(x \mid \text{bottom}, \neg BT) &= \begin{pmatrix} 0 \\ 1 \end{pmatrix}. \end{aligned}$$

Now sum the *correctly matched* cases, namely the no-barrier distributions, since the both-paths-open experiment has no barriers:

$$P(x \mid \neg BB, \neg BT) = \frac{1}{2} P(x \mid \text{top}, \neg BB) + \frac{1}{2} P(x \mid \text{bottom}, \neg BT) \quad \begin{pmatrix} 0 \\ 1 \end{pmatrix} = \frac{1}{2} \begin{pmatrix} 0 \\ 1 \end{pmatrix} + \frac{1}{2} \begin{pmatrix} 0 \\ 1 \end{pmatrix}.$$

⁷Deutsch, *The Fabric of Reality*, ch. 2.

⁸T. Maudlin, *Philosophy of Physics: Quantum Theory* (Princeton: Princeton University Press, 2019), ch. 1.

Additivity holds exactly, provided we stay among distributions conditioned on the same experimental configuration. Hence it is simply not the case that the Mach–Zehnder interferometer, and by extension the double slit, requires us to suppose that the particle takes both paths as a wave. The statistics are explained statistically, once we properly condition on the presence or absence of a barrier on the path the photon does not traverse. In each individual run, the photon takes one path; the *law* governing its behavior at the second beam splitter consults more than the photon’s location. This is not news to a reader of Chapter 5, but we can now do better than analogy. Why should the presence of an untouched barrier matter? It seems arbitrary. It is not, and the machinery of Chapter 6 tells us exactly why.

8.7 Where the base case comes from

The single bit of Section 8.2 is a convenience, and the next step needs a physical description instead. That encoding works because the experiment was restricted to one photon traveling on one of two paths, and under that restriction the arrangement has one bit of freedom left: which path. Nothing in the interferometer enforces the restriction. Both input ports accept light, and nothing stops a photon entering each of them.

Describing the arrangement rather than the restriction means asking one question of each path, *does this path carry a photon*, which takes two bits. The gain is immediate. A photon absorbed by a barrier is on *neither* path, a configuration the one-bit encoding has no room for and the two-bit encoding names as 00.

The barrier needs a bit of its own, and the reason is bookkeeping that the formalism already enforces. Every gate is unitary and therefore reversible, so a circuit destroys no information and creates none: run it backward with Section 3.3’s time-reverse and the input returns. A photon that is absorbed has not left the account, then. It has deposited its energy somewhere, and that somewhere is the barrier, whose particles are altered by the encounter. Writing the barrier down is what keeps the books balanced, and one question with two answers, disturbed or not, is one more bit. Leaving it out is what makes the barrier’s effect look like a collapse imposed on the state from outside, rather than something the circuit does. So we expand to three bits: the first two for the presence of a photon on the top and bottom paths, the third for whether the barrier has been perturbed.

The beam splitter, now acting on two path-occupancy bits, is represented by the *Givens operator* with angle $\pi/4$:

$$G(\theta) = \begin{pmatrix} 1 & 0 & 0 & 0 \\ 0 & \cos \theta & -\sin \theta & 0 \\ 0 & \sin \theta & \cos \theta & 0 \\ 0 & 0 & 0 & 1 \end{pmatrix},$$

a rotation confined to the one-photon configurations 01 and 10, since photons are neither created nor destroyed, merely moved between paths. The Givens operator is not a new primitive: it can be constructed from the SWAP, CX, and R_Y gates of Chapter 3, as shown in the circuit below, which the reader may verify by multiplying out the matrices (the chapter’s code checks it at nine angles spanning a full turn):

$$\begin{array}{ccccccc} \mathbf{e}_0 & \text{SW} & \overline{\text{CX}} & R_Y(-\theta) & \overline{\text{CX}} & R_Y(\theta) & \overline{\text{CX}} \quad \perp \\ \mathbf{e}_0 & \text{SW} & \text{CX} & & \overline{\text{CX}} & & \overline{\text{CX}} \quad \perp \end{array}$$

The barrier is a SWAP between the bottom-path bit and the barrier bit: if a photon is on the barrier's path, the 1 is swapped out of the path and into the barrier, so no photon remains on either path, and the barrier is perturbed. If not, nothing happens. The full circuit, with the photon injected on the top path:

$$\begin{array}{ccccccc} \mathbf{e}_1 & G(\pi/4) & & G(\pi/4) & \perp \\ \mathbf{e}_0 & \overline{G}(\pi/4) & \text{SW} & \overline{G}(\pi/4) & \perp \\ \mathbf{e}_0 & & \text{SW} & & \perp \end{array}$$

Computing the full statistics (the chapter's code checks every line):

$$|(G(\pi/4) \otimes \mathbf{I}_2)(\mathbf{I}_2 \otimes \text{SW})(G(\pi/4) \otimes \mathbf{I}_2)(\mathbf{e}_1 \otimes \mathbf{e}_0 \otimes \mathbf{e}_0)|^2 = \begin{pmatrix} 0 \\ 0.5 \\ 0.25 \\ 0 \\ 0.25 \\ 0 \\ 0 \\ 0 \end{pmatrix}.$$

A 50% chance of 001 (no photon on either path and a perturbed barrier, the absorption cases) and 25% each of 010 and 100, the photon exiting on one or the other path: exactly the laboratory statistics of Section 8.2.

Now freeze the circuit at the intermediate time, after the barrier column and before the second beam splitter, and interrogate it with Chapter 4's tools. The joint state there is $\frac{1}{\sqrt{2}}(\mathbf{e}_{100} - \mathbf{e}_{001})$. Marginalizing out the barrier, the path bits are 00 or 10, each with probability 0.5: either the photon is gone (blocked) or it is on the top path. It is certainly not on the bottom path, which is the barrier's path. Compute the density matrix and take partial traces, and the quantum mutual information between the top path and the barrier comes out

$$I(\text{top} : \text{barrier}) = S(\rho^{\text{top}}) + S(\rho^{\text{barrier}}) - S(\rho^{\text{top,barrier}}) = 1 + 1 - 0 = 2 :$$

maximal. The barrier has perfectly recorded whether the photon is on the top path: absorbing the photon when it came its way is a recording, and the barrier does not interact with the photon again. Both conditions of Chapter 6 are met: *the moment after the barrier is a valid base case* for the marginal statistics of the paths. (The bottom-path bit is in a degenerate distribution at this moment, which is the other kind of base case; every bit the second splitter touches has had its history restarted.)

And a base case licenses unistochastic dynamics from that moment on: the marginal statistics after the second beam splitter should be the plain Markov step $|G(\pi/4)|^2$ applied to the marginal statistics before it. Check:

$$|G(\pi/4)|^2 \begin{pmatrix} 0.5 \\ 0 \\ 0.5 \\ 0 \end{pmatrix} = \begin{pmatrix} 1 & 0 & 0 & 0 \\ 0 & 0.5 & 0.5 & 0 \\ 0 & 0.5 & 0.5 & 0 \\ 0 & 0 & 0 & 1 \end{pmatrix} \begin{pmatrix} 0.5 \\ 0 \\ 0.5 \\ 0 \end{pmatrix} = \begin{pmatrix} 0.5 \\ 0.25 \\ 0.25 \\ 0 \end{pmatrix},$$

which is precisely the exact answer computed from the full circuit. With the barrier present, the second beam splitter is a genuine coin, because the barrier resets the photon's history, exactly

as the recording CX resets the interrupted Hadamard pair in Chapter 6. The interferometer with a barrier *is* the H–CX–H circuit; the interferometer without one *is* the H–H circuit.

Now remove the SWAP, eliminating the barrier, and run the same test. The intermediate marginal statistics of the paths are now $(0, 0.5, 0.5, 0)$: no photon is ever absorbed, so the two one-photon configurations split the probability evenly between them. But no recording has occurred: the QMI between either path and the rest of the world is zero, no base case exists at the intermediate time, and applying $|G(\pi/4)|^2$ to those intermediate marginals predicts $(0, 0.5, 0.5, 0)$ again, since the uniform mixture of the two one-photon configurations is a fixed point of the splitter’s unistochastic matrix. The true answer, from the full circuit, is that the photon exits deterministically on one port:

$$\mathbf{p}_{\text{paths}}^{[3]} = \begin{pmatrix} 0 \\ 1 \\ 0 \\ 0 \end{pmatrix} \neq |G(\pi/4)|^2 \mathbf{p}_{\text{paths}}^{[2]}.$$

Interference, in our terms, is nothing more than this inequality: the failure of the single-step Markov guess at a moment that is *not* a base case. Chapter 5’s second summation is alive and steering, because no interaction has altered the history it feeds on.

So the question “why must we condition on the barrier?” has a precise answer. The defining characteristic of a base case for a bit is that it shares maximal QMI with some other system it will not meet again, or that it is in a degenerate distribution. The latter is really the former, with the observer playing the role of the recorder. The barrier’s presence satisfies the condition; its absence does not. The presence or absence of the barrier therefore changes which law governs the photon’s statistics at the second beam splitter (Markov step or interference step), and so it must be conditioned on, whether or not the photon ever touches it.

One consequence of that condition deserves stating on its own: nothing in it mentions measurement. What licenses the base case is maximal QMI with a system the photon will not meet again. A barrier supplies it by absorbing, a detector supplies it by recording, and the two destroy the interference in the same way and for the same reason. A measuring device earns no special standing here. It is one more part of the arrangement that comes to share everything with the photon and then keeps away, and any part of the arrangement that does as much has the same effect.

8.8 Holism

QMI is a *de-localized* property of a system. Particles can establish quantum mutual information by locally interacting, but the Mach–Zehnder shows something stranger: a *non*-interaction can also impart information. If the photon interacts with the barrier, the barrier’s state is altered; if it does not, the barrier’s state is unaltered, and therefore an unaltered barrier certifies the open path just as an altered one certifies absorption. The correlation between photon and barrier is established by the structure of the whole arrangement, not by any collision. One could, in principle, measure the barrier’s particles precisely enough to learn whether the photon passed, even though the photon never went near it.

There is thus a holistic aspect to the structure of quantum information. Quantum information evolves in a way that depends upon this structural property: upon what is

correlated with what, across the entire arrangement. It thus evolves in a way that requires taking into account the whole system's configuration, not merely what a particle directly does or does not touch. Even when the particle does not interact with the barrier, the barrier's presence establishes QMI between the particle and itself, and thereby alters which law the particle's statistics follow. The presence or absence of a distant, untouched object matters, and so we must condition on it.

If "the law consults the history" was Part II's temporal thesis, this is its spatial twin: *the law consults the arrangement*. Neither requires waves, neither requires the particle to be in two places, and neither is optional; both are read directly off the mathematics of quantum theory by the machinery of Chapters 4 through 6. What they jointly retire is a picture: the picture of a particle whose statistics are fixed by its own route and nothing else. What they install instead is a picture we have been assembling since the magician: particles always somewhere definite, and laws that consult both the history and the whole arrangement.

Distributions have been doing more work than they do in ordinary probability, where they describe an ensemble and not the individual system in front of you. The magician behaved as though the system knew which distribution its input was drawn from, and a law that consults the arrangement behaves the same way. It is tempting to read that as evidence that the distribution is something the particle carries about with it, and from there to arrive at the picture of a particle taking every available route at once. An analogy will show why that step is unnecessary.

Give someone a robot's complete code and ask which route it will take through a maze. They cannot say, however perfect the listing, unless you also give them the maze. The code says what the robot does at a junction; it says nothing about where the junctions are. Code and maze together fix the route, and neither fixes it alone.

The circuit is the maze. A photon arriving at a beam splitter leaves with its path distribution uniform, and what makes it uniform is the beam splitter. The photon has no property of its own that would divide its chances in two; place the same photon in a different arrangement and the distribution differs, while nothing about the photon has changed.

The distribution therefore says nothing beyond the nature of the particle, the arrangement it is placed in, and the history of how it came to be there. Those three are enough to compute it, so the distribution is a summary of them rather than a fourth thing standing alongside. Read that way, the holism stops being strange. A distribution that depends on a distant untouched barrier is puzzling while one takes it to be a property of the particle; as an accounting of the whole arrangement, it is what one would expect.

Section 5.7 had already established this, before any interferometer was on the table. The rule derived there evolves a bare distribution, with no phase on the bit at any step, and it reproduces the interference all the same, because the only phases in it are the ones in U . Take the wavefunction out and the gate is what remains to explain the interference. A gate's phases do nothing on their own, though. A bit sent through a single Hadamard out of a base case comes out uniform, and the minus sign plays no part in that, the statistics for the step coming from $|H|^2$, which has squared the sign away. Put a second Hadamard after the first and that same sign is what returns the bit to the value it started with. What the phases record is how the gates stand in relation to one another, and what the dynamics become once they are composed, which is the holism of this section. It also says why the phases were never as troubling as the wavefunction. A phase belonging to a beam splitter stays with the beam

splitter, and nobody needs it riding along with the photon.

Nothing stops the polar decomposition from being applied to the gates themselves, and doing that makes the relation between them explicit. Let time 0 be a base case and t any later time, and cut the circuit at a time s in between. Everything before the cut composes into one operator and everything after it into another, which is the move Section 6.2 opens with, and each composed stretch decomposes into a Γ and a Θ exactly as a single gate does. The rule assembling the two into the circuit's own pair is Section 5.6's with the gates in place of the distribution:

$$\begin{aligned}\Gamma_{ik}^{[t\leftarrow 0]} &= \sum_j \Gamma_{ij}^{[t\leftarrow s]} \Gamma_{jk}^{[s\leftarrow 0]} + \sum_{j < l} 2\sqrt{\Gamma_{ij}^{[t\leftarrow s]} \Gamma_{jk}^{[s\leftarrow 0]} \Gamma_{il}^{[t\leftarrow s]} \Gamma_{lk}^{[s\leftarrow 0]}} \\ &\quad \times \cos\left((\Theta_{ij}^{[t\leftarrow s]} - \Theta_{il}^{[t\leftarrow s]}) + (\Theta_{jk}^{[s\leftarrow 0]} - \Theta_{lk}^{[s\leftarrow 0]})\right), \\ \Theta_{ik}^{[t\leftarrow 0]} &= \arctan2\left(\sum_j \sqrt{\Gamma_{ij}^{[t\leftarrow s]} \Gamma_{jk}^{[s\leftarrow 0]}} \sin(\Theta_{ij}^{[t\leftarrow s]} + \Theta_{jk}^{[s\leftarrow 0]}), \right. \\ &\quad \left. \sum_j \sqrt{\Gamma_{ij}^{[t\leftarrow s]} \Gamma_{jk}^{[s\leftarrow 0]}} \cos(\Theta_{ij}^{[t\leftarrow s]} + \Theta_{jk}^{[s\leftarrow 0]})\right).\end{aligned}$$

Each Γ above is the $|U|^2$ of the stretch of circuit its label names. Two of them are laws of the process as well, the two whose right index is the base case, by Section 6.2; $\Gamma^{[t\leftarrow s]}$ is that stretch's matrix and nothing further, since nothing says the distribution at s is a base case. The first summation multiplies the two factors together, which is the whole of what a classical account of the pair would predict, and the second is the difference between that product and $\Gamma^{[t\leftarrow 0]}$, which is what Section 6.2 meant in calling the rule not composable. The indices read as they did in Section 2.3, where a stochastic matrix was a statistical truth table with the inputs down the columns and the outputs along the rows. So $\Gamma_{jk}^{[s\leftarrow 0]}$ takes input k to output j , and $\Gamma_{ij}^{[t\leftarrow s]}$ takes that j as its own input and gives output i , the shared index being the one the matrix product sums over. A route therefore leaves configuration k at 0, occupies j at s , and arrives at i at t , and the correction sets that route against one occupying l at s instead, the cosine's argument being the two routes' phase difference.

The gates' phases enter nowhere else, and neither gate's contribution to that difference is fixed on its own. Section 8.4 replaces every gate by $DUD^\dagger$, which multiplies the entry in row j and column k by $e^{i\theta_j}e^{-i\theta_k}$. For the earlier gate the phase on k drops out of $\Theta_{jk}^{[s\leftarrow 0]} - \Theta_{lk}^{[s\leftarrow 0]}$, both routes leaving from the same k , and the difference shifts by $\theta_j - \theta_l$. For the later gate the phase on i drops out of $\Theta_{ij}^{[t\leftarrow s]} - \Theta_{il}^{[t\leftarrow s]}$, both routes arriving at the same i , and the difference shifts by $\theta_l - \theta_j$. The two shifts are one quantity with opposite signs, because j and l are configurations at the single time s that the two gates share, the one the earlier gate arrives at and the later gate leaves from. Their sum is therefore unmoved, and that sum is the cosine's entire argument.

A second redefinition is invisible in the same way: reversing the sign of every phase in the circuit leaves each Γ exactly as it was, so no prediction moves, and it sends the sum to its negative. Each of the four phases can be set to any value, and so can either of the two differences. What has a value of its own is the absolute value of the sum, which is what the

cosine reads. So the phases record a relation between the gates and no property of either one: the adjustable part of a gate's phase can be adjusted only by moving the other gate's phase the opposite way.

The cut could have been made anywhere, and nothing above depended on where it fell, so the cancellation repeats at every configuration two routes pass through between a shared start and a shared finish. That leaves any single route's phase free to take any value and the difference between two of them fixed. That is Section 8.4's closed circuit, out along one route and back along the other, and it settles what kind of thing a phase is. The bits carry none, as Section 5.7's expansion already showed. No single gate holds one either, by the cancellation just given. That is the test Section 8.4 put the wave to: a quantity any redefinition can move is no property of the thing it was attached to. A phase is a property of how the gates compose together in the arrangement. It belongs to a pair of routes through that arrangement, and a pair of routes is the smallest thing that can have one.

Chapter 6's construction keeps them exactly there. Each step's transition matrix is built from $|U|^2$, which carries no phase at all, and the phases enter only in fixing which distributions the step has to bridge. The arrangement sets the marginals; the bits hop between configurations under a matrix of plain probabilities; and nothing anywhere carries a phase through the circuit on the particle's behalf.

What invites the claim that the particle takes every route at once is that where the particle is found depends on which routes were open to it. Close one slit and the pattern on the screen changes. Each particle that lands came through the slit that stayed open, so nothing it touched was altered, and yet where it lands has changed. The tempting conclusion is that the experiment can only be explained if the particle went down every route available to it. But the information those routes carry is already held in two other places: in the object's *disposition*, what it does under a given circumstance, and in the arrangement it is placed in. Loosely, and only as a way of fixing the shape of the claim:

$$\text{disposition} + \text{arrangement} = \text{path}.$$

The robot had a disposition in that sense: its code said what to do at a junction, and the maze supplied the junctions. Let it decide at random instead and no single path follows, but the equality survives in a modified form,

$$\text{propensity} + \text{arrangement} = \text{possible paths},$$

where a *propensity* is a disposition with probabilities in it, the likelihood of doing one thing rather than another under a given circumstance. That is the case a particle is in, and the possible paths are what a distribution is a distribution over. Unistochastic dynamics is that equality made into a formalism. Multiply a circuit's unitary operators together into a single U and take $|U|^2$, as Section 6.2's multi-step shortcut does, and the statistics follow with no quantum state appearing anywhere. What those operators encode is the arrangement itself. In the three-bit circuit of Section 8.7 the Givens operator G is a beam splitter and the SWAP is the barrier, moving the photon onto the barrier's own bit when it comes that way. Putting the barrier in means putting that SWAP in, and the marginal statistics of where the photon ends up change accordingly. The whole layout is thereby settled before any statistics are computed at all, which is what it means for a law to consult the arrangement.

The magician of the Introduction shows it again. Had he been told which program you ran, he would have known at once whether the bit came out uniform or degenerate, because the program fixes that. The distribution is already contained in the program and the machine, and nothing further has to be carried along by the bit. One look at the code that produced a single bit tells him which distribution that bit belongs to, and looking back at what produced it is exactly the non-Markovian dependence Part II reconstructed.

So any account that explains interference by having the particle take every route can be replaced by one that reads the same information off the arrangement, and that is what this chapter has been doing. Because the laws are non-Markovian, the arrangement that counts is not only the present one: what the experiment did to the particle earlier belongs to it as well, which is why the history and the arrangement have to be given together. The double slit and the Mach-Zehnder were offered for a century as proof that the particle goes both ways and that a wave is needed to say how. The structure of the whole experiment does that work instead, and in the interferometer that structure includes a barrier the photon never touched, sitting on the path it did not take. QMI is what puts a number on that structural fact, which is why this section opened by calling it a de-localized property: it belongs to the arrangement and to no part of the arrangement on its own.

The stakes of that retirement were named a long time ago. In 1948, in a short essay for the journal *Dialectica*, Einstein set out what he regarded as the working condition of physics itself: things lying in different parts of space must claim an existence independent of one another, such that an external influence on A has no immediate effect on B . This is his “principle of local action,” and his defense of it was methodological rather than metaphysical. Without such independence of spatially distant things, he wrote, “physical thinking in the sense familiar to us would not be possible,” since one does not see “how physical laws could be formulated and tested” without a clean separation; and complete abandonment of the principle “would make the idea of the existence of (quasi-)closed systems, and with it the establishment of empirically testable laws in the sense familiar to us, impossible.”⁹ Notice the shape of the worry. Physics studies things by cutting them out of the world, and a world that cannot be cut leaves nothing to study. Non-locality repelled Einstein because it dissolves the very objects on which lawful, testable, deterministic science was to be practiced.

The most pointed reply came from the Soviet physicist Dmitry Blokhintsev, whom Chapter 9 will meet again. In his 1951 critique of the idealist readings of quantum theory, he quotes Einstein’s sentence back at him with the qualifier underlined: *in the sense familiar to us*. That non-closed systems would wreck the familiar formulation of natural law, he answers, does not show that nature refrains from non-closedness; “it follows only that these laws would have to be sought in a form unfamiliar to us.”¹⁰ Quantum mechanics, on his reading, is that unfamiliar form, found. “In nature there is no absolute division into microworld and macroworld”; because of the atomism of action (the finiteness of h), “closed, isolated microsystems do not exist, and every quantum ensemble includes within itself the connection of microsystems with macrosystems.” The inversion is complete. Einstein rejected non-separability because it would end the deterministic prediction of isolated pieces; Blokhintsev

⁹A. Einstein, “Quanten-Mechanik und Wirklichkeit,” *Dialectica* 2, no. 3–4 (1948), 320–324, at 324. Translations ours.

¹⁰D. I. Blokhintsev, “Kritika idealisticheskogo ponimaniya kvantovoy teorii” [Critique of the idealist understanding of quantum theory], *Uspekhi Fizicheskikh Nauk* 45, no. 2 (1951), 195–228, at 207–210. Translations ours.

accepts non-separability and reads quantum statistics as exactly that ending, made lawful. A system that cannot be fully cut out of the world is a system whose future cannot be computed from its own state alone, and probability is what lawfulness looks like under that condition. (Chapter 2 taught the small version of this lesson: the COIN operator, which looks like pure chance, is the deterministic XC gate on two bits, seen by someone who can track only one of them.)

Blokhintsev held, moreover, that even classical determinism had always been quietly conditional on the same idealization:

a deterministic prediction of the future is itself conditional: the future of a mechanical system may be predicted only if we can be sure that the system is isolated. The guarantee required here is not implied by the equations of motion but is an additional condition, which produces a great reduction in the reliance on determinism. A vast and depressing “if” arises in the way of the prophet who sets out to predict the future of a real mechanical system.¹¹

Classical physics could treat that “if” as a practical nuisance: a system like the solar system is isolated well enough that the condition holds for any working purpose. The finiteness of h makes it a matter of principle instead, because it means no system is ever isolated exactly. We do not adopt Blokhintsev’s ensemble interpretation wholesale (Chapter 9 relies on him for different purposes), but on the present point his position and our mathematics say the same thing: the statistics of a bit never belong to that bit alone, and the demand that they should is not a finding of physics but an assumption physics has had to give up.

Einstein’s condition was that the world come already divided into things that hold their properties on their own. Bohm’s answer is that the division is ours. Fragmentation, he argued, comes of “the almost universal habit of taking the content of our thought for ‘a description of the world as it is’”; since thought is pervaded with distinctions, the habit leads us “to look on these as real divisions,” and the world is then “seen and experienced as actually broken up into fragments.” We then act on the picture and arrange things to match it, which supplies “an apparent proof of the correctness of his fragmentary self-world view,” while overlooking that it is the thinker “who has brought about the fragmentation that now seems to have an autonomous existence.”¹² A theory, on his reading, is “primarily a form of insight, i.e. a way of looking at the world, and not a form of knowledge of how the world is.”

That is the distinction Chapter 9 will draw with the philosopher Jocelyn Benoist, between a concept, which is ideal, and what the concept is applied to, which is real; and the philosopher François-Igor Pris makes the corresponding point about the applications themselves, that they are “not autonomous” but presuppose a context.¹³

The word *context* clearly has a holistic character. Benoist denies that a thing has a character settled independently of the circumstances it is met in: if the red of a car is real, “its bronze colour in the tunnel is no less ‘real,’” even though by one standard that is not the car’s real color.¹⁴ There is no lighting-free version of the car waiting behind the two to be the true one. And the work he gives the context is to settle what is there rather than to loosen

¹¹D. I. Blokhintsev, *The Philosophy of Quantum Mechanics* (Dordrecht: D. Reidel, 1968), pp. 10–11.

¹²D. Bohm, *Wholeness and the Implicate Order* (London: Routledge, 1980), ch. 1.

¹³I. Pris, *Contextual Realism and Quantum Mechanics* (Chisinau: Scholars’ Press, 2020), §1.

¹⁴Benoist, *Toward a Contextual Realism*, “(Perceptual) Things Being What They Are.”

it: “the role of the context is to fix the content, not to make it unstable.”¹⁵ Context is what fixes what gets realized, and it takes in the whole arrangement the thing is met in, which is holism under another name. The photon’s distribution is fixed by the interferometer it is placed in, and there is no interferometer-free version of it waiting behind to be the real one. A particle with a distribution of its own, carried about independently of any arrangement, is an autonomous object in exactly that sense: a piece of the accounting mistaken for a piece of the world. Stop making that mistake and the holism of this section is the consequence.

The same pattern appears outside physics. Blokhintsev’s account of chance is quoted at length from the philosopher Friedrich Engels’s *Dialectics of Nature*, whose fragment on chance and necessity attacks the determinism that “passed from French materialism into natural science” on exactly the ground at issue here, the interconnection of things. Trace any individual fact backward and its causal web fans out without limit; a single pea-pod, Engels observes, “would already provide more causal connections for following up than all the botanists in the world could solve.” A necessity asserted across such a web explains nothing (call it “the eternal decree of God, or Kismet,” call it necessity: for science it is all one), and if the count of peas in a pod is placed on the same footing as the law of motion of the solar system, then “chance is not elevated into necessity, but rather necessity degraded into chance.”¹⁶ Absolute determinism, on this argument, was already broken on the interconnectedness of things before quantum mechanics arrived to make the breakage exact.

The philosopher Louis Althusser, writing in the same tradition a century later, gave the thesis its name: *overdetermination*. He worked it out on social formations, where an event is never the product of one isolable contradiction, since the contradiction is “inseparable from the total structure of the social body in which it is found,” is “determining, but also determined in one and the same movement,” and is therefore “overdetermined in its principle.”¹⁷ Where every cause is conditioned by the whole in which it acts, no single deterministic cause can be extracted, and prediction acquires an ineliminably statistical character.

The thesis reaches further than that. Engels was writing about nature, which is why the pea-pod and the law of motion of the solar system sit in the same fragment. Resnick and Wolff, restating overdetermination decades after Althusser, state the premise with no domain attached to it: everything is related to everything else, “the conditions of the existence of any event are infinite in number and variety,” and each of those conditions has conditions of its own, which puts complete explanation of anything out of reach.¹⁸ What is left is a choice about how to proceed. One may take some conditions to be the essential ones and build on those, which is how a system gets treated as isolated, or one may hold that there are no essential causes and work with explanations that stay partial by construction. The finiteness of h makes the first option unavailable here, and the mathematics of Chapter 4 says why: the correlations that would have to be dropped are the ones carrying the phase.

We flag the lineage and leave it. What our own argument requires, we have supplied: a mathematical sense in which the law of a particle’s statistics refuses to factor into private,

¹⁵Benoist, *Toward a Contextual Realism*, “Contextualism or Relativism?”

¹⁶F. Engels, *Dialectics of Nature*, trans. C. Dutt (Moscow: Progress Publishers, 1954), the notes fragment “Chance and Necessity.” Blokhintsev quotes the passage at length in the 1951 critique.

¹⁷L. Althusser, “Contradiction and Overdetermination,” in *For Marx*, trans. B. Brewster (London: Verso, 2010), pp. 100–101.

¹⁸S. Resnick and R. Wolff, “Everythingism, or Better Still, Overdetermination,” *New Left Review* 1/195 (September–October 1992), 124–126, at 125.

local pieces, owing to the histories those particles share, which run back to the Big Bang.

That last clause can be put on a firmer footing. Buniy and Hsu argue that Big Bang cosmology leaves a typical particle entangled with an enormous number of others, including particles far outside our horizon, since regions now out of causal contact were once in equilibrium with each other.¹⁹ Their second result reads at first like a retraction of the first: two randomly chosen particles are unlikely to be entangled with each other at all. The reduced density matrix of the pair is *separable*, a mixture of products with no entanglement in it, and even its classical correlations are slight, since that reduced state sits close to maximally mixed.

The two results are one result, and decoherence is what joins them. Obtaining the pair's description means tracing out everything else, which is the partial trace of Chapter 4, and entanglement divided among an enormous number of partners leaves none for any particular pair. Separability here is emergent, produced by the universality of the entanglement rather than by any lack of it. This is also why the world admits local description at all. Every system physics has ever treated as isolated was decohered by a history it shares with everything else, and Section 6.3's apparent base cases are that fact stated for a circuit. Blokhintsev's "if" holds in practice for the same reason it fails in principle.

8.9 When the statistics fix the law

Section 6.5 said that two constraints pin the single-step matrices down almost completely, and Section 6.6 supplied a rule for what they leave open. Run the interferometer without the barrier and they leave nothing open at all.

Start with what has to be explained. The beam splitter's unistochastic matrix is the coin, and two splitters composed give something else:

$$|B|^2 = \begin{pmatrix} 0.5 & 0.5 \\ 0.5 & 0.5 \end{pmatrix}, \quad |B B|^2 = \begin{pmatrix} 0 & 1 \\ 1 & 0 \end{pmatrix} \neq |B|^2 |B|^2.$$

Section 3.6 gave the Hadamard pair as the minimal witness of that inequality, which is the same witness in another operator, and the correction term of Section 8.8 is the difference in general. Whatever matrix the second splitter carries, it is not the one the first carries.

Now take the two splitters in turn. Send the photon into the first one on the top port and it leaves on either path with equal probability; send it in on the bottom port and the same holds; so both columns of the first splitter's matrix are the uniform distribution.

Measure between the splitters and the photon is on one path or the other, half the time each. A measurement does disturb the experiment, and with a detector in place the photon no longer leaves the second splitter on one port with certainty. The disturbance falls on what the second splitter does with the photon, though. The statistics the detector records are the ones a single splitter gives on its own, which Section 8.2 ran separately: a photon in, either path out, half the time each. So the first splitter has done what a first splitter does, and short of retrocausality a detector placed after it has no way to reach back and alter that. (Assuming the photon had a real position independent of measurement at all.)

So both of the second splitter's columns are exercised, and what comes out is one port with certainty. A stochastic matrix has no negative entries, so an outcome of probability zero

¹⁹R. V. Buniy and S. D. H. Hsu, "Everything is Entangled," *Physics Letters B* 718 (2012), 233–236.

can receive nothing from either of two columns that both carry probability, which puts a zero at both entries of the empty port's row and leaves the certainty in each column. Had the incoming distribution been degenerate, only one column would have been exercised and the other would have stayed open.

Both matrices are settled, then, and no transition law was consulted to settle them:

$$\Gamma^{[1 \leftarrow 0]} = \begin{pmatrix} 0.5 & 0.5 \\ 0.5 & 0.5 \end{pmatrix}, \quad \Gamma^{[2 \leftarrow 1]} = \begin{cases} \begin{pmatrix} 0 & 0 \\ 1 & 1 \end{pmatrix} & \text{if the photon entered on the top port} \\ \begin{pmatrix} 1 & 1 \\ 0 & 0 \end{pmatrix} & \text{if it entered on the bottom port.} \end{cases}$$

These are Section 5.1's Γ_{coin} , Γ_{one} and Γ_{zero} , which is no surprise, Section 8.2 having identified the interferometer with the Hadamard pair. The second matrix is not the first, and that is the inequality above read off the circuit: a beam splitter with another in front of it does not do what a beam splitter with nothing in front of it does. Section 8.8's holism, stated for gates, is that fact.

The zero is what does the work. Where the outgoing distribution has a zero in it, non-negativity carries that zero into the matrix along a whole row, and enough such zeros leave one matrix standing. Where the outgoing distribution has no zeros, nothing of the sort happens.

Put the barrier back and the same circuit leaves both matrices open. Before the barrier the photon is on the top path or the bottom with equal probability; after it, the photon is on the top path or absorbed, again with equal probability; and after the second splitter the three outcomes that occur carry a half, a quarter and a quarter. No outcome that happens has probability zero at either step, so no row is forced to zero and nothing is settled. The constraints leave a one-parameter family of matrices at the barrier and a two-parameter family at the splitter after it.

At the barrier the incoming distribution is uniform over the two paths and the outgoing one is uniform over absorbed and top, so the statistics cannot tell a barrier that absorbs from the bottom path and leaves the top one alone from a barrier that absorbs from the top path and moves the other photon across. Section 6.6 picks $|U|^2$ out of that family, and something has to.

Section 7.10 arrived at a law-free answer by a different route. The branches of the teleportation circuit never meet, so the interference correction of Section 5.6 has nothing to act on and the polar decomposition hands over $|U|^2$ at every step with no rule invoked. The statistics can therefore settle the matter in two separate ways: once because the circuit never interferes, and once because interference and non-negativity between them leave a single candidate standing. Neither way holds generally, which is what the rule of Section 6.6 is for, and any rule worth adopting has to return what quantum mechanics has already fixed in cases like these two. Ours does.

8.10 Simulating the double slit

Everything in this chapter has been argued on the interferometer, chosen because two paths are easier to compute with than a continuum of screen positions. The argument transfers to the double slit, and we close by watching it do so.

The pictures below come from a simulation built out of the machinery presented so far and nothing else. The particle's position on a grid of 256 by 512 cells is held in two coordinate registers, one of eight bits and one of nine, seventeen bits in all, and one unitary is applied per time interval. The bits are given values at every interval, carried from one interval to the next by the iterative proportional fitting of Chapter 6 and the stochastic hops of Section 2.7, exactly as in the H, T, H circuit computed there, so at every moment the particle is somewhere. That is what makes the pictures possible at all. Each picture is many runs laid over one another, one faint line drawn per run, and each line is that run's trajectory: a definite path through a definite sequence of cells. The dark regions are where the most trajectories overlap.

The unitary is a hopping step, and like everything else here it is written as a circuit. Across all seventeen bits at once, the step is

$$U_{\text{step}} = X_e X_o Y_e Y_o^{(2\alpha)} Y_e X_o X_e,$$

where, for an n_x -bit x -register and an n_y -bit y -register,

$$\begin{aligned} X_e &= I^{\otimes(n_x-1)} \otimes R \otimes I^{\otimes n_y}, & Y_e &= I^{\otimes n_x} \otimes I^{\otimes(n_y-1)} \otimes R, \\ X_o &= (D_x \otimes I) X_e (D_x^\dagger \otimes I), & Y_o &= (I \otimes D_y) Y_e (I \otimes D_y^\dagger), \end{aligned}$$


with D the increment circuit that sends $\mathbf{e}_j$ to $\mathbf{e}_{j+1}$, and

$$R = R_X(-2\alpha) = \begin{pmatrix} \cos \alpha & i \sin \alpha \\ i \sin \alpha & \cos \alpha \end{pmatrix}.$$

The angle α is the only number in the construction that can be chosen; it sets how far the particle moves per interval.

The rotation sits on a register's low-order bit because the cells $2k$ and $2k+1$ differ in that bit alone, which covers half the neighboring pairs; the other half straddle a carry, and are reached by the shift, rotate and shift back that X_o and Y_o perform. The middle factor carries twice the angle because the palindrome places two identical factors together there and they merge, and running the four as a palindrome rather than in sequence is what leaves an error second order in α instead of first.


The operator above is exact on a lattice that wraps around, so a hard wall has to be imposed separately, by removing the single neighboring pair $\mathbf{e}_{N-1}$ with $\mathbf{e}_0$. It contains no barrier at all: to block a cell, each R becomes a controlled rotation, conditioned on the rest of the register not naming a blocked cell. That is how the slits themselves are made, and for an arbitrarily drawn barrier there is no compact closed form for it.



The left run has one slit and the right has two. The particles are the same, the unitary is the same, the hopping angle is the same; only the barrier differs. Both patterns spread out, because a particle that passes through a narrow opening does not carry on in a straight line, and that spreading is diffraction. The thing to look at is the dense core of each picture, the region beyond the barrier where the trajectories crowd together most thickly. In the left picture that core is solid: a particle is about as willing to take any path through it as any other. In the right picture the same core is divided, with two thin regions running through it that hardly any trajectory enters. Those are the minima produced by interference.

A blocked cell and an open cell make different arrangements; the rotation is omitted on pairs that touch a block and applied on pairs that do not; so the law the particles hop under is not the same law in the two runs. A particle going through the upper slit of the right picture is hopping under a law fixed by a barrier that has a second opening in it, and that is why there are places it does not go. The minima are where the law sends fewer particles, and the law is a function of the arrangement.

There is no wave anywhere in this, and the pictures show where the appearance of one comes from. Below, we run the simulation once and plot that single run:



That is the complete record of one execution: one particle, one definite path, through one slit. There are no fringes in it, and there could not be. A minimum is a statement about where many particles land, and one particle lands in one place. The interference pattern is a property of the ensemble, and it exists only once the ensemble has been assembled.

Einstein read the quantum state exactly this way:²⁰

The attempt to conceive the quantum-theoretical description as the complete description of the individual systems leads to unnatural theoretical interpretations, which become immediately unnecessary if one accepts the interpretation that the description refers to ensembles of systems and not to individual systems. In that case the whole “egg-walking” performed in order to avoid the “physically real” becomes superfluous.

²⁰A. Einstein, “Remarks Concerning the Essays Brought together in this Co-operative Volume,” in P. A. Schilpp (ed.), *Albert Einstein: Philosopher-Scientist* (New York: MJF Books, 2001), p. 672.

The strangeness that has to be walked around, the century of disputes about what is physically real, the standing advertisement that nobody understands quantum mechanics: none of it follows from the statistics. It follows from insisting that ψ is something the individual system carries. Grant that ψ describes an ensemble and the “egg-walking” stops, because there is nothing left to walk around. A distribution is not the sort of thing an individual system has.

Blokhintsev made the same point empirically, in terms of what a single run of an experiment on a single particle actually shows, which is no wave at all:²¹

[I]t is stated in many textbooks on quantum mechanics that the wave function is a characteristic of the state of a single particle. If this were so, it would be of interest to perform such a measurement on a single particle (say an electron) which would allow us to determine its own individual wave function. No such measurement is possible.

None of which is strange once a wave is asked what it is a wave *of*. A wave on water is a pattern in the motion of water molecules; a sound wave is a pattern in the motion of air. A light wave is a pattern in the arrival of photons, and the interference pattern of an electron beam is a pattern in the arrival of electrons. In every case the wave is something a collection *does*. One water molecule does not wave, and, in the picture above, one particle is not sufficient to produce the interference pattern. There is no wave composed of nothing, as nothing cannot *do* anything.

The analogy breaks down in one respect. A wave on water is a pattern in a motion that is itself deterministic: the molecules obey ordinary mechanics, and the statistics are a summary kept because tracking every molecule is impractical. Statistical theories are usually like that, emergent descriptions of some underlying deterministic process. What is unusual about the laws reconstructed here is that nothing has been supplied beneath them. The transitions are stochastic at the level of the law itself, and no deeper machinery is posited to produce them.

That the account supplies nothing beneath the laws does not mean there is nothing there. Ballentine, setting out the statistical interpretation in 1970, was careful on the same point: the interpretation “does not prejudice the possibility of introducing hidden variables which would determine (in principle) the outcome of each individual measurement.”²² The position here is the same. Nothing in the construction requires a deterministic theory underneath it, and nothing in it forbids one.

The restraint here is a matter of parsimony. To hold that quantum theory needs an underlying deterministic process before its statistics can be accounted for would leave realism, which Part III takes up, waiting on the completion of a theory that has yet to reproduce the whole of modern physics. It waits on no such thing: the bits hold definite values already, under the laws reconstructed here.


Section 8.6 wrote the additivity assumption as a sum over distributions measured under two different arrangements. Write *BB* for a barrier blocking the bottom slit and *BT* for a barrier blocking the top. The single-slit experiments of Section 8.1 block one slit each. With the bottom slit blocked, every transmitted particle passed through the top slit, and

²¹D. I. Blokhintsev, *The Philosophy of Quantum Mechanics* (Dordrecht: D. Reidel, 1968), p. 50.

²²L. E. Ballentine, “The Statistical Interpretation of Quantum Mechanics,” *Reviews of Modern Physics* 42 (1970), 358–381, at 378.


the histogram of final positions over those runs defines $P(x \mid \text{top}, BB)$: the distribution of screen positions given passage through the top slit with the bottom slit blocked. With the top slit blocked, every transmitted particle passed through the bottom slit, and the histogram over those runs defines $P(x \mid \text{bottom}, BT)$: the distribution of screen positions given passage through the bottom slit with the top slit blocked. The sum that predicts the pattern with both slits open needs the same two distributions with the barriers removed, $P(x \mid \text{top}, \neg BB)$ and $P(x \mid \text{bottom}, \neg BT)$. In the interferometer these were computed from the circuit.

The first chart shows the two blocked-slit distributions and their average. Each blocked-slit distribution has a single broad local maximum at the position of the open slit. The average is smooth, and the deep minima of the double-slit pattern (the dark bands where the particles are unlikely to be found at all) are absent from it. This is the failed sum of Section 8.1, produced by the simulation, and it fails for the reason Section 8.6 identified: the two distributions in the sum were measured with a barrier present.



As this is a simulation, we can play the demon of Section 6.5 and measure $P(x \mid \text{top}, \neg BB)$ and $P(x \mid \text{bottom}, \neg BT)$ directly, even with both barriers removed. Every run passes through one slit, so the runs can be sorted into separate buckets according to which slit the particle passed through, and the final positions within each bucket give the corresponding conditional distribution. A laboratory could not sort this way. The sort requires observing which slit the particle takes. The observation is a recording, and a recording creates a base case that removes the interference, as Section 8.7's detector does: the laboratory's buckets would hold the smooth single-slit distributions of the first chart, with the fringes gone. The simulation includes none of the physical disturbance that an observation carries, so the interference survives: the complete trajectories can be simulated with both slits open and sorted afterwards. With both slits open, the distribution of final positions in each bucket is roughly the double-slit distribution cut in half at the midline, the line halfway between the two slits: the particle almost never crosses it.

The second chart shows the two bucket distributions and their average. Each has the same shape as the double-slit pattern and comes close to zero at its minima, including minima where the corresponding blocked-slit distribution of the first chart is large.



Expressed in the conditional notation, each conditional is, to a good approximation, the double-slit pattern restricted to one side of the midline,

$$P(x \mid \text{top}, \neg BB) \approx \begin{cases} 2 P(x \mid \neg BB, \neg BT) & x \text{ on the top side of the midline} \\ 0 & x \text{ on the bottom side} \end{cases}$$

$$P(x \mid \text{bottom}, \neg BT) \approx \begin{cases} 0 & x \text{ on the top side of the midline} \\ 2 P(x \mid \neg BB, \neg BT) & x \text{ on the bottom side.} \end{cases}$$

The factor of 2 restores normalization: one side of the midline carries half of the double-slit pattern's probability, and doubling the restricted pattern makes it sum to one again. The average of the two bucket distributions reproduces the double-slit pattern to sampling error,

$$P(x \mid \neg BB, \neg BT) = \frac{1}{2} P(x \mid \text{top}, \neg BB) + \frac{1}{2} P(x \mid \text{bottom}, \neg BT).$$

This is the corrected sum of Section 8.6, and both of its terms are now measured distributions.

The approximation signs come from the lattice. The midline falls between two cells. On the bond between those cells the two Born probabilities are equal, so $|U|^2$ already transports them correctly and the fitting of Section 6.6 returns it unchanged. The step across the midline is a symmetric exchange between the two cells, so a walker can cross the midline by a few cells. A small part of each population ends up on the other side, almost all of it within a few cells of the midline, and that part shrinks when the time intervals are shortened. In the continuum limit no trajectory crosses the midline.

The two charts show the dependence on the arrangement directly. A trajectory through the top slit ends in the fringed distribution of the second chart when the bottom slit is open and in the smooth blocked-slit distribution of the first chart when the bottom slit is closed. The particle passed through the same slit in both cases. The transition law depends on the whole arrangement, as Section 8.8 stated, so closing the bottom slit changes the distribution of every particle that passes through the top slit.

The deepest minimum, near the midline, repeats the interaction-free measurement of Section 8.5. A trajectory through the top slit that ends there never approached the bottom slit. With both slits open, arrivals in that minimum are rare. With the bottom slit closed, arrivals there are several times more frequent. A particle detected in the minimum is evidence that the bottom slit is closed, on a path the particle never took. The evidence is statistical, since a slit of finite width does not produce zero minima; the interferometer of Section 8.5 gives the same evidence with certainty.

8.11 Wheeler’s delayed-choice experiment

In Wheeler’s *delayed-choice* arrangement, the second beam splitter of a Mach–Zehnder interferometer is inserted or withdrawn at the last picosecond, after the photon is already in flight.²³ With it in place the output port is certain; with it removed the two ports come up uniformly. The decision is made too late to be communicated to a photon choosing, at the first splitter, whether to behave as a wave or as a particle. Wheeler put the provocation as deciding “whether the photon ‘shall have come by one route, or by both routes’ after it has ‘already done its travel.’”

That phrasing is where the popular reading of the experiment comes from. The photon makes its trip either as a wave, taking both routes, or as a particle, taking one, and which it did is settled by whether interference appears at the end. In the Mach–Zehnder interferometer, interference appears only when the second splitter is in place. So the photon’s choice is made at the *first* splitter while depending on what is done at the *second*, and if that second decision is delayed until the photon is already between them, the choice appears to reach back and rewrite a trip already under way.

On the reading built here the puzzle does not arise, because no such decision is ever taken. The photon is on one path, as it always is. Its distribution over the two paths after the first splitter is uniform, and it remains uniform whether or not a second splitter is later installed, since the law reaches backward and not forward. What the late choice settles is which gate acts next. Inserting the second splitter changes the statistics, but only from that point on. It settles nothing about the trip already taken.

Wheeler drew the same line, though from the opposite side. His claim was never that the late choice reaches back and alters the photon’s past, but that it has “an unavoidable effect on what we have a *right to say* about the already past history of that photon,” since “we have no right to say what the photon is doing in all its long course from point of entry to point of detection.” On his account there is no history to be altered; on ours there is one, and the late choice does not touch it. Neither account needs anything to travel backward in time, and the retrocausal reading belongs to neither.

The same reading covers the *delayed-choice quantum eraser*, where which-path information is recorded and then discarded after the photon has already been detected.²⁴ There too nothing is retrocausal: the full set of detections shows no interference at all, and the fringes appear only in subsets sorted afterwards by a correlation that was established when the pair was made.

²³J. A. Wheeler, “The ‘Past’ and the ‘Delayed-Choice’ Double-Slit Experiment,” in A. R. Marlow (ed.), *Mathematical Foundations of Quantum Theory* (New York: Academic Press, 1978), 9–48; the passages quoted here are from Wheeler’s later “Law Without Law,” in J. A. Wheeler and W. H. Zurek (eds.), *Quantum Theory and Measurement* (Princeton: Princeton University Press, 1983).

²⁴M. O. Scully and K. Drühl, “Quantum eraser: A proposed photon correlation experiment concerning observation and ‘delayed choice’ in quantum mechanics,” *Physical Review A* 25 (1982), 2208–2213; Y.-H. Kim, R. Yu, S. P. Kulik, Y. Shih and M. O. Scully, “Delayed ‘Choice’ Quantum Eraser,” *Physical Review Letters* 84 (2000), 1–5. The reading given here is argued at length in R. E. Kastner, “The ‘Delayed Choice Quantum Eraser’ Neither Erases Nor Delays,” *Foundations of Physics* 49 (2019), 717–727, and in T. Qureshi, “The Enigma of Delayed Choice Quantum Eraser,” *Quanta* 14 (2025), 66–74.


8.12 Albert's moving pointer

Section 6.7 adopted a law for the interval between base cases, and left one objection standing until the double slit had been introduced. David Albert's charge is that a theory which fixes the probability of each configuration, and says nothing about the path between them, has not yet said enough to contain observers.²⁵ Fill the gap the simplest way, by drawing each moment afresh from the distribution belonging to it. He calls that the *minimal completion*, and it is the resampling of Section 6.7 under another name: a moment drawn independently of the one before it is a moment whose transition matrix carries the target distribution in every column.

The records such a world contains agree with one another at every instant while indicating nothing whatever about the past. Albert's example is a book about Napoleon, found on a shelf. The distribution that put it there makes other Napoleon books likely, and paintings, and memories, all of them elsewhere in the world *at that same moment*. Ask the same distribution what the world held *earlier* and the book on the shelf counts for nothing, since each moment is drawn on its own. What governs that earlier moment is the unconditioned distribution, spread across the whole space of configurations, and the configurations holding an earth are a vanishing fraction of it. Albert's verdict goes further than saying we cannot tell: the probability of there having been an earth in 1804, or ten minutes ago, or a nanosecond ago, comes out astronomically low. The records tally with one another because one moment's configuration is drawn all at once.

Nothing in such a world has a reliable disposition, and a measuring device is nothing if not a reliable disposition, so it is fair to doubt that any measurement happens there at all.

He argues it through an experiment, and the choice is a good one: a particle passing two slits, with a pointer that records which slit it took. Everything at issue sits in one arrangement. The record is made there, the interference the record destroys can be seen there, and the stretch between the recording and the particle's arrival at the screen is exactly the interval he says goes unmodeled.



The pointer rests at C and moves to the box marked A or the box marked B according to which slit the particle passed, and once it has moved it is left alone: every later gate acts on the particle only. The particle runs on a ring of sixty-four sites, far smaller than Section 8.10's grid and for a reason: the whole arrangement, particle and pointer together, is then small enough that the transition matrix can be written down and fitted at every step rather than

²⁵D. Z. Albert, "Quantum mechanics without a wave-function," *Synthese* 208 (2026), article 141. The charge is his §6; the arrangement with the recording pointer is his §4.

sampled one run at a time, which is what lets the record be watched directly. The code for this chapter runs it under both laws.

Take the arrangement first, before any question of a transition law. The position distribution at the screen can be compared with the classical sum of the two single-slit distributions, and the gap between them is *interference*:

	distance from the classical sum
no pointer	0.33
pointer recording	0

So the arrangement behaves: the recording removes the interference, and removes it completely. That much is the quantum state's doing and holds whatever law fills the interval.

Now the interval. After the pointer has moved, the two laws are asked how the configuration gets from one moment to the next:

per step	KL-minimal	resampling
record changes	0	0.50
particle moves	0.8 sites	10.6 sites

The first row is the answer to Albert. Under the law of Section 6.6 the probability that the record changes is zero, and zero exactly rather than merely small: a gate that leaves the pointer alone has $|U|^2$ equal to zero on every entry that would move it, and Section 6.6's fitting cannot put probability where the unitary put none. Under resampling the record is a coin, flipped afresh every step.

The second row answers a question he does not ask but would. A law could hold the record steady and still have the particle appear at the far side of the ring from one instant to the next. It does not. The unitary permits single hops of up to seven sites, and the fitted law uses less than one, because the same property that pins the record pins the positions: the gate's own $|U|^2$ is massed near the diagonal and the fitting keeps it there. The particle takes a path. It is not the Bohmian path, since the steps are stochastic rather than guided, and it is not far off one either.

What remains is the thing Albert wanted and doubted could be had. The record holds, so a system with a reliable disposition can exist. The particle keeps to a path, so there is something for the record to be a record of. And the interference that ought to go, goes.

To reiterate what Section 6.7 said: quantum theory does not fix the law from its own structure, and Barandes is right that what it leaves is an equivalence class of processes agreeing on the first-order data. We cannot be certain that a law of this kind is the correct one, and it is put forward to illustrate. Albert's criticisms do not disprove Barandes' position either. Schrödinger thought he had settled such a picture by ridicule, retorting that he could not imagine an electron hopping about like a flea,²⁶ and for all anyone knows the flea is what an electron is. What criticisms of this kind do is place reasonable further constraints, narrowing the class, and it was concerns of just that sort that brought us to this law in the first place, as Section 6.7 set out.

²⁶E. Schrödinger, quoted in C. Rovelli, *Helgoland: Making Sense of the Quantum Revolution*, trans. E. Segre and S. Carnell (New York: Riverhead Books, 2021), ch. 1, who takes it from M. Kumar, *Quantum: Einstein, Bohr and the Great Debate about the Nature of Reality* (London: Icon Books, 2008), p. 220.

This closes Part II. The double slit, the delayed-choice variant of Section 8.11, and the recording pointer just considered were the last exhibits customarily offered as *proof* that no statistical, definite-values account of quantum mechanics can exist. It has now been given exactly such an account, one mathematically equivalent to the quantum-mechanical description. What remains are the arguments that such an account, though possible, is somehow philosophically inadmissible: collapse, Wigner's friend, no-go theorems, and the rest of the interpretive objections. Part III takes them in order.

Code: this chapter in Octave

```

% --- Section 8.2: the bare interferometer ---
B = (1/sqrt(2)) * [1 1i; 1i 1];
abs(B*B*[1;0]).^2          % -> [0; 1]: always the vertical port
abs(B*[1;0]).^2            % one splitter alone: [0.5; 0.5]

% --- Section 8.7: three-bit circuit (top path, bottom path, barrier) ---
I2 = eye(2);  c = cos(pi/4);  s = sin(pi/4);
G = [1 0 0 0; 0 c -s 0; 0 s c 0; 0 0 0 1];  % beam splitter (Givens)

% Givens decomposition into SWAP, CX, RY (circuit order left to right:
% SW, CX(1->2), RY(-t) on B1, CX(2->1), RY(t) on B1, CX(1->2)):
SWg = [1 0 0 0; 0 0 1 0; 0 1 0 0; 0 0 0 1];
CX12 = [1 0 0 0; 0 1 0 0; 0 0 0 1; 0 0 1 0];
CX21 = [1 0 0 0; 0 0 0 1; 0 0 1 0; 0 1 0 0];
RY = @(t) [cos(t/2) -sin(t/2); sin(t/2) cos(t/2)];
err = 0;                                     % check it at several angles
for t = linspace(-pi, pi, 9)
    Gt = [1 0 0 0; 0 cos(t) -sin(t) 0; 0 sin(t) cos(t) 0; 0 0 0 1];
    Gdec = CX12 * kron(RY(t),I2) * CX21 * kron(RY(-t),I2) * CX12 * SWg;
    err = max(err, norm(Gdec - Gt));
end
err                                     % -> 0: exact at every angle
SW = [1 0 0 0; 0 0 1 0; 0 1 0 0; 0 0 0 1];  % barrier: path <-> barrier
e100 = zeros(8,1);  e100(5) = 1;           % photon on the top path
paths = kron(eye(4), [1 1]);               % marginalize the barrier

mid = kron(I2,SW) * kron(G,I2) * e100;      % after the barrier column
fin = kron(G,I2) * mid;                    % after the 2nd splitter
abs(fin).^2                                % [0 .5 .25 0 .25 0 0 0]
pmid = paths * abs(mid).^2                 % [0.5; 0; 0.5; 0]
pfin = paths * abs(fin).^2                 % [0.5; 0.25; 0.25; 0]
abs(G).^2 * pmid                           % = pfin: base case -> Markov works

% --- QMI(top : barrier) at the intermediate time ---
rho = mid * mid';
trB2 = @(r,keep) keep * r * keep';        % generic sandwich partial trace
e0=[1;0]; e1=[0;1];
% trace out the bottom-path bit (bit 2 of 3):
K0 = kron(kron(I2,e0'),I2);  K1 = kron(kron(I2,e1'),I2);
rTB = K0*rho*K0' + K1*rho*K1';           % rho^{top,barrier}
S = @(r) (@(v) -sum(v(v>1e-12).*log2(v(v>1e-12))))(real(eig(r)));
trR = @(r) [r(1,1)+r(2,2), r(1,3)+r(2,4); r(3,1)+r(4,2), r(3,3)+r(4,4)];
trL = @(r) [r(1,1)+r(3,3), r(1,2)+r(3,4); r(2,1)+r(4,3), r(2,2)+r(4,4)];
QMI = S(trR(rTB)) + S(trL(rTB)) - S(rTB)  % -> 2: maximal, base case

```

```

% --- no barrier: no base case, Markov guess fails = interference ---
nb = kron(G,I2) * kron(G,I2) * e100;
paths * abs(nb).^2 % [0; 1; 0; 0]
abs(G).^2 * paths * abs(kron(G,I2)*e100).^2 % [0; .5; .5; 0]: wrong

% --- Section 8.12: Albert's two-slit scenario with a recording pointer ---
% The KL-minimal fit of Section 6.6, restricted to the support since these
% distributions contain zeros.
function G = fit(A, p0, p1, iters)
    n = length(p0); G = zeros(n);
    c = find(p0 > 1e-13); r = find(p1 > 1e-13);
    As = A(r,c); v = ones(numel(c),1);
    for k = 1:iters
        u = p1(r) ./ (As * (v .* p0(c))); v = 1 ./ (As' * u);
    end
    G(r,c) = diag(u) * As * diag(v);
    G(:,c) = G(:,c) ./ max(sum(G(:,c),1), 1e-300);
end
% Particle on a ring of N sites; pointer with three states C, A, B.
% The pointer swaps C->A if the particle is at slit A, C->B if at slit B.
N = 64; dt = 0.5; nsteps = 24;
K = zeros(N);
for j = 1:N
    K(j, mod(j,N)+1) = -1; K(mod(j,N)+1, j) = -1;
end
Ustep = expm(-1i*K*dt); Uf = kron(Ustep, eye(3));
SA = [0 1 0; 1 0 0; 0 0 1]; SB = [0 0 1; 0 1 0; 1 0 0];
Urec = zeros(3*N);
for x = 1:N
    Pm = eye(3);
    if x >= 25 && x <= 32, Pm = SA; elseif x >= 33 && x <= 40, Pm = SB; end
    Urec(3*x-2:3*x, 3*x-2:3*x) = Pm;
end
slit = @(s) circshift([1; zeros(N-1,1)], s-1);
psi = zeros(3*N,1); psi(3*28-2) = 1; psi(3*36-2) = 1; psi = psi/norm(psi);

% ring distance, and the pointer's value in each joint configuration
xs = (1:N)'; dx = min(abs(xs-xs'), N-abs(xs-xs')); dd = kron(dx, ones(3,3));
ptr = repmat((1:3)',N,1); flips = (ptr != ptr');

fkl = []; frs = []; dkl = []; drs = [];
seq = [{Urec}, repmat({Uf}, 1, nsteps)];
for k = 1:numel(seq)
    U = seq{k}; p0 = abs(psi).^2; psi = U*psi; p1 = abs(psi).^2;
    G = fit(abs(U).^2, p0, p1, 1500); % KL-minimal
    Rs = p1 * ones(1, 3*N); % resampling

```

```

if k > 1                                     % after the record has formed
    fkl(end+1) = sum(sum(flips .* (G .* p0')));
    frs(end+1) = sum(sum(flips .* (Rs .* p0')));
    dkl(end+1) = sum(sum(dd .* (G .* p0')));
    drs(end+1) = sum(sum(dd .* (Rs .* p0')));
end
end
max(fkl)                                     % -> 0: the record never changes
mean(frs)                                    % -> 0.5 under resampling
mean(dkl)                                    % -> 0.8 sites per step
mean(drs)                                    % -> 10.6 sites per step
max(max(dx(abs(Ustep).^2 > 1e-12)))          % -> 7: the furthest hop |U|^2 allows

% interference, against the classical sum of the two slits
Ut = Ustep*nsteps;  a = Ut*slit(28);  b = Ut*slit(36);
Pcoh = abs((a+b)/sqrt(2)).^2;  Pmix = (abs(a).^2 + abs(b).^2)/2;
Pptr = kron(eye(N), ones(1,3)) * abs(psi).^2;
tv = @(u,v) 0.5*sum(abs(u-v));
tv(Pcoh, Pmix)                               % -> 0.328: interference present
tv(Pptr, Pmix)                               % -> 0 to machine precision: recording removes it

```

Part III

Restoring Realism

Chapter 9

Realism

Part II established a mathematical fact: quantum mechanics is equivalent to a statistical theory over bits with definite values, governed by non-Markovian laws. Part III defends a philosophical claim: that this is also the *right way to read* quantum mechanics. It is, indeed, very nearly the only way left standing once the alternatives are examined with the care they demand. We conduct the examination. It establishes three requirements that any acceptable interpretation must meet, and then watches the major interpretations fail them.

1. A physical theory must give its systems definite configurations prior to an observation being made; otherwise it cannot explain what is observed at all.
2. If a physical mechanism of “collapse” is what creates those configurations, the theory containing it inherently deviates from the statistical predictions of quantum theory.
3. The configurations of *other observers* (their sense organs, their memories, their brains) must be not merely definite but *invariant*: the same for every observer who describes them. Otherwise one cannot admit that another observer’s experience of the world is as definite as one’s own, and the position collapses into solipsism.

9.1 Theories that point

How does a mathematical theory manage to be *about* anything?

Imagine being handed a dictionary of a language you do not speak, one written entirely in that language. You look up a word, and its definition is made of other words you do not know. You look those up, and their definitions are made of still more. You can chase entries through the dictionary forever, and at the end of any such chase you will possess an enormous amount of structure (which words connect to which) and no meaning whatsoever. Symbols defined solely in terms of more symbols yield an infinite regress in which nothing ever touches the world.

What breaks the regress is nothing inside the dictionary. Words mean something because they are *used*: woven into the practical, real-world activities of the people who speak them: asking for things and receiving them, following directions and arriving somewhere, reporting what one sees and being checked against what is there. Ludwig Wittgenstein called such a practice a *language game*: an activity in which words and deeds are interwoven, and in

which whether a word has been applied correctly is settled within the practice itself, by what speakers do, encounter, and accept, rather than by further definitions.¹

Wittgenstein's advice for finding out what a word means follows from this, and it is not to look the word up. It is "don't think, but look!"² Stop hunting for the definition, and attend instead to what is actually *going on* when the word is used. Someone says *look at that dog over there*. What "dog" picks out on that occasion is not settled by a definition anyone is consulting, but by the whole situation in which the words are said, taken in at once: the animal, the direction of the speaker's gaze, the street, what the hearer does next. The word is made concrete by that situation, and it is that concrete situation, rather than any definition, that the word is answerable to. Remove the situation and nothing has been picked out at all.

This is what it will mean, for the rest of the chapter, to say that a theory *points*: not that its symbols are carefully defined, but that somewhere in the practice of using them one runs up against something that is not another symbol. A word refers to something in the world when there is an established practice of applying it in which that something can actually be met with: referenced in practice, encountered, observed. This is also the force of Wittgenstein's *rule-following problem*: no collection of symbols fixes, by itself, how the symbols are to be applied, because any rule for applying a rule is only more symbols, and the regress ends not in further rules but in application, that is, in what is actually done.³

Benoist calls the position that results *contextual realism*. Concepts, rules, and theories are *ideal*; what they are applied to is *real*; and a concept has content only where there is a context of use in which reality answers to it, meeting it or failing to meet it. The reality that does the answering is not produced by the concept and is not itself a concept.⁴ Benoist's *context* is the concrete situation just described: what is going on in the practice at the moment the term is applied, and what gives the term something to answer to. The word does more work here than its everyday sense of surrounding circumstances, though that is certainly part of it.

Pris has carried this frame into physics: a physical theory, quantum mechanics included, functions as a Wittgensteinian rule, a norm for measuring reality, and a measurement is an application of that norm within the language game of a concrete experimental context. The measurement problem, Pris argues, has exactly the structure of the rule-following problem: it is what the gap between a rule and its application looks like when the rule is a physical theory.⁵

Mathematics, too, is a language, and a mathematical model of the world is a description of the world in the language of mathematics. What is true of language in general is true here: in any mathematical model of reality, some symbols must not simply be defined in terms of additional mathematical symbols, but must be applied, in practice, to something which can actually be observed: they must participate in the language game of the laboratory, not merely in the web of definitions. These anchor symbols have a special name: *observables*. If an observable possesses a particular value, this implies that such a value can actually be

¹L. Wittgenstein, *Philosophical Investigations*, trans. G. E. M. Anscombe (Oxford: Blackwell, 1953), §§7, 23, 43.

²Wittgenstein, *Philosophical Investigations*, §66.

³Wittgenstein, *Philosophical Investigations*, §§143–242, especially §201. The problem is developed at book length in S. A. Kripke, *Wittgenstein on Rules and Private Language* (Cambridge, MA: Harvard University Press, 1982).

⁴J. Benoist, *Toward a Contextual Realism* (Cambridge, MA: Harvard University Press, 2021).

⁵F.-I. Pris, "The Real Meaning of Quantum Mechanics," arXiv:2107.10666 (2021).

found in the world: the particle's position is an observable, and if its position has a particular value, then we can be assured that if we were to observe the particle, we would indeed find it at that position.

Of course, many observables require a measurement apparatus, and cannot be read with the naked eye. This does not particularly matter: the observation of the readout of a measuring device is still an observation of something in the world, since the pointer is as real as the particle, and a theory whose observables anchor to pointer readings is still a theory that points. What a theory cannot survive is having *no* anchors at all: a physical theory in which every symbol is defined in terms of more symbols, and none refers to anything observable, is a formalism about nothing.

Quantum mechanics has an interpretation that embraces the condition deliberately. There are those who hold that the quantum state is not a description of statistics but the entirety of what exists: that the universe simply *is* ψ , evolving unitarily, forever, with no collapse and no further ingredients. This is the Many Worlds interpretation, and Chapter 5 already convicted it once: a hidden Markov memory is by definition hidden, measurement only ever reveals bits, and a theory whose sole ontology never appears in any observation has made measurement impossible rather than explained it.

That argument can now be made in full, because we now know exactly what kind of object ψ is. Its polar decomposition is a probability distribution and a sufficient statistic over the history of that distribution, both of them *statistical* objects, neither of them assigning a definite value to any observable. To say the world is only ψ is to say the world is only statistics: a distribution with nothing distributed, ignorance with no fact to be ignorant of. Every symbol in such a theory is defined in terms of further symbols; nothing points; the dictionary has closed on itself.

Nature cannot consist of an entirely invisible mathematical object, and a reading that says it does owes a method for connecting the object back to definite values for the observables. Consider a simple experiment: a bit is placed into a superposition of states by the Hadamard gate, and then it is measured. The measurement returns a definite value on the computational basis, every time. Explaining that outcome requires a physical mechanism by which the fundamental components of the theory, the things it says constitute physical reality, give rise to what is observed in the experiment, and so to a definite value on the computational basis.

In Many Worlds the mechanism is usually argued for in terms of a branching structure. Introduce a second bit for the arrangement of the physical matter in Alice's brain that gives rise to a mental state of what she believes she saw: a value of 0 is a brain holding a memory of having perceived 0, and a value of 1 is a memory of having observed 1. The physical interaction that writes the memory is a recording, and the abstract operator for it is the CX of the earlier chapters, the Hadamard-and-CX program of Section 3.5, copying the measured bit onto her mental state. After the two gates the pair is in the state

$$\frac{1}{\sqrt{2}} (\mathbf{e}_{00} + \mathbf{e}_{11}),$$


and on the branching reading the state splits along its two terms into two worlds: in one, the bit holds 0 and Alice remembers seeing 0; in the other, the bit holds 1 and Alice remembers seeing 1. Both worlds are in the state, and each holds a definite outcome, which is what the experiment demanded.

That is one carving of the state, and the state does not settle it. The same vector admits

other readings at once,

$$\frac{1}{\sqrt{2}}(\mathbf{e}_{00} + \mathbf{e}_{11}) = \frac{1}{\sqrt{2}}(\mathbf{He}_0 \otimes \mathbf{He}_0 + \mathbf{He}_1 \otimes \mathbf{He}_1) = \frac{1}{\sqrt{2}} \left(\frac{1}{2} \begin{pmatrix} 1 \\ 1 \\ 1 \\ 1 \end{pmatrix} + \frac{1}{2} \begin{pmatrix} 1 \\ -1 \\ -1 \\ 1 \end{pmatrix} \right),$$

the same vector written three ways, as the chapter's code confirms. The third form is the second with its two terms multiplied out. Each term is spread over all four configurations. Adding them cancels the middle two entries and leaves the outer two, which returns the left-hand side. Branches have to be mutually *orthogonal*, meaning their overlap $\psi^\dagger \varphi$ is zero, so a system in one of them is never found in the other. That is what fixes the count at two. A vector can always be chopped into more pieces than that, but pieces that are not orthogonal cancel against one another where they meet, and cancellation is interference, which is the one thing branches must not do. Both carvings above are orthogonal pairs. In the first each branch is an eigenstate of the computational basis; in the second neither one is. Each side splits it cleanly into two, and the two splittings describe Alice differently:



Either way each branch contains a brain in some configuration, and what differs is the configuration. On the left every branch has the measured bit holding one value and Alice's bit holding the matching record, which is the picture the interpretation is usually sold on. On the right the same vector is cut so that in each branch the measured bit is spread across both values and Alice's memory is spread across both records: ask either Alice what she saw and the answer is 0 or 1 with equal probability. The two carvings do not disagree about *which* outcome Alice saw. They disagree about whether anything in the branch has a value at all.

The carving therefore settles what a world contains. On the left each branch holds the intuitive picture: amplitude 1 on the value Alice takes the bit to have and 0 on the other, with a second branch where the two are exchanged. On the right each branch holds amplitude $1/\sqrt{2}$ on both values at once, and Alice's memory in the same shape, the two branches differing only in a sign. Section 8.4 argued that a wave spread through space cannot be what the state describes; on the right-hand carving every branch contains exactly such a spread. Nor are these the only two carvings. Rotating both bits through any angle splits the state just as cleanly, and out of that continuous family the one aligned with the record basis is the only carving whose branches hold definite values at all.

This is the *preferred basis problem*, and what makes it bite here is the interpretation's own selling point. ψ is supposed to be everything: no collapse, no added variables, nothing outside the state. The branches must therefore be structures in ψ and in nothing else, and the state displayed above does not contain them.

The usual answer is that decoherence settles the matter: brains and their environments interact in the record basis, so that basis is dynamically singled out. The interaction is real physics and nobody chose it.

But notice what the answer requires, and where the burden comes to rest. A record basis is defined by splitting the world into a system and an environment that records it. Any description that makes such a split can be combined into a single description that makes none: write one wavefunction over the configurations of system and environment together. The environment, moreover, is everything left out of the experiment, so the combined description covers everything there is, and the wavefunction it carries is a wavefunction of the whole universe. That is the *universal wavefunction*, and it is no import of the critics: it is what any laboratory description becomes once the split is removed. The question the decoherence answer must face is therefore a question set inside the universal wavefunction itself. To use a split at all, the account must show that the split has physical meaning there, that ψ itself distinguishes a system on one side from an environment on the other. If it does not, the split is arbitrary, and an arbitrary split can be drawn in many ways. Galvan states the difficulty plainly: defining pointer states requires decomposing the global system into system and environment, “and it is not at all clear how to perform such a decomposition when the global system is the universe.”⁶

A split can of course be drawn by hand, and Section 9.3 will call that what it is: a decision to stop tracking. But a decision is not part of ψ , and doing without anything beyond ψ was the undertaking. And whoever draws it already has in mind which divisions are the classical ones, and a definite classical reality is what the account set out to derive. What is needed is a criterion fixing the split from ψ and its evolution alone, and Zurek, whose work made decoherence the standard account, grants that none is in hand: the question of what the “systems” are is “looming big, as a foundation of the whole decoherence program.”⁷

Kent stated the general form of the difficulty before decoherence became the standard reply, and the standard reply does not touch it: “no preferred basis can arise, from the dynamics or from anything else, unless some basis selection rule is given.” Dynamical processes that destroy phase information are available in abundance, and none of them carries a system from a state expandable in any basis whatever to a claim about one basis in particular. That step has to be supplied as a rule. Kent then allowed the record basis as such a rule, framed by reference to measuring devices, and declined the result on a ground that applies here as well: a rule of that kind makes complicated macroscopic objects, “whose precise definition must ultimately be arbitrary,” into fundamental quantities. His reasons there, he notes, are Bell’s.⁸

⁶B. Galvan, “On the preferred-basis problem and its possible solutions,” arXiv:1008.3708 (2010), §3. Galvan separates the problem of decomposing a given state, which is the one at issue here, from the further problem of singling out a preferred basis for the whole space.

⁷W. H. Zurek, “Decoherence, einselection and the existential interpretation (the rough guide),” *Philosophical Transactions of the Royal Society of London A* 356 (1998), 1793–1821, at 1820. His later synthesis does not supply one: it makes the division an axiom, “The Universe consists of systems” (“Relative States and the Environment: Einselection, Envariance, Quantum Darwinism, and the Existential Interpretation,” arXiv:0707.2832 (2007), §I). M. Schlosshauer, “Decoherence, the Measurement Problem, and Interpretations of Quantum Mechanics,” *Reviews of Modern Physics* 76 (2004), 1267–1305, §III.A, reports the same absence: there is “no general criterion for how the total Hilbert space is to be divided into subsystems.”

⁸A. Kent, “Against Many-Worlds Interpretations,” *International Journal of Modern Physics A* 5 (1990), 1745–1762. Stoica has since argued the corresponding negative result formally: wherever a factorization, a space structure, or a preferred basis does emerge from the state vector and its evolution alone, it emerges in many physically distinct versions rather than one (“3D-Space and the preferred basis cannot uniquely emerge from the quantum structure,” *Advances in Theoretical and Mathematical Physics* 26 (2022), 3895–3962). That last claim is under dispute, a counterexample having been offered and answered while this book was being written.

Everettians have conceded as much about the branching itself. Sebens and Carroll, whose derivation of the Born rule from self-locating uncertainty is the interpretation's standing defense of its probabilities, write: "The branching structure of the wave function in Everettian quantum mechanics is not uniquely defined by the theory, which is just a state space and a dynamical law; like coarse-graining in classical statistical mechanics, there is a subjective element in how it is constructed. We must decide what kind (and what degree) of separation is required for two putative branches to count as distinct (this involves choosing an appropriate system/environment split, or, in a consistent-histories approach, choosing which coarse-grained observables to use)."⁹

The analogy does not survive inspection. In classical statistical mechanics the coarse-graining groups microstates into macrostates, and the grouping is indeed a choice; the choice is harmless because the fine grain already holds the facts. A particle has its position and its momentum whatever the grouping, and a macrostate is a summary of definite properties that stand on their own. The branching of ψ is no summary. The grouping determines the branches, the branches are what hold the values, and under a wrong grouping no branch holds a definite value for the observable the observer measured. A grouping that arranges already-definite facts can be drawn as arbitrarily as its user likes; a grouping that supplies the only definite facts in the theory cannot. If the facts of observation depend on the grouping, then the grouping is part of the physics, it must be fixed by something inside ψ , and until it is, the reading makes no empirical prediction at all.

Every account of the branching on offer is about an agent. Decoherence depends upon a split between a system and an environment, and neither half of the dependence is given a physical account: the theory fixes neither which branches there are, as Sebens and Carroll concede above, nor what counts as the system and what counts as the environment, as Galvan's objection records. The split is a description that has stopped following part of the world, as Section 6.3 put it, and nothing physical marks where the following stopped. The decision-theoretic program asks what a rational agent within a branch ought to bet. Self-locating uncertainty asks what such an agent does not know about which branch they occupy. An agent's impression of a definite outcome was the thing to be accounted for. Kent raises the objection against the last of them. Self-locating uncertainty is meant to be uncertainty about an objective fact, "either I am the observer on the up branch or the observer on the down branch, but I'm not sure which," and he asks "how can these be possible facts about reality, when the choice of branches is agreed to be not objective, but rather an arbitrary mathematical choice which — remember — can be made in many generally incompatible ways in generic models?"¹⁰

A last retreat routes the connection through consciousness. What is observed, some of the interpretation's advocates say, just is experience, and experience is the subject of the problem of consciousness, which no one has solved. Penrose puts the demand in its sharpest form: "a theory of consciousness would be needed before the many-worlds view can be squared with what one actually observes."¹¹ Chalmers, who has offered his own theory of consciousness in

⁹C. T. Sebens and S. M. Carroll, "Self-Locating Uncertainty and the Origin of Probability in Everettian Quantum Mechanics," *British Journal for the Philosophy of Science* 69 (2018), 24–75, §2.2.

¹⁰A. Kent, "Does it Make Sense to Speak of Self-Locating Uncertainty in the Universal Wave Function? Remarks on Sebens and Carroll," *Foundations of Physics* 45 (2015), 211–217.

¹¹R. Penrose, *The Emperor's New Mind: Concerning Computers, Minds, and the Laws of Physics* (Oxford: Oxford University Press, 1989), quoted in D. J. Chalmers, *The Conscious Mind: In Search of a Fundamental*

the interpretation's support, agrees: of Everett's assumption that a superposed brain holds distinct subjects of experience, the matter "depends crucially on a theory of consciousness."¹²

Aaronson, entering the reading on its own terms, states the move in its plainest form: "If you wanted me to have an ontology that was as simple as possible, then I would say it's going to look like a wave function that is evolving by unitary transformation. And the hard part, of course, is how do I get out of that my subjective experience. . . . But, at that point, I feel like that's going to be entangled with enormous questions about what is an observer, what is experience, what is consciousness?"¹³

The retreat is refused here, and not because the problem is unsolved. No physical theory has ever had to solve the problem of consciousness to make a single prediction: the symbols that anchor to observation anchor to it by definition, which is where this section began. An observation is a configuration of a physical instrument and a physical brain, and connecting configurations to the rest of a physical theory is the theory's own job. A reading that hands that connection to an unsolved problem has not explained the observation; it has renamed the failure to explain it.

Two descriptions are left with no stated relation between them: the universal wavefunction on one side, and on the other the definite world of the living room and the kitchen, with nothing in the first that fixes the second. That is a kind of dualism. Ilyenkov, following how the Cartesians handled a thinking substance and an extended one with no attribute in common, records the answer they settled on: "God knows, answered Descartes, and Malebranche, and Geulincx; from our point of view it is inexplicable. Only God can explain this fact. He makes the two opposing worlds agree."¹⁴ Keeping ψ and the definite world both, and supplying no physical process that joins them, arrives at the same place by another road.

Many Worlds is offered as taking on nothing beyond ψ and its unitary evolution, and it is that claim the argument refutes. Recovering a memory means recovering a definite configuration of a brain, which configurations the branches hold depends on how ψ is carved, and an assumption about the carving has to be supplied from somewhere. A version that supplies one is a repair of a repair, Bell's and then this, and is no longer the interpretation as advertised.

Needing such an assumption is not by itself a defect, and Chapter 10 will find that any realist reading needs one. The difference is where it goes. Chapter 10 postulates a preferred basis at the outset rather than hoping to derive one as something emergent, and once it is postulated, the systems have definite configurations from the start. Nothing then recommends having stripped them of those configurations in the first place. On that basis the values can be assigned consistently, which does away with the need for a multiverse to begin with.

The point is not that Many Worlds is extravagant, though Bell noted that too: "this multiplication of universes is extravagant, and serves no real purpose in the theory, and can simply be dropped without repercussions." The point is that it is *incomplete in the one place*

Theory (New York: Oxford University Press, 1996), p. 348.

¹²D. J. Chalmers, *The Conscious Mind: In Search of a Fundamental Theory*, p. 348; the offer of support is at p. xv.

¹³S. Aaronson, in discussion with J. Barandes, hosted by C. Jaimungal, "Harvard Scientist Rewrites the Rules of Quantum Mechanics | Scott Aaronson & Jacob Barandes," video (3 March 2025). Quoted from the author's transcription.

¹⁴E. V. Ilyenkov, *Dialectical Logic: Essays on its History and Theory*, trans. H. Campbell Creighton (Moscow: Progress Publishers, 1977), Essay One.

a physical theory cannot be incomplete: at the interface with observation.

The incompleteness is total, and it is worth saying exactly why. The definite outcome was the appearance the interpretation owed an account of, and the branching that was to supply the account is, on the interpretation's own admission, a subjective carving. An illusion has been grounded in an illusion, and a theory built that way has no connection left to the empirical world: nothing in it states, as an objective fact, how what it declares physical gives rise to what any observation finds. The mechanism the experiment demanded at the outset is nowhere in it: no physical process gives rise to the branches, and without the branches there is no objective way for the measured bit to come to hold a definite value. A theory with no such connection cannot make a single empirical prediction. Even the predictions quantum mechanics already had are gone from it, for the Born statistics were predictions only because they were probabilities over configurations that readings could find, and the configurations are what the reading removed. What remains is a formalism about nothing, in the exact sense of this section: a dictionary with no entry that touches the world. The first requirement of this chapter's list, systems with definite configurations prior to observation, is met nowhere in it.

9.2 The need for systems

A *system* is a thing described by a configuration, and its configuration is the definite values it holds. The observables are what a configuration supplies: some are read directly off it, and the rest arise from it through the physical process of measuring, a division Chapter 10 will make precise. The dependence runs one way only. Fix the configuration and the observables follow; there is no fixing the observables and hoping a configuration will follow from them.

Now let a system's configuration change randomly, as it does in a probabilistic computer and, if we are right, in the world. Then it is impossible to track which configuration it is in at all times; we can track, at best, a probability distribution over them. The mathematical model now contains no definite configurations, only statistics, and this raises the question of how a statistical model connects back to the world at all.

The resolution is one our reader has already built with their own hands. We make a distinction between the *ontological* state of a system and our *epistemic* description of it. In the real world, systems possess definite configurations; we do not know which configuration a system is in; the distribution expresses that ignorance while remaining anchored, through the configurations it ranges over, to what observation would find. Chapter 6 shows by explicit construction that this reading is available for quantum mechanics: assign the bits definite values, evolve them by stochastic hops at every time interval, and the ensemble reproduces the Born statistics at every moment. In each individual run the bits really have values; over many runs the statistics emerge; when we observe, we merely learn which value was there. A statistical theory read this way points at the world through its sample space.

That this was Einstein's actual concern (and not, as the popular telling has it, a squeamishness about randomness) is captured in an anecdote recounted by the physicist Abraham Pais.¹⁵ Walking home from the Institute for Advanced Study around 1950, Einstein suddenly stopped, turned to him, and asked whether Pais really believed that the moon exists only

¹⁵A. Pais, "*Subtle is the Lord...*": *The Science and the Life of Albert Einstein* (Oxford: Oxford University Press, 1982), p. 5.

when one looks at it. The conversation, Pais stresses, was not loose metaphysics; it was about what is doable and knowable in the sense of physical observation.

“God does not play dice” is quoted far more often, but the moon question is the deeper one: dice at least have faces before they are read. A theory whose moon has no position until someone looks has subtracted the world.

The letter that made the ‘dice’ phrase famous makes the same point. Writing to Max Born in December 1926, Einstein followed the remark about dice with a complaint about where Schrödinger’s waves live:¹⁶

Quantum mechanics is very impressive. But an inner voice tells me that it is not yet the real thing. The theory produces a good deal but hardly brings us closer to the secret of the Old One. I am at all events convinced that He does not play dice. Waves in $3n$ -dimensional space whose velocity is regulated by potential energy (e.g., rubber bands)...

The last two sentences are one thought. A wave in $3n$ -dimensional space is not a wave anywhere in space: for n particles it is a single object in a space of configurations, and the arrangement of things in three dimensions has dropped out of the description. It is hard to see how something that can be placed nowhere in three-dimensional space exists anywhere in the world, and that is the moon question in another form: a moon with no position and a wave with no location are the same difficulty, raised a quarter-century apart.

One reading of quantum mechanics embraces that subtraction, and Section 9.1 has already declined it: on the Many Worlds interpretation the universe simply *is* ψ , the worlds are carvings with no physical mark, and nothing in it points. This section follows a second dependence, which the interpretation disowns and nevertheless has: it needs systems, in definite configurations, at every point where it touches observation.

Maudlin has discussed this objection at length, and two of his points are relevant here.¹⁷

The first concerns what the Born-rule numbers are *for*. It is not enough, Maudlin argues, that a theory supply probabilities for *something physical* to happen; they must be probabilities for the *right kind of thing*, and the right kind of thing is what he calls *manifest* variables: quantities, like the configuration of a laboratory pointer, whose values people can reliably determine just by looking. A theory whose probabilities attach to something invisible and intangible, he notes, “could not explain what we want to explain, viz., things that we take to happen in laboratories.”

There is an irony in Maudlin’s choice of words, which the physicist John Bell had already noticed decades earlier. The definite-valued quantities of statistical interpretations were saddled, historically, with the name *hidden* variables, yet they are precisely the manifest ones, the only entries in any formalism that anyone has ever seen. Writing on the pilot-wave theory, whose added variable X is nothing but the configuration of things, Bell was blunt: “Although Ψ is a real field it does not show up immediately in the result of a single ‘measurement,’ but

¹⁶Pais, “*Subtle is the Lord...*”: *The Science and the Life of Albert Einstein*, p. 443, translating Einstein to M. Born, 4 December 1926. Pais gives the phrase correctly as $3n$ -dimensional. The canonical English edition of the Born–Einstein correspondence drops the n .

¹⁷T. Maudlin, “Can the World be Only Wavefunction?” in S. Saunders, J. Barrett, A. Kent, and D. Wallace (eds.), *Many Worlds? Everett, Quantum Theory, and Reality* (Oxford: Oxford University Press, 2010), pp. 121–143.

only in the statistics of many such results. It is the de Broglie–Bohm variable X that shows up immediately each time. That X rather than Ψ is historically called a ‘hidden’ variable is a piece of historical silliness.”¹⁸ Elsewhere he proposed the accurate terminology: it “would be appropriate to refer to the x s as ‘exposed variables’ and to ψ as a ‘hidden variable.’ It is ironic that the traditional terminology is the reverse of this.”¹⁹

What is hidden (perfectly, permanently, in every individual run) is the wavefunction, and our Part II explains why the backwards naming nevertheless felt natural: the phase is a statistic of the history, and statistics of histories are not the kind of thing that *could* appear in a single glance.

The second point concerns a quiet cheat in the monistic vocabulary. The wavefunction is standardly said to live on *configuration space*, and a configuration space is, by definition, a space whose every point represents a configuration of things in ordinary space. The word smuggles the missing ontology in through the name. If the world contains only ψ , there are no things and no configurations of them, and each point of the high-dimensional space represents nothing at all. In Maudlin’s words: “in a monistic theory, all of this, from beginning to end, is a cheat: there are no configurations at all in the ontology, and so nothing else for a particular state of the wavefunction to ‘represent’ or ‘be associated with.’” Two wavefunction states may differ in where their amplitude is massed over that space, and the theory may assign the two lumps different probabilities, yet nothing in the monistic ontology says which lump is *the pointer reading 1* and which is *the pointer reading 2*, because there is no pointer. Our vocabulary makes the same diagnosis from the constructive side: the entries of ψ are indexed by configurations of bits (that is what Chapter 1 built), and a theory that keeps the indexing while deleting the bits has kept a filing system and discarded the files.

That diagnosis reaches past the vocabulary and into the mathematics. A state vector is a list, and a list means nothing until it is indexed: this entry belongs to that configuration of those systems, and without the assignment there are numbers and no state. The assignment is nowhere written in ψ ; it has to be in hand before ψ can be written down at all. So the wavefunction makes sense only if there are already systems which can be in definite configurations, and the notation presupposes them rather than describing them. Every reading that starts from ψ alone has to bring them back, as a postulate either implicit or declared, and the rest of this section is a survey of where each one does it.

The physicist Carlo Rovelli, in his book *Helgoland*, reaches the same verdict from inside physics:²⁰

The gigantic, universal ψ wave that contains all the possible worlds is like Hegel’s dark night in which all cows are black: it does not account, per se, for the phenomenological reality that we actually observe. In order to describe the phenomena that we observe, other mathematical elements are needed besides ψ : the individual variables, like X and P , that we use to describe the world. . . . It

¹⁸J. S. Bell, “On the impossible pilot wave,” *Foundations of Physics* 12 (1982), 989–999, at 993; reprinted in J. S. Bell, *Speakable and Unspeakable in Quantum Mechanics* (Cambridge: Cambridge University Press, 1987).

¹⁹J. S. Bell, “Quantum mechanics for cosmologists,” in C. J. Isham, R. Penrose, and D. W. Sciama (eds.), *Quantum Gravity 2* (Oxford: Clarendon Press, 1981); reprinted in *Speakable and Unspeakable in Quantum Mechanics*, pp. 117–138, at 128.

²⁰C. Rovelli, *Helgoland: Making Sense of the Quantum Revolution*, trans. E. Segre and S. Carnell (New York: Riverhead Books, 2021).

is not enough to know the ψ wave and Schrödinger's equation in order to define and use quantum theory: we need to specify an algebra of observables, otherwise we cannot calculate anything and there is no relation with the phenomena of our experience.

Bell examined the Everett program at length in “Quantum mechanics for cosmologists,” and his diagnosis is followed here step by step, because it uncovers a failure beyond the missing observables: the theory also loses the *past*.²¹ Bell begins with a repair. The official theory contains only ψ : “the classical variables x do not appear in Everett and De Witt. However, it is taken for granted there that meaningful reference can be made to experiments having yielded one result rather than another.” Everettians describe experiments as having yielded results, and an ontology of ψ alone contains nothing that could have yielded them. So Bell supplies the missing ingredient on Everett's behalf, and studies “the pilot-wave theory without trajectories. Thus instantaneous classical configurations x are supposed to exist, and to be distributed in the comparison class of possible worlds with probability $|\psi|^2$.” The question mark he attaches to the name, “Everett (?)” concedes that the repaired theory may be more his than Everett's.

The criticism that follows is therefore aimed at the strongest version available, one granted the first requirement outright: definite configurations at every instant, Born distributed exactly as ours are. What the repaired theory still refuses is a law pairing the instants: “no pairing of configurations at different times, as would be effected by the existence of trajectories, is supposed.” The pairing is what our transition matrices supply. Everything the pilot-wave theory strings into a continuous history, Many Worlds leaves as unstrung instants; keeping the configurations while “discarding the trajectory,” Bell says, “is the essential (in my opinion) of the theory of Everett.”

Consider what that discards. Walk from your living room to your kitchen. You have a clear memory of the traversal: the doorway passed, the floor crossed, one continuous path through space. In Many Worlds no such path exists. The configuration that includes *you in the living room* and the later one that includes *you in the kitchen* are joined by no fact whatever; no trajectory was actually traversed. So the memory is a memory of something that never happened. Everett saw this and embraced it: the theory, he held, need not deliver the past at all, only explain why each branch *contains memories*, and the memories, being present configurations of brains and notebooks, are all the theory needs to supply. Bell states the position exactly: “in our interpretation of the Everett theory there is no association of the particular present with any particular past. And the essential claim is that this does not matter at all. For we have no access to the past. We have only our ‘memories’ and ‘records.’ But these memories and records are in fact *present* phenomena.”

That is a brain-in-a-vat scenario promoted to fundamental physics: all of your memories are, strictly, lies: accounts of trajectories no one ever took. The theory's explanatory job ends with explaining why the lies are consistent. Bell's verdict: “Everett's replacement of the past by memories is a radical solipsism — extending to the temporal dimension the replacement of everything outside my head by my impressions, of ordinary solipsism or positivism. Solipsism cannot be refuted. But if such a theory were taken seriously it would hardly be possible to take anything else seriously.” He then recalled the last theory to demand that present

²¹Bell, “Quantum mechanics for cosmologists,” pp. 133–136, from which the quotations in this section and the next three paragraphs are drawn.

records not be taken to imply any past: the creation of the world in 4004 BC, defended in the 18th century on the grounds that God would naturally have created a going concern: trees complete with annular rings for years that never elapsed, rocks bearing fossils of creatures that had never lived. The comparison is harsh, and it fits. Where the argument lands matters, because we will meet the destination again: Many Worlds is standardly advertised as the sober realist alternative, the interpretation that takes the formalism seriously without collapse, and yet, even after Bell's repair, it terminates in a solipsism of the present moment. We will meet it again at the third requirement below.

That bears on the retreat described above, and not in its favor. Everett's position surrenders the past and keeps the memory: the record is a definite present configuration of a brain, and what it lacks is anything it is a record of. The carvings threaten the half that was kept. Whether a branch holds a definite record at all depends on how the state is cut, as Section 9.1 showed, and ψ does not do the cutting. Retreating to present phenomena is therefore not retreating far enough to find ground, since the present phenomena were to be definite and that is the thing not secured.

Each account here puts the systems back, and what distinguishes them is where. Everett's needs a basis; a basis needs the world divided into a part that records and a part recorded; and Galvan's objection is that no such division is available once the world is all there is. Zurek, whose work made decoherence the standard account of how a basis gets singled out, does not derive the division. He postulates it, setting down as an axiom that the universe consists of systems and observing that without them the measurement problem cannot be posed at all.

Rovelli postulates the systems too, and says so in as many words, though he withholds the configurations until an interaction supplies them. Relational quantum mechanics begins by assuming "that the world can be decomposed (possibly in a large number of ways) into a collection of systems, each of which can be equivalently considered as an observing system or as an observed system," with information passing between them through physical interactions.²² Bell's repair keeps the memories, which is to say it keeps brains in definite configurations holding definite mental states, and surrenders the past to do it. The assumption arrives in the indexing, or in an axiom, or in a repair, and it always arrives.

The observables (the algebra of them, the anchor symbols) are exactly what the "only ψ " ontology declines to supply, and Rovelli's remark says why they cannot be recovered from ψ alone: what has to be added is an algebra of observables. By the dependence set out at the start of this section, that is the same thing as saying what has to be added is systems in definite configurations, since the observables come from the configuration and from nowhere else. Standard Everett supplies no such thing, and even Bell's repair, as we have already seen, terminates in a solipsism of the present moment.

The carvings of Section 9.1 are that same gap, met where the basis has to be chosen. Reaching definite records meant supplying a basis, which is to say supplying by hand the very thing the ontology declined to include. That is the shape every repair here has taken: the interface with observation is where the account runs out, closing it takes an assumption from outside ψ , and an interpretation recommended for needing nothing outside ψ has little left to recommend it once the assumption is in.

Notice what is absent in each case, because it is the same thing every time. The "only ψ "

²²C. Rovelli, "Relational quantum mechanics," *International Journal of Theoretical Physics* 35 (1996), 1637–1678, §3.1.

ontology leaves the systems with no configuration to be in. Standard Everett gives a branch none until a basis is supplied from outside. Bell’s repair keeps the records and empties them of anything recorded. Each account fails at the moment it has to say what the measurement found, and it fails because there was nothing there for the measurement to find. What a physical theory cannot do without is systems holding definite configurations.

The same conclusion arrives from what a probability distribution is. A distribution is a statement of ignorance: it says the system is in one of these configurations and declines to say which. Take the configurations away and the distribution is no longer ignorance about anything. It is a rule for producing numbers, and the question of what the numbers are about has lost its answer.

The demand is narrower than it may appear. It does not ask that every observable carry a value, which Chapter 10 will show cannot be had. By the dependence set out at the start of this section, the configuration is what the observables come from. Some may be read straight off it. Others may be produced from it at the moment of measurement, by a real physical process, which is the account Section 10.5 gives. Either way, what no theory can do without is the configuration they are produced from.

9.3 Wigner’s friend

Suppose one accepts the first requirement but tries to satisfy it minimally: a system *acquires* a definite configuration, but only at measurement; before that, it has none. This last clause is the doctrine of *value indefiniteness*: the claim that an observable has no definite value until a measurement confers one, so that between measurements there is no fact of the matter about what the observable is. It is the position the first requirement was framed to deny, and the rest of Part III is in large part an argument against it. All that exists inside a running quantum computer is the quantum state; when the memory is read out, the quantum state “collapses” into bits with definite values. This is the popular picture, the Copenhagen interpretation as transmitted by the Dirac–von Neumann axioms, and it is usually offered with a shrug: the collapse happens at measurement; by the time anyone looks, the values are there, so where is the problem?

The problem, first pointed out by the physicist Eugene Wigner, is that observers can observe *each other*, and Einstein’s moon question becomes considerably sharper when the moon is replaced by a person. Wigner imagined a friend who performs a measurement inside a closed laboratory while he himself waits outside.²³ Call the friend Alice, and the observer outside Bob.

We can pose the scenario with circuits, since a mental state is, for our purposes, information: Alice’s brain records what she saw, and one bit suffices to model the record. A bit is prepared in a superposition and measured by Alice inside the sealed lab; the CX records the bit’s value onto Alice’s memory:

$$\begin{array}{c} \mathbf{e}_0 \\ \mathbf{e}_0 \end{array} \quad \begin{array}{c} \text{H} \\ \\ \end{array} \quad \begin{array}{c} \overline{\text{CX}} \\ \text{CX} \end{array} \quad \psi = \frac{1}{\sqrt{2}} \begin{pmatrix} 1 & 0 & 0 & 1 \end{pmatrix}^T.$$

²³E. P. Wigner, “Remarks on the mind–body question,” in I. J. Good (ed.), *The Scientist Speculates* (London: Heinemann, 1961); reprinted in E. P. Wigner, *Symmetries and Reflections* (Bloomington: Indiana University Press, 1967), pp. 171–184.

From Alice's perspective, she measured the bit and perceived a definite value; her mental state is definite, 0 or 1. But Bob, outside the sealed lab, knows the experimental protocol and not the outcome. If Bob describes the lab by the axioms, with unitary evolution until *he* measures, his description is the entangled superposition above: a state in which Alice's memory bit has no definite value. Asked "what is Alice's mental state?", Bob's description answers: there is no fact about it yet.

One bit is a caricature of a person. A real Alice interacts with the air in the room, with the light, with the apparatus, and her memory is not one bit but an enormous number of them. None of that changes the argument. Each further interaction is one more unitary entangling one more system into the same state, and Section 4.10 established what that does: the joint state stays pure, and the superposition grows to take in whatever has been drawn into it. Tracing those extra degrees of freedom out is a decision to stop tracking them. It is an operation on the description, and the system is untouched: keep them in the account and the superposition is there in full. Bob's description of the sealed laboratory is a superposition still, and a larger one. Enlarging the system never produces a definite value.

So the collapse doctrine faces a dilemma. Either the collapse is *variant*: it happened for Alice but not for Bob, each observer's description being valid "for them"; or it is *invariant*: Alice's measurement collapsed the state for everyone, Bob's entangled description is simply wrong, and the lab is really in one of the two definite configurations. Each option is a requirement violated, and we take them in turn.

9.4 Variant collapse is solipsism

Take the first option: the state of the lab, including Alice's mental state, is variant, definite from Alice's perspective and indefinite from Bob's. This is the considered position of the relational interpretation of quantum mechanics, in which the values of all physical quantities are relative to the observer interacting with them,²⁴ and the implicit position of any reading on which "collapse" is just each observer's bookkeeping.

That position is answerable to Alice's mind. Her mental state is not the workings of an ethereal soul, exempt from the physics. It is the physical substrate of her experience, made of particles whose configurations are bits like any others, and so whatever the theory says about bits it says about her mind. If there is no invariant fact about which value her memory holds, there is no invariant fact about *what she is experiencing*: whether she saw a 0 or a 1, whether she is presently believing, feeling, reporting one thing or another. Her inner life becomes perspectival: definite in her own frame, smeared in Bob's, with no fact of the matter about which is the truth. It is a basic principle of realism, perhaps *the* basic principle, that the minds of other observers definitely exist, regardless of who perceives them. Each of us knows, from the inside, that our own experience is definite; to refuse that same definiteness to others, making the very existence of their mental states depend on whether or not we have looked at them, is to make ourselves the only fully real observer in the world. However carefully the position is stated, it is difficult, in the end, to distinguish from solipsism.

The charge of solipsism can be escaped, at a price. One may refuse that name by granting every observer a world of their own, each furnished with its own facts, none of them the facts

²⁴C. Rovelli, "Relational quantum mechanics," *International Journal of Theoretical Physics* 35 (1996), 1637–1678.

of any other. Each observer is then fragmented into a world of their own, and each of those worlds is solipsistic in the strict sense: its holder is the only true observer in it, and every other observer a fiction within their account. Solipsism is avoided only by granting that this holds for all observers alike, each fragmented from all the rest, none ever able to reach another. Section 11.12 analyzes a reading that accepts exactly this.

This, at bottom, was the danger Dmitry Blokhintsev, whom Chapter 8 left rebutting Einstein on separability, warned against in the same 1951 critique of the idealist readings of quantum theory (a warning we regard as having aged remarkably well).²⁵ The statistical description, he argued, is *subjective*: it reflects an observer's knowledge, as any probability distribution does. Treat that subjective description as an objective feature of the world, promoting the observer's state of knowledge into the state of reality, and reality itself becomes indexed to observers: what exists depends on who has looked, and the observer's knowledge dictates what is real. The conversion of the subjective into the objective does not elevate knowledge; it demolishes the world. One ends, by careful steps, in idealism, and idealism about *other people's minds* is a form of solipsism. Hence the third requirement: *the configurations of other observers must be invariant*. Any interpretation that makes Alice's memory variant fails it, however the failure is described.

Einstein put the same commitment positively. It is where natural science begins.²⁶

The belief in an external world independent of the perceiving subject is the basis of all natural science.

And it is what gives the word *physical* its content.²⁷

Physics is an attempt conceptually to grasp reality as it is thought independently of its being observed. In this sense one speaks of "physical reality."

An interpretation on which what exists depends on who has looked gives up both.

9.5 Invariant collapse is a different theory

Take the second option: the collapse is invariant. When Alice measures, the state collapses for everyone; the lab really is in a definite configuration; Bob's entangled description is simply mistaken. This option, at least, is realist: requirement three is satisfied. The trouble is that it is *physics*, and physics has consequences.

If Alice's measurement really collapsed the state, then after her measurement the lab (the bit plus Alice's memory) is in one of the two configurations at random: a statistical mixture. If no collapse occurred, the lab is in the entangled superposition. Chapter 4 taught us that

²⁵D. I. Blokhintsev, "Kritika idealisticheskogo ponimaniya kvantovoy teorii" [Critique of the idealist understanding of quantum theory], *Uspekhi Fizicheskikh Nauk* 45, no. 2 (1951), 195–228.

²⁶A. Einstein, "Clerk Maxwell's Influence on the Evolution of the Idea of Physical Reality" (1931), the essay's opening sentence; reprinted in *The World As I See It* (New York: Covici-Friede, 1934), p. 60.

²⁷A. Einstein, "Autobiographical Notes," in P. A. Schilpp (ed.), *Albert Einstein: Philosopher-Scientist* (New York: MJF Books, 2001), p. 81.

these two descriptions are different density matrices:

$$\rho_{\text{uncollapsed}} = \begin{pmatrix} 0.5 & 0 & 0 & 0.5 \\ 0 & 0 & 0 & 0 \\ 0 & 0 & 0 & 0 \\ 0.5 & 0 & 0 & 0.5 \end{pmatrix} \quad \rho_{\text{collapsed}} = \begin{pmatrix} 0.5 & 0 & 0 & 0 \\ 0 & 0 & 0 & 0 \\ 0 & 0 & 0 & 0 \\ 0 & 0 & 0 & 0.5 \end{pmatrix},$$

identical on the diagonal, so no direct readout distinguishes them, yet different in the corners. And Chapter 4 also taught us what different density matrices mean: *some* further evolution distinguishes them empirically. Here one suffices to apply a Hadamard to both bits and then measure:

$$\begin{aligned} (\mathbf{H} \otimes \mathbf{H}) \rho_{\text{uncollapsed}} (\mathbf{H} \otimes \mathbf{H})^\dagger &= \rho_{\text{uncollapsed}} \longrightarrow \begin{pmatrix} 0.5 \\ 0 \\ 0 \\ 0.5 \end{pmatrix} \\ (\mathbf{H} \otimes \mathbf{H}) \rho_{\text{collapsed}} (\mathbf{H} \otimes \mathbf{H})^\dagger &= \frac{1}{4} \begin{pmatrix} 1 & 0 & 0 & 1 \\ 0 & 1 & 1 & 0 \\ 0 & 1 & 1 & 0 \\ 1 & 0 & 0 & 1 \end{pmatrix} \longrightarrow \begin{pmatrix} 0.25 \\ 0.25 \\ 0.25 \\ 0.25 \end{pmatrix}. \end{aligned}$$

If no collapse occurred, Bob's perturbed measurement finds the two bits perfectly correlated; if an invariant collapse occurred, he finds them uncorrelated. The two hypotheses make *different statistical predictions*.

Deutsch sharpened the same point into a full thought experiment: let the observer be one whose sense organs and memory are precisely characterized (he suggests an artificial intelligence, whose internal dynamics its designers know), have it measure a spin, and then reverse the measurement interaction exactly. Unitary quantum mechanics predicts the reversal succeeds, interference and all, with certainty; an invariant collapse, being irreversible, predicts it fails half the time. In his words: “Now, finally, the empirical difference between the ‘interpretations’ has become accessible.”²⁸ (Deutsch deploys the argument on behalf of Many Worlds, which Section 9.2 has already declined; but the argument itself cuts only against physical collapse, and cuts deep.)

Performing such an experiment on a human-scale observer is far beyond any foreseeable technology, because the coherence requirements are astronomical; the physicists Scott Aaronson, Yosi Atia, and Leonard Susskind have made that caveat quantitative.²⁹ In practice, the two hypotheses are indistinguishable. In *principle*, they are not, and principle is what interpretation is about.

Why can no one patch the collapse to hide inside quantum mechanics? The obstruction is structural, and Blokhintsev identified it: linearity. Every process quantum theory allows is linear in the quantum state; the reduction of a superposition to a single definite outcome is not a linear map. No composition of the theory's own machinery can produce it, and any machinery added from outside must, somewhere, at some scale, produce statistics the unmodified theory does not. As Blokhintsev put it, reviewing decades of such attempts.³⁰

²⁸D. Deutsch, “Quantum theory as a universal physical theory,” *International Journal of Theoretical Physics* 24, no. 1 (1985), 1–41.

²⁹S. Aaronson, Y. Atia, and L. Susskind, “On the hardness of detecting macroscopic superpositions,” arXiv:2009.07450 (2020).

³⁰D. I. Blokhintsev, *The Philosophy of Quantum Mechanics* (Dordrecht: D. Reidel, 1968), pp. 43–44.

These difficulties are common to all attempts to formulate quantum mechanics in the language of quantities non-linear with respect to the wave function; throw the latter out of the door and it will come back through the window. This persistence of the wave function is an expression of the fact that the language of the linear theory is that of the essential nature of quantum mechanics; translation into any other tongue will always cause the poetry of this linearity to sound weaker than in the original.

And if one accepts the deviation and tries to make collapse an explicit postulate, one runs into the demand made unforgettable by John Bell in his paper “Against ‘measurement’”: say *when*.³¹ The textbook axioms trigger the collapse at “measurement,” and Bell’s charge is that this word is doing no real work, that it “anchors there the shifty split of the world into ‘system’ and ‘apparatus’,” a split that every author draws in a different place and none can define:

What exactly qualifies some physical systems to play the role of ‘measurer’? Was the wavefunction of the world waiting to jump for thousands of millions of years until a single-celled living creature appeared? Or did it have to wait a little longer, for some better qualified system . . . with a PhD?

A theory “fully formulated in mathematical terms, with nothing left to the discretion of the theoretical physicist” (Bell’s standard) must replace the word “measurement” with a physical condition: a scale, a rate, a trigger. It can be done. The Ghirardi–Rimini–Weber theory does it with spontaneous localization at a fixed tiny rate;³² the Diósi–Penrose model does it with a gravitational threshold.³³ These are serious, careful, realist proposals, and they are *different theories*, which deviate from quantum mechanics in specified regimes and are being hunted by experiment in exactly those regimes. Roger Penrose, to his credit, says so plainly: he believes quantum mechanics is incorrect and expects it to fail at the gravitational threshold.

That may even be true. But it is a prediction about future physics, not an interpretation of the physics we have. Our approach is *interpretive*. It does not propose a new theory; it proposes a way of understanding quantum mechanics exactly as it stands, and it is therefore reluctant to adopt any model that fails to reproduce the statistical predictions of quantum theory under all circumstances. Invariant collapse models fail that test by design. Whatever their merits as speculative physics, they are not available as readings of the theory we set out to read. Hence the second requirement, now earned in full: *if a physical collapse mechanism is what creates the values, the resulting theory inherently deviates from quantum mechanics*, detectably in principle, by Bob’s Hadamards, by Deutsch’s reversal, and by GRW’s spontaneous flashes, *and so collapse cannot be part of any interpretation of quantum mechanics as it is*.

³¹J. S. Bell, “Against ‘measurement’”, *Physics World* 3, no. 8 (August 1990), 33–40.

³²G. C. Ghirardi, A. Rimini, and T. Weber, “Unified dynamics for microscopic and macroscopic systems,” *Physical Review D* 34 (1986), 470–491.

³³L. Diósi, “Models for universal reduction of macroscopic quantum fluctuations,” *Physical Review A* 40 (1989), 1165–1174; R. Penrose, “On gravity’s role in quantum state reduction,” *General Relativity and Gravitation* 28 (1996), 581–600.

9.6 What remains

John Bell compressed the measurement problem into a single sentence: “Either the wavefunction, as given by the Schrödinger equation, is not everything, or it is not right.”³⁴ It offers two escapes: add something to the ontology beyond ψ , or change the dynamics of ψ . Maudlin’s paper is, in large part, a warning that the two are not symmetric: changing the dynamics alone accomplishes nothing, because a collapsed wavefunction is exactly as configuration-free, as unmanifest, and as closed on itself in the way Section 9.1 described, as an uncollapsed one. Our three requirements sharpen the asymmetry into a verdict.

They also settle the question Section 9.2 left open. That section established that there must be systems holding definite configurations and said nothing about when. There are three ways to answer, and two of them are now closed. The configurations might never exist, which Section 9.2 ruled out, since then ψ indexes nothing and the theory has nothing from which a value for any observable could be computed, leaving it unable to predict anything at all. They might arise at observation and vary with the observer, which Section 9.4 ruled out, since other minds then hold no settled contents and realism about them goes. Or they might arise at observation and be invariant for everyone, which takes a physical mechanism to bring them into being, and Section 9.5 ruled that out as a different theory rather than a reading of this one, answerable to experiments quantum mechanics does not predict. One answer is left standing, and it is the one the first requirement states: the configurations are there before anybody observes them.

Many Worlds refuses both of Bell’s escapes (ψ is everything *and* the linear dynamics is right) and accordingly fails the first requirement: no system ever holds a configuration, nothing points, and observation itself goes unexplained. And by Bell’s own accounting it forfeits the third as well, along the temporal axis rather than the interpersonal one: with no trajectories, every memory is a record of a past that never occurred, and the theory collapses into the “radical solipsism” of a present moment furnished with consistent lies. The Copenhagen interpretation, read with a variant collapse, fails the third: other minds lose their invariant existence, and realism gives way, however unintentionally, to solipsism. Read instead with an invariant, physical collapse, it takes the “not right” escape and fails the second, and so do GRW and Diósi–Penrose, openly and honorably: they are alternative theories awaiting experimental verdicts, not interpretations of quantum mechanics.

One reading satisfies all three requirements at once, and it takes Bell’s “not everything” escape in the least exotic way imaginable: what is added to ψ is nothing but the configuration itself: Maudlin’s manifest variables, the bits we have tracked since page one. The bits possess definite, invariant values at all times; the statistics are our knowledge of those values, not the values themselves; measurement is discovery, its “collapse” the Bayesian update of Chapter 3; and no mechanism creates the values because nothing needs to; they were never absent. This is the statistical interpretation, the reading Einstein and Blokhintsev defended in the century’s great argument, and the one we have been assembling from the start. Part II supplied what the century’s argument lacked: the explicit mathematics (base cases, history-dependent laws, transition matrices at every interval) showing that the reading is not wishful but *constructible*, down to a working simulation.

³⁴J. S. Bell, “Are there quantum jumps?” in C. W. Kilmister (ed.), *Schrödinger: Centenary Celebration of a Polymath* (Cambridge: Cambridge University Press, 1987), pp. 41–52; reprinted in *Speakable and Unsayable in Quantum Mechanics*.

What remains against it is a family of celebrated theorems (Kochen–Specker, Bell, Frauchiger–Renner, Pusey–Barrett–Rudolph) widely advertised as proving that no such reading can survive. Each advertisement, we will find, oversells. The theorems are correct, and what they actually constrain is not the existence of definite values but the *shape of the laws* that govern them, which must be non-Markovian and non-local, precisely as Part II already concluded for independent reasons. The remaining chapters take the theorems one at a time, translate each into circuits, and compute what it does and does not forbid.

Code: this chapter in Octave

```
% --- the two laboratory descriptions of Section 9.5 ---
H = (1/sqrt(2))*[1 1; 1 -1]; HH = kron(H,H);
runc = [0.5 0 0 0.5; 0 0 0 0; 0 0 0 0; 0.5 0 0 0.5]; % entangled lab
rcol = [0.5 0 0 0; 0 0 0 0; 0 0 0 0; 0 0 0 0.5]; % collapsed lab

diag(HH * runc * HH') % -> [0.5; 0; 0; 0.5] correlated
diag(HH * rcol * HH') % -> [0.25; 0.25; 0.25; 0.25] uniform

% Same diagonals before the perturbation: no direct readout can tell
diag(runc), diag(rcol)

% --- Section 9.2: one state, two branch readings ---
e0 = [1;0]; e1 = [0;1];
bell = (kron(e0,e0) + kron(e1,e1))/sqrt(2);
alt = (kron(H*e0,H*e0) + kron(H*e1,H*e1))/sqrt(2);
norm(bell - alt) % -> 0: the same vector, split two different ways

% every rotation of both bits splits it in two just as cleanly
for th = linspace(0, pi/2, 5)
    R = [cos(th) -sin(th); sin(th) cos(th)];
    c = kron(R,R)' * bell;
    abs(c(2)) + abs(c(3)) % -> 0: no cross terms, at every angle
end
```

Chapter 10

Quantum Contextuality

Chapter 9 ended with a promise: to take the celebrated no-go theorems one at a time, translate each into circuits, and compute what it does and does not forbid. We take the first: the claim that definite values are not merely unfashionable but *mathematically impossible*, that no assignment of values to a quantum system’s observables can be consistent with the theory’s predictions. The claim is usually credited to the Kochen–Specker theorem, whose proof is a combinatorial construction that takes real effort to follow. We will instead reach the same conclusion through an experiment where the algebra fits on half a page: the Greenberger–Horne–Zeilinger experiment.¹ The GHZ experiment is usually taught as a demonstration of non-locality. It is that too, and Chapter 11 will spell out the non-local implications in detail; but the cleanest lesson it teaches is a different one, called *contextuality*, and that lesson is our subject.

The conclusion, stated in advance so the reader knows what is and is not at stake: quantum mechanics really does forbid assigning definite values to *all* of a system’s observables simultaneously. What it does not forbid, and what we have been doing since Chapter 1, is assigning definite values to one basis of observables at a time. The theorem constrains how many value-assignments a realist may carry at once, and the statistical interpretation, it will turn out, never needed more than one.

The experiment is performed on electron spin, and the apparatus is the Stern–Gerlach device of Section 3.4: the electron passes through and is deflected one way or the other relative to how the device is turned, giving one bit, with 0 assigned to an upward deflection and 1 to a downward one. Turning the device picks out which of the three axes the deflection is measured along, and an unadorned measurement is the Z basis as always. We write X_i , Y_i , Z_i for the bit value obtained by measuring particle i on the corresponding basis, so that a run is named by which basis each of the three particles was read on.

10.1 Four runs and a bookkeeping certainty

The GHZ experiment is run four times. In each run, three electrons are entangled together, using the same preparation every time, and then each is measured once, on a basis that

¹D. M. Greenberger, M. A. Horne, and A. Zeilinger, “Going beyond Bell’s theorem,” in M. Kafatos (ed.), *Bell’s Theorem, Quantum Theory and Conceptions of the Universe* (Dordrecht: Kluwer Academic, 1989), pp. 69–72.

depends on the run. In the first run, particle 1 is measured on the X basis and particles 2 and 3 on the Y basis. The second and third runs rotate the roles: the X measurement moves to particle 2, then to particle 3, with the other two measured on Y. In the fourth run, all three particles are measured on the X basis.

Consider the first three runs, and in each run take the three measured bits and XOR them together, using $\oplus$, addition modulo two, the exclusive-or of Chapter 1. Experiment finds that the result is *always* 1: in every single run, either all three bits read 1, or exactly one of them does. Three experimental facts, then:

$$X_1 \oplus Y_2 \oplus Y_3 = 1$$

$$Y_1 \oplus X_2 \oplus Y_3 = 1$$

$$Y_1 \oplus Y_2 \oplus X_3 = 1.$$

Now suppose, since this is the assumption on trial, that each particle simply *has* an X value and a Y value, and each measurement passively reads off the value that is there. Then the six quantities $X_1, X_2, X_3, Y_1, Y_2, Y_3$ are six definite bits, and the three equations above are three facts about them. XOR the three equations together. Any quantity XORed with itself cancels, because $b \oplus b = 0$, so we combine the left-hand sides and hunt for pairs:

$$(X_1 \oplus Y_2 \oplus Y_3) \oplus (Y_1 \oplus X_2 \oplus Y_3) \oplus (Y_1 \oplus Y_2 \oplus X_3) = 1 \oplus 1 \oplus 1.$$

Each Y appears exactly twice: Y_1 in the second and third equations, Y_2 in the first and third, Y_3 in the first and second. So all six Y terms cancel pairwise, and the $1 \oplus 1$ on the right cancels too:

$$X_1 \oplus X_2 \oplus X_3 = 1.$$

We have predicted the outcome of the fourth run without performing it. If the three particles have X values at all, the XOR of the three must be 1: either all three bits read 1, or exactly one does.

Run the fourth experiment, and the real world answers: the XOR is *always* 0. Either all three bits read 0, or exactly one does, which is the exact opposite of the prediction, in every single run:

$$X_1 \oplus X_2 \oplus X_3 = 0.$$

There is no noise to hide in and no statistics to argue about. The prediction was not that some average would come out one way; it was a parity constraint holding with certainty, derived from nothing but the assumption that the six values exist. The assumption is false. It is mathematically impossible to assign definite values to $X_1, X_2, X_3, Y_1, Y_2, Y_3$ simultaneously such that the first three equations give 1 and the fourth gives 0, because the first three algebraically force the fourth to be 1. This is precisely the kind of impossibility that the Kochen–Specker theorem establishes in general,² but the GHZ experiment needs no combinatorics: the reader has just verified it with XOR.

²S. Kochen and E. P. Specker, “The problem of hidden variables in quantum mechanics,” *Journal of Mathematics and Mechanics* 17 (1967), 59–87. The original proof is a coloring argument involving 117 distinct measurement directions; the GHZ experiment compresses the same moral into four lines of algebra, which is why we teach it this way. A well-known presentation of the GHZ argument in this style is N. D. Mermin, “Quantum mysteries revisited,” *American Journal of Physics* 58 (1990), 731–734.

10.2 The experiment as circuits

Following our standing rule, the experiment must now be translated into an explicit circuit and its statistics computed, both to confirm that stock quantum mechanics actually predicts these outcomes and because the translation itself will teach us what a measurement on a rotated basis *is*.

An analogy first, since the recasting is easier to see than to state. A globe sits on a table. Looking at it from where you stand, you see a particular country: call it A. Walk around the table by some angle and you see a different country, B. Now return to your starting position and, instead of walking, rotate the *globe* by the same angle in the opposite direction: you see country B again. Rotating the context of an observation and counter-rotating the object observed are empirically equivalent.

The quantum state is a vector of unit length (it must be, to conserve probability), and a unitary operator preserves that length, so the only thing a unitary can change is the vector's direction: every unitary is a rotation of the quantum state. Rotating the Stern–Gerlach apparatus is therefore mathematically equivalent to leaving the apparatus alone and counter-rotating the quantum state with the corresponding unitary. Concretely, there exists a unitary that rotates the Z observable into the X observable, and another that rotates Z into Y, and both are gates the reader already knows:

$$U_{Z \rightarrow X} Z U_{Z \rightarrow X}^\dagger = X \quad U_{Z \rightarrow Y} Z U_{Z \rightarrow Y}^\dagger = Y$$

$$U_{Z \rightarrow X} = H \quad U_{Z \rightarrow Y} = SH.$$

(Both identities can be checked by direct multiplication with the gate definitions of Chapter 3.) To recast a rotated measurement as a standard one, we rotate the state in the opposite direction, and since the opposite of a unitary rotation is its Hermitian transpose:

$$U_{Z \rightarrow X}^\dagger = H^\dagger = H \quad U_{Z \rightarrow Y}^\dagger = (SH)^\dagger = HS^\dagger.$$

A measurement on the X basis is the gate H followed by a measurement on the Z basis. A measurement on the Y basis is the gate $S^\dagger$, then H, then a measurement on the Z basis.

The entangled preparation shared by all four runs is the *GHZ state*: a Hadamard on the first bit, then a CX recording it onto the second, then a CX recording the second onto the third, producing

$$\psi_{\text{GHZ}} = \frac{1}{\sqrt{2}} (1 \ 0 \ 0 \ 0 \ 0 \ 0 \ 0 \ 1)^\top :$$

three bits that read all-0 or all-1, half and half: the three-bit analogue of the Bell pair. The four runs are then the four circuits below, one per measurement pattern, with each X measurement recast as H-then-measure and each Y measurement recast as $S^\dagger$ -H-then-measure:

$X_1 Y_2 Y_3$								$Y_1 X_2 Y_3$							
$\mathbf{e}_0$	H	$\overline{\text{CX}}$		H	$\perp$	$\mathbf{e}_0$	H	$\overline{\text{CX}}$		$S^\dagger$	H	$\perp$			
$\mathbf{e}_0$		$\overline{\text{CX}}$	$\overline{\text{CX}}$	$S^\dagger$	H	$\perp$	$\mathbf{e}_0$		$\overline{\text{CX}}$	$\overline{\text{CX}}$	H	$\perp$			
$\mathbf{e}_0$			$\overline{\text{CX}}$	$S^\dagger$	H	$\perp$	$\mathbf{e}_0$			$\overline{\text{CX}}$	$S^\dagger$	H	$\perp$		
$Y_1 Y_2 X_3$								$X_1 X_2 X_3$							
$\mathbf{e}_0$	H	$\overline{\text{CX}}$		$S^\dagger$	H	$\perp$	$\mathbf{e}_0$	H	$\overline{\text{CX}}$		H	$\perp$			
$\mathbf{e}_0$		$\overline{\text{CX}}$	$\overline{\text{CX}}$	$S^\dagger$	H	$\perp$	$\mathbf{e}_0$		$\overline{\text{CX}}$	$\overline{\text{CX}}$	H	$\perp$			
$\mathbf{e}_0$			$\overline{\text{CX}}$		H	$\perp$	$\mathbf{e}_0$			$\overline{\text{CX}}$	H	$\perp$			

Computing the four output distributions (the chapter’s code listing does it in a few lines) gives, for each of the first three circuits,

$$\mathbf{p}_{X_1Y_2Y_3} = \mathbf{p}_{Y_1X_2Y_3} = \mathbf{p}_{Y_1Y_2X_3} = (0 \quad 0.25 \quad 0.25 \quad 0 \quad 0.25 \quad 0 \quad 0 \quad 0.25)^\top,$$

supported on the configurations 001, 010, 100, 111 (exactly the four whose XOR is 1), each with probability 0.25. And for the fourth,

$$\mathbf{p}_{X_1X_2X_3} = (0.25 \quad 0 \quad 0 \quad 0.25 \quad 0 \quad 0.25 \quad 0.25 \quad 0)^\top,$$

supported on 000, 011, 101, 110, the four whose XOR is 0. Stock quantum mechanics predicts the parity constraints exactly as the experiment finds them: XOR of 1 with certainty in the first three runs, XOR of 0 with certainty in the fourth. Note what the individual outcomes look like, because it will matter below: within each run, every *individual* bit is uniform, like a 50/50 coin, and only the parity of the trio is fixed.

10.3 Contextuality

What exactly has been refuted? Not definite values as such. The refuted assumption was stated before the algebra: that each particle has an X value *and* a Y value at the same time, each measurement passively revealing its own. The technical name for what fails is *non-contextuality*: the doctrine that a measured value does not depend on the context of its measurement, so that X_1 may be treated as one and the same bit whether particle 1 is measured alongside Y_2Y_3 or alongside X_2X_3 . It was precisely this treatment that licensed XORing the four equations together as simultaneous facts about six quantities. The algebra then convicts the doctrine: values for *all* the bases at once, waiting context-independently to be read off, cannot exist. A measurement is not always a passive act of looking.

If the four lines cannot all hold at once, they belong to different experiments carried out under different conditions. Yet the bits were prepared identically in every one of them, and the only difference between the runs is which axes the experimenters chose to measure them upon. The choice of measurement therefore plays a part in fixing the values found on those axes. And that choice is made after the preparation is complete, which leaves no way for the measured value to be a straightforward reading of a property the particle already carried. The value arises, at least in part, in the act of measuring, and thus depends upon the configuration of the measuring device. That configuration is what *context* means here: the value is produced under one particular measuring arrangement, and means nothing apart from it. This is the heart of contextuality.

In 1932 von Neumann published a proof that hidden variables are impossible, and it stood for three decades. Its essential assumption was that any real linear combination of observables is again an observable, and that expectation values combine the same way, demanded not only of the quantum averages but of each hypothetical underlying state. Bell took it apart in 1966.³ Measuring Z takes a magnet at one orientation, X a second, and $Z + X$ a third, and the values do not combine even when the averages do: the first two return ± 1 while the third returns $\pm\sqrt{2}$. Bell’s verdict on what the proof had actually established:

³J. S. Bell, “On the Problem of Hidden Variables in Quantum Mechanics,” *Reviews of Modern Physics* 38 (1966), 447–452. Von Neumann’s proof is in *Mathematische Grundlagen der Quantenmechanik* (Berlin: Springer, 1932).

It was not the objective measurable predictions of quantum mechanics which ruled out hidden variables. It was the arbitrary assumption of a particular (and impossible) relation between the results of incompatible measurements, either of which might be made on a given occasion but only one of which can in fact be made.

Von Neumann assumed non-contextuality. A value belongs to an observable whatever arrangement is used to find it, so the value of $Z + X$ is settled by the values of Z and X even though measuring it takes a third orientation. The GHZ argument reaches a contradiction from that same assumption, so the proof that rested on it does not stand.

Contextuality does not make the observer special. The context is a physical arrangement and not a state of anyone's knowledge: a magnet turned to one angle rather than another, a gate applied rather than omitted. Whatever fixes the value does so by interacting with the particle, and an apparatus left running in an empty room, with no conscious observer present at all, fixes it the same way. What the GHZ argument establishes is a fact about measuring devices, not about minds.

The reader should notice what the conviction does *not* cover. Nothing here has ever assigned a particle an X value and a Y value and a Z value simultaneously. The bits of Chapters 1 through 8 carry one value each: a configuration in a single designated basis, the computational basis. The statistics, histories, and transition matrices were all statements about *those* values. The GHZ experiment convicts a position we never held.

It does, however, force us to say something we have not yet had to say out loud: what happens when someone measures on a basis *other* than the designated one, as the GHZ experimenters do in every run. The circuits above already contain the answer, and the next section draws it out.

10.4 Observables and emergeables

Look again at any of the four circuits. Each “ X measurement” is, in circuit terms, a perturbation (the unitary H) followed by an ordinary measurement in the computational basis. Each “ Y measurement” is a perturbation (the unitary $HS^\dagger$) followed by the same. Nothing else in the formalism distinguishes a rotated-basis measurement from a plain one. A measurement on a basis orthogonal to the designated basis is not a passive reading; it is a physical disturbance of the system followed by a passive reading. The value it returns did not exist before the procedure. It was brought into being by the perturbation, exactly as the output of any stochastic gate in Chapter 2 is brought into being by the gate.

Barandes has recommended a vocabulary for this distinction. The quantities whose values really are there (the configuration in the designated basis) he calls the system's *beables*, borrowing Bell's word for what a system can physically *be*. Quantities associated with other bases he calls *emergeables*: observables that are emergent patterns in the dynamics of measurement processes and, in his words, “do not have a transparent interpretation at the level of the underlying configuration space.”⁴ Measuring a beable reveals what was

⁴Barandes, “The Stochastic-Quantum Correspondence,” §4. The term *beable* is Bell's: see J. S. Bell, “Subject and object,” in *The Physicist's Conception of Nature* (Dordrecht: Reidel, 1973), reprinted in *Speakable and Unspeakable in Quantum Mechanics*.

there. Measuring an emergeable manufactures a value through a physical process: lawfully, reproducibly, through a perturbing interaction whose form is fixed by the orientation of the measuring device relative to the designated basis.

This asymmetry between bases is interpretative, not mathematical. The mathematics of quantum theory treats every basis symmetrically; it is the realist reading that must single out one basis as the one whose values constitute the system's real configuration. Call that designated basis the *ontic basis*, the basis associated with what objectively exists. The physicist Gerard 't Hooft has likewise argued that a realist model of quantum mechanics requires such a choice (his term is the “ontological basis”),⁵ and the physicist David Bohm built the best-known realist model by choosing the position basis.⁶ Since our purposes are quantum information, where the systems are bits, the natural choice is the computational basis, which Section 3.4 identified with the Pauli-Z basis, and that is the choice every chapter so far has made.

That reduction has a limit. Choosing an ontic basis makes every other quantum variable an emergeable, so in Bohm's model momentum, spin and polarization all come down to different ways of measuring position. The limit is that a particle's mass and charge are not among them. Contextuality is a claim about quantities that can stand in superposition, and charge cannot: no preparation puts an electron in a superposition of charges, and no rotation of a measuring device turns charge into something else. These properties are therefore always beables: intrinsic physical properties of the system, whichever ontic basis is chosen.

Rovelli draws the same boundary around his relational reading. Of all a system's physical properties, he applies that reading only to “those that are variables: that is, those described by functions on the phase space, not the invariant properties such as the nonrelativistic mass of a particle.”⁷

10.5 Measurement as a function of what is really there

Realism does not require that every measurement passively reveal a pre-existing value. It requires something weaker and more precise: that what is observed be a lawful physical function of what is really there. We can state the criterion formally. Let X be the set of possible real states of the system, Y the set of possible measurement results, and $\Delta(Y)$ the set of probability distributions over Y . Then:

$$\forall y \in Y \exists f : X \rightarrow (Y \cup \Delta(Y)) \exists x \in X : (f(x) = y) \vee (f(x) \in \Delta(Y) \wedge f(x)(y) > 0).$$

In words: for any result y that a measurement can return, there must be a physical description of the measurement process, a function f of the system's real state x , that either produces y outright or produces a probability distribution in which y has non-zero weight. If the measurement is a passive reading, f is as simple as $f(x) = x$: what is measured is what is there. But the criterion deliberately allows more: some properties may be emergent, observable

⁵G. 't Hooft, *The Cellular Automaton Interpretation of Quantum Mechanics* (Cham: Springer, 2016).

⁶D. Bohm, “A Suggested Interpretation of the Quantum Theory in Terms of ‘Hidden’ Variables,” I and II, *Physical Review* 85 (1952), 166–179 and 180–193.

⁷C. Rovelli, *Helgoland: Making Sense of the Quantum Revolution*, trans. E. Segre and S. Carnell (New York: Riverhead Books, 2021), note 54.

only through a physical process that returns a *derivative* of the real state, deterministically if f is deterministic and statistically if f is a stochastic perturbation.

Written this way, the criterion states exactly what Section 9.2 argued in plain terms. Every result a measurement can return is $f(x)$ for some real state x , so a theory that returns results owes an x to evaluate f at. The demand falls on x and not on y . Nothing here asks that each observable carry a value in advance, which is what Section 10.1 showed cannot be had; the emergeables are values of f , and a function needs no second function to supply its argument. What a theory cannot do is offer the f and withhold the x , which is what a distribution over nothing comes to, and what taking the probabilities as primitive comes to as well.

Chapter 9's third requirement is the further demand that x be the same for everyone. Were x to depend on who was asking, two observers measuring one system would evaluate their functions at different arguments, and no comparison of the results would mean anything. Section 9.4 gives the reason that matters, which concerns what becomes of other people's minds when their configurations are allowed to vary with the describer.

The criterion sorts the interpretations of Chapter 9 efficiently. One might offer collapse as satisfying it: the real state is the quantum state, and $f(x) = |x|^2$ hands each result its Born weight. But for the criterion to be satisfied, f must describe a *physical process*; a non-physical f explains nothing about how the real state translates into the observed one, and simply re-poses the measurement problem. A collapse taken to be physical does satisfy the criterion, though by Chapter 9's second requirement it is then a different theory. Many Worlds posits no f at all, physical or otherwise.

For us, f is the machinery already built in Chapter 6. A measurement at one time interval is a function of the system's real configuration at the immediately preceding one. It depends not on the statistical distribution but on the definite values, because a stochastic matrix is a statistical truth table whose columns are indexed by definite input configurations. The single-step transition matrix for the measurement is computed by the same iterative proportional fitting as always, with one substitution: the unitary being fitted is the basis rotation $U^\dagger$ that recasts the measurement onto the ontic basis,

$$\Gamma^{[t+1 \leftarrow t]} = \lim_{k \rightarrow \infty} \text{diag}(\mathbf{u}^{(k)}) \left| U^\dagger \right|^2 \text{diag}(\mathbf{v}^{(k)}),$$

and the measurement function is simply that matrix applied to the real configuration:

$$f(x) = \Gamma^{[t+1 \leftarrow t]} \mathbf{e}_x.$$

If the measurement device is aligned with the ontic basis, the recasting unitary is the identity, the distribution it transports is unchanged, the fitted matrix collapses to the identity, and

$$f(x) = \mathbf{e}_x :$$

the measurement is passive and merely reveals what was there. If the device is misaligned, the matrix does not reduce to the identity, and the perturbation is a real one: the device genuinely coins something at random, though the coin is a lawful function of the real value, and the observed result always has non-zero weight in $f(x)$.

Both cases are on display in the GHZ circuits. For a single bit fresh from a base case with value 0, an X measurement is the perturbation H: the transition matrix is $|H|^2$, and

$$f(0) = |H|^2 \mathbf{e}_0 = \begin{pmatrix} 0.5 \\ 0.5 \end{pmatrix} :$$

the X value is manufactured on the spot, a fair coin, though a lawful function of the bit that was really there. For the entangled trio in the fourth run, the fitting runs from the GHZ distribution $(0.5, 0, 0, 0, 0, 0, 0, 0.5)$ to the all-X output distribution, with default matrix $|\mathbf{H} \otimes \mathbf{H} \otimes \mathbf{H}|^2$, and the resulting matrix answers the question the reader should now be asking: what do the definite values do in an individual run? In each run the trio's real configuration is either 000 or 111, and the fitted matrix gives

$$f(000) = f(111) = (0.25 \quad 0 \quad 0 \quad 0.25 \quad 0 \quad 0.25 \quad 0.25 \quad 0)^T :$$

whichever configuration the trio really occupies, the three X values are coined at the measurement, individually uniform yet jointly constrained, with every configuration of odd parity carrying zero weight. The parity certainty of the GHZ experiment is enforced by the *law*, not by pre-existing X values. That is exactly as it must be, since the algebra proved that pre-existing X and Y values cannot exist. The values that do exist at all times are the ontic-basis values; the X and Y outcomes are emergeables, manufactured by lawful stochastic perturbation of what really is. Contextuality, which arrives as a threat to realism, ends up as a description of the laws: the function f depends on which measurement is performed, that is, on the context, and nothing about definite values ever needed it not to.

One more observation, and it points at the next chapter. The three particles of the GHZ trio may be spatially separated when the measurements occur, and the parity constraint binds their three coined outcomes *jointly* no matter the separation. A law that coins values under a joint constraint across distance is a non-local law. That is not news to us, since Chapter 8's holism said as much, but the GHZ experiment sharpens it, and the Frauchiger–Renner argument of Chapter 11 will sharpen it further, into a statement about what simultaneity itself can mean.

10.6 The underdetermination problem

One question is still unsettled. The ontic basis had to be *chosen*. The mathematics did not choose it for us: in Werner Heisenberg's original matrix formulation of quantum mechanics, all bases enter symmetrically, and the theory's predictive structure supplies no empirical reason to crown one basis over another. Rotate every state and every operator by the same unitary and the statistics are untouched.

We must be precise about what kind of problem this is. It is not that realism is *ruled out*: the GHZ experiment ruled out all-bases-at-once realism, and one-basis realism passed untouched, constructively, with transition matrices. The problem is that one-basis realism is *underdetermined*: there are too many viable versions, one per choice of ontic basis, and quantum theory alone cannot adjudicate among them. A theory being ruled out and a theory being underdetermined are opposite predicaments. The first says no candidate survives; the second says too many do.

Several responses are available, and we regard each of them as legitimate.

The first is conventionalism. Choose whatever ontic basis is convenient for the experimental context at hand (for a quantum computer, its computational basis), while acknowledging the choice as a convention: a working guess about the underlying ontology, revisable in another context.

What conventionalism does not license is holding two such choices at once. Each is internally consistent, and no experiment distinguishes them, since the formalism is basis-symmetric. But two laboratories that designate different ontic bases for the same system are not offering compatible descriptions to be set side by side. They are making rival claims about what its real configuration is, and they cannot both be right: the assumption refuted in Section 10.1 was exactly this one, that a system carries definite values in two bases at the same time. A realist who wants to cover both laboratories in a single account must therefore adopt one of the two conventions throughout, or find a third basis and describe both in that. The convention is revisable, but it is held one at a time.

Whether any ontology has to cover every domain is itself in dispute. Pris treats a physical theory as a norm for identifying reality within the domain where it is applied, and holds that classical mechanics is true in its own domain without its objects reducing to quantum ones; reduction is something to be claimed locally, context by context, and not in general.⁸ The absence of one ontology spanning everything is then no scandal, since no ontology has been used or tested outside the domains it was built for.

Our situation is not quite his, since the rival ontic bases here compete within a single domain rather than belonging to separate ones. The moral survives the difference. A basis chosen to suit the experiment at hand, which the experiment then answers to, has captured what there was to capture, and describing two experiments together is a further context with a choice of its own. A single basis could be imposed on every context by fiat, but nothing would pick out that basis rather than another, and the imposition could make later work harder, since a basis that suits one experiment may obscure the next.

The second is to argue for a particular basis on grounds independent of empirical evidence, and here the strongest candidate is *position*. If any property of a particle deserves to be called fundamental, it is position, because a particle is a geometric point in three-dimensional space, and the defining characteristic of a point is *where it is*. Position is also continuous, and the continuum yields something remarkable: the polar decomposition of the quantum state on the position basis can be rearranged into the classical Hamilton–Jacobi equation plus a single non-linear correction term (the “quantum potential”), and the Hamilton–Jacobi equation is *deterministic*.

The position basis thus offers to convert quantum mechanics from a stochastic theory into a deterministic one: this is exactly Bohm’s pilot-wave construction, point particles with well-defined positions at all times, moving on deterministic trajectories. The rearrangement depends on the continuum, so a basis whose values are discrete offers nothing comparable. The construction is mathematically equivalent to standard quantum mechanics: a boon, since that is what makes it a viable interpretation rather than a rival theory; and a bane, since it means no experiment can ever certify that the position basis was the right choice.

The alternatives also run into conceptual difficulties quickly. To be an object at all seems to involve being *somewhere*, and being somewhere *else* is how one object is told apart from another. Spatial separability is thus a good part of what makes objects count as distinct in the first place. A basis that demotes position to an emergeable therefore demotes the one feature by which the world is ordinarily carved into separate objects, and a theory whose values are not located anywhere becomes correspondingly hard to state in ordinary language: there is no longer any saying which thing the values belong to, or which of them are near each

⁸F.-I. Pris, “On Quantum Fundamentalism,” arXiv:2307.12992 (2023), §2.

other, or what is going on over *here* as opposed to over *there*.

Another possibility is to take the configuration of the fields as the beable rather than the positions of particles. A field configuration is specified at every point of space and holds a definite value at every moment, so the world it describes is still carved into here and there. Position becomes an emergeable: no particle has a definite position at all times, and where one is found emerges from the field configuration. Nikolić has shown that a model of this kind reproduces the measurable predictions of relativistic quantum field theory in a realist fashion without much difficulty.⁹

Taking the field as the beable also changes what counts as non-local. The barrier of Section 8.5 influences a photon that never touched it, which is action at a distance only if the photon was on one path. A field is on both paths and meets the barrier locally, so on that reading nothing acts at a distance at all. Catani, Leifer, Schmid and Spekkens have shown as much with a worked classical field theory.¹⁰ Bell's inequalities are another matter. They constrain the statistics rather than any picture of what carries them, so no choice of beable relieves the violation, and a field model inherits the non-locality with everything else. The one apparatus holds both kinds. With a barrier in place, the photon's path and the barrier's condition are a maximally entangled pair between the two splitters, and measuring each of them on two bases would in principle break the bound of Section 11.6; the interaction-free measurement reads one basis and breaks nothing. So the field description is useful whether or not it is fundamental, since it separates a non-locality that depends upon which basis was chosen as ontic from one that survives every choice.

A third possibility stands between conventionalism and an argued-for basis. A convention tight enough to force the ontic basis from the structure of the experiment itself would leave less to taste than picking whichever basis suits the work at hand, since some physical property of the experiment would settle it. Taken one at a time, such a convention might assign incompatible ontic bases to two experiments built differently. Brought together under one description, the two choices would merge into a single basis covering both, and a description of the whole universe would force an ontic basis for the whole universe. It stays conventional even so, since a universal ontic basis is thereby implied to exist and any smaller description may be working in the wrong one. 't Hooft has suggested such an approach.¹¹ His criterion appeals to how a quantum system is physically realized, which an abstract description in terms of quantum information does not record, so it is not available to us here.

Here is where the realist stands at the end. Giving every bit a definite value on one basis can be done, and Chapter 6 does it. Giving every bit a definite value on all bases at once cannot be done, and four lines of XOR show why. Which basis to pick is left open, with a conventionalist answer, an argued-for answer, and a convention fixed by physical structure all on the table, and no experiment able to decide among them.

We adopt the conventionalist posture with the computational basis, and will flag the places where the choice matters. Underdetermination recurs throughout the realist reading:

⁹H. Nikolić, "Relativistic QFT from a Bohmian perspective: A proof of concept," arXiv:2205.05986 (2022), §3.2.

¹⁰L. Catani, M. Leifer, D. Schmid and R. W. Spekkens, "Why interference phenomena do not capture the essence of quantum theory," *Quantum* 7 (2023), 1119, §4.1; they grant at §5.1.3 that Bell's theorem lies beyond the model.

¹¹G. 't Hooft, "Explicit construction of Local Hidden Variables for any quantum theory up to any desired accuracy," arXiv:2103.04335 (2021), §5.

it first appeared in Chapter 6 as the choice among transport-compatible transition matrices, it returns here as the choice of ontic basis, and it will return in Chapter 11 as the choice of a preferred slicing. In every case the lesson is the same: the theorems leave the realist reading standing, and show where the choices are.

Code: this chapter in Octave

```

% --- the gates, and the GHZ preparation ---
I2 = eye(2); H = (1/sqrt(2))*[1 1; 1 -1]; Sd = [1 0; 0 -1i];
CX = [1 0 0 0; 0 1 0 0; 0 0 0 1; 0 0 1 0];
e0 = [1;0]; psi0 = kron(e0, kron(e0, e0));

% GHZ preparation: H on bit 1, CX(1->2), CX(2->3)
ghz = kron(I2,CX) * kron(CX,I2) * kron(H,kron(I2,I2)) * psi0;
ghz' % -> (1/sqrt(2)) [1 0 0 0 0 0 0 1]

% Rotated measurements recast onto the ontic (Z) basis:
% X basis -> apply H, Y basis -> apply S' then H (matrix H*Sd)
X = H; Y = H*Sd;
pXYX = abs( kron(X, kron(Y,Y)) * ghz ).^2 % [0 .25 .25 0 .25 0 0 .25]
pYXY = abs( kron(Y, kron(X,Y)) * ghz ).^2 % same
pYYX = abs( kron(Y, kron(Y,X)) * ghz ).^2 % same
pXXX = abs( kron(X, kron(X,X)) * ghz ).^2 % [.25 0 0 .25 0 .25 .25 0]

% Parity certainties: XOR = 1 for the first three runs, 0 for the fourth
par = mod(sum(dec2bin(0:7,3) - '0', 2), 2);
[sum(pXYX(par==1)), sum(pYXY(par==1)), sum(pYYX(par==1)), ...
 sum(pXXX(par==0))] % -> [1 1 1 1]

% --- IPF: the single-step transition matrix for the all-X run ---
A = abs(kron(H,kron(H,H))).^2; % default matrix |H(x)H(x)H|^2
pt = abs(ghz).^2; pt1 = pXXX;
v = ones(8,1);
for k = 1:5000
    u = pt1 ./ (A*(v.*pt)); u(~isfinite(u)) = 0;
    v = 1 ./ (A'*u); v(~isfinite(v)) = 0;
end
G = diag(u)*A*diag(v);
G * pt % transports correctly -> pXXX
G(:,1)' % f(000): [.25 0 0 .25 0 .25 .25 0] - even-parity support only
G(:,8)' % f(111): identical - the law, not the values, fixes the parity

% Aligned measurement: recasting unitary is the identity
% -> fitted matrix is the identity -> f(x) = e_x (passive revelation)

```

Chapter 11

Quantum Non-Locality

Chapter 10 left something unfinished. The GHZ trio’s measurement outcomes are coined at the moment of measurement: individually random, yet jointly constrained across whatever distance separates the three particles. A law that assigns values under a joint constraint at a distance is a non-local law. We ask what follows from that non-locality: how it resolves a celebrated paradox, what it means for the standard reading of special relativity, and why it can never be used to send a signal, despite everything.

The vehicle is a thought experiment published by the physicists Daniela Frauchiger and Renato Renner under the deliberately provocative title “Quantum theory cannot consistently describe the use of itself.”¹ Their protocol involves four entangled systems: two “friends” making measurements inside sealed labs, and two “Wigners” measuring the labs from outside. It builds a chain of inferences that ends, with probability 1/12 per run, in two agents holding contradictory conclusions about the same past. The four-system bookkeeping is heavy, but the paradox has a two-bit core, and we work with the core. Everything essential survives the distillation (including, as we will see, the probability 1/12). Section 11.6 then does the same for Bell’s 1964 theorem, which reaches the chapter’s conclusion by the older route of ensembles and an inequality.

11.1 A two-bit distillation

Prepare two bits, giving the first to Alice and the second to Bob, with the following circuit. The rotation angle is chosen so that $\cos(\theta/2) = 1/\sqrt{3}$ (that is, $\theta = 2 \arccos(1/\sqrt{3}) \approx 1.911$), and CH is the *controlled Hadamard*, which works exactly like CX except that the gate it conditionally applies to the target is H rather than X:

$$\begin{array}{cc} \mathbf{e}_0 & R_Y(\theta) \quad \overline{\text{CX}} \quad \overline{\text{CH}} \\ \mathbf{e}_0 & \text{CX} \quad \text{CH} \end{array} \quad \text{CH} = \begin{pmatrix} 1 & 0 & 0 & 0 \\ 0 & 1 & 0 & 0 \\ 0 & 0 & \frac{1}{\sqrt{2}} & \frac{1}{\sqrt{2}} \\ 0 & 0 & \frac{1}{\sqrt{2}} & -\frac{1}{\sqrt{2}} \end{pmatrix}.$$

The rotation puts the first bit’s amplitudes at $(1/\sqrt{3}, \sqrt{2/3})$; the CX records the first bit onto the second; the CH then, conditionally on the first bit being 1, mixes the second. The

¹D. Frauchiger and R. Renner, “Quantum theory cannot consistently describe the use of itself,” *Nature Communications* 9, 3711 (2018).

resulting quantum state and its Born distribution are

$$\psi_0 = \frac{1}{\sqrt{3}} \begin{pmatrix} 1 \\ 0 \\ 1 \\ -1 \end{pmatrix} \quad \mathbf{p}_0 = \begin{pmatrix} 1/3 \\ 0 \\ 1/3 \\ 1/3 \end{pmatrix}.$$

We will write A for the outcome of a plain, ontic-basis measurement of Alice's bit, and B for the same on Bob's. Reading the Born distribution: the outcomes $AB = 00, 10$, and 11 each occur with probability $1/3$, and

$$AB = 01 \text{ is impossible.}$$

That single impossible outcome, $AB = 01$, is the engine of everything that follows.

Now recall Chapter 10's lesson about rotated measurements: applying H to a bit and then measuring is a measurement of an emergeable, which is a lawful stochastic perturbation followed by an ontic-basis reading. We consider three variations on the experiment, distinguished by who perturbs before reading.

Case A: Alice perturbs, Bob reads plainly. Apply H to Alice's bit before measurement. The value Alice now reads is not A , because her perturbation manufactured it; call it α . The state and distribution:

$$(H \otimes I) \psi_0 = \frac{1}{\sqrt{6}} \begin{pmatrix} 2 \\ -1 \\ 0 \\ 1 \end{pmatrix} \quad \mathbf{p}_A = \begin{pmatrix} 2/3 \\ 1/6 \\ 0 \\ 1/6 \end{pmatrix},$$

so $\alpha B = 00, 01, 11$ are possible, and $\alpha B = 10$ is impossible.

Case B: Bob perturbs, Alice reads plainly. Apply H to Bob's bit instead; his manufactured value is β :

$$(I \otimes H) \psi_0 = \frac{1}{\sqrt{6}} \begin{pmatrix} 1 \\ 1 \\ 0 \\ 2 \end{pmatrix} \quad \mathbf{p}_B = \begin{pmatrix} 1/6 \\ 1/6 \\ 0 \\ 2/3 \end{pmatrix},$$

so $A\beta = 10$ is impossible.

Case C: both perturb. Apply H to both bits:

$$(H \otimes H) \psi_0 = \frac{1}{2\sqrt{3}} \begin{pmatrix} 1 \\ 3 \\ 1 \\ -1 \end{pmatrix} \quad \mathbf{p}_C = \begin{pmatrix} 1/12 \\ 3/4 \\ 1/12 \\ 1/12 \end{pmatrix},$$

and now *all four* outcomes are possible.²

²The four cases are a CHSH test: two parties, two settings each, four joint distributions. Score every reading as $+1$ for 0 and -1 for 1 , and the correlations for $\mathbf{p}_0$, $\mathbf{p}_A$, $\mathbf{p}_B$ and $\mathbf{p}_C$ come to $\frac{1}{3}$, $\frac{2}{3}$, $\frac{2}{3}$ and $-\frac{2}{3}$. Clauser, Horne, Shimony and Holt showed that the combination $\frac{1}{3} + \frac{2}{3} + \frac{2}{3} + \frac{2}{3}$ of those four reaches 2 at most on any local account, and here it is $\frac{7}{3}$. J. F. Clauser, M. A. Horne, A. Shimony and R. A. Holt, "Proposed Experiment to Test Local Hidden-Variable Theories," *Physical Review Letters* 23 (1969), 880–884. No single distribution over all four readings reproduces them, which is Section 11.3's conclusion arrived at from the other end.

11.2 The inferences, and the contradiction

The impossibilities license inferences, and the inferences are pure textbook, since each one is exactly the reasoning the Dirac–von Neumann axioms endorse.

Suppose Case A is run and Alice happens to read $\alpha = 1$. Since $\alpha B = 10$ is impossible, the only live outcome with $\alpha = 1$ is $\alpha B = 11$: Alice knows $B = 1$ without anyone touching Bob's bit. And since the original preparation makes $AB = 01$ impossible, $B = 1$ requires $A = 1$. Alice concludes: *before my perturbation, the pair was $AB = 11$.*

Suppose instead Case B is run and Bob happens to read $\beta = 0$. Since $A\beta = 10$ is impossible, $\beta = 0$ forces $A = 0$; and since $AB = 01$ is impossible, $A = 0$ forces $B = 0$. Bob concludes: *before my perturbation, the pair was $AB = 00$.*

Each inference, on its own, is sound. Now run Case C, where both perturb. With probability $1/12$ (the entry $\mathbf{p}_C$ assigns to the configuration 10, and the same $1/12$ as in Frauchiger and Renner's original protocol), the outcome is $\alpha\beta = 10$. Alice, holding $\alpha = 1$, runs her Case-A reasoning and concludes the pair *was* 11. Bob, holding $\beta = 0$, runs his Case-B reasoning and concludes the pair *was* 00. Both cannot be right.

One might hope to shrug the contradiction off: both parties perturbed the pair, the original values are gone, and no one can go back and check, so who cares? A pragmatist may not. A realist must. Chapter 9's first requirement is that the pair of bits held some definite configuration before the perturbations, and the stakes sharpen in the usual way: nothing prevents the two bits from being two bits *of two brains*. In Frauchiger and Renner's original telling the systems under measurement include the memories of two friends, so the contradictory inferences are contradictory claims about *what two observers experienced*. By the third requirement, there must be a fact of the matter. Something in the reasoning has to give.

11.3 Why the inferences cannot be local

When Alice perturbs her bit, is the perturbation's effect a function of her bit alone?

If it were, the change from the preparation's distribution $\mathbf{p}_0$ to Case A's distribution $\mathbf{p}_A$ would be effected by some stochastic matrix acting on Alice's bit only: a 2×2 column-stochastic Γ with

$$(\Gamma \otimes \mathbf{I}) \mathbf{p}_0 = \mathbf{p}_A.$$

Let us write out the candidate entry by entry, $\Gamma = \begin{pmatrix} \gamma_{00} & \gamma_{01} \\ \gamma_{10} & \gamma_{11} \end{pmatrix}$. Nothing whatever is being assumed about these four numbers. We are supposing only that *some* local perturbation of Alice's bit does the job, and writing its entries as unknowns to see what they would have to be. Working out $(\Gamma \otimes \mathbf{I}) \mathbf{p}_0$ and setting it equal to $\mathbf{p}_A$ gives four equations in those four unknowns:

$$(\Gamma \otimes \mathbf{I}) \mathbf{p}_0 = \frac{1}{3} \begin{pmatrix} \gamma_{00} + \gamma_{01} \\ \gamma_{01} \\ \gamma_{10} + \gamma_{11} \\ \gamma_{11} \end{pmatrix} = \begin{pmatrix} 2/3 \\ 1/6 \\ 0 \\ 1/6 \end{pmatrix}.$$

They can be solved in order, each one pinning down a single unknown. The second component gives $\gamma_{01}/3 = 1/6$, so $\gamma_{01} = 1/2$. The fourth gives $\gamma_{11}/3 = 1/6$, so $\gamma_{11} = 1/2$. The first then reads $(\gamma_{00} + 1/2)/3 = 2/3$, so $\gamma_{00} = 3/2$. The third reads $(\gamma_{10} + 1/2)/3 = 0$, so $\gamma_{10} = -1/2$.

Four equations, four unknowns, full rank: no freedom is left anywhere. There is exactly one matrix of the assumed form, and it is

$$\Gamma = \begin{pmatrix} 3/2 & 1/2 \\ -1/2 & 1/2 \end{pmatrix}.$$

Substituting it back reproduces $\mathbf{p}_A$ exactly. The requirement determines this matrix and no other.

The number $\gamma_{00} = 3/2$ is the probability that Alice's bit, found at 0, is left at 0, and $\gamma_{10} = -1/2$ is the probability that the same bit is flipped to 1. A probability of $3/2$ and a probability of $-1/2$ are not probabilities: the first says the outcome happens more often than always, the second that it happens less often than never. Notice also which requirement is *not* what fails. Both columns sum to 1, automatically, as they had to once $\mathbf{p}_0$ and $\mathbf{p}_A$ were normalized; the candidate is not convicted of failing to conserve probability. It is convicted because two of its entries lie outside $[0, 1]$, and no assignment of chances to outcomes can look like that.

That is the whole argument, and note how little went into it. The bounds on a probability were never imposed on the solve: the four equations have full rank, so they settle all four entries on their own, and the violation is what they settle on rather than a conflict arranged in advance. We assumed only that *some* local law exists, let the four equations say what it would have to be, and they returned a single candidate that is not a law at all. *No local stochastic perturbation of Alice's bit alone can account for what the H gate does to the pair's statistics.* The perturbation's law consults Bob's bit as an input, exactly as Chapter 8's beam splitter consulted a barrier the photon never touched. Under the law of Section 6.6 it consults without disturbing, which is the *pseudo-locality* of Section 6.7: the default $|U \otimes I|^2$ is zero on every entry that would change Bob's value, and the fitting cannot put probability where the unitary put none, so the value of his bit is left where it was. That much is a property of this law rather than of quantum mechanics. Section 11.9 proves the companion statement, which holds whatever fills the interval: nothing Alice does moves Bob's statistics.

One escape covers every case in which Alice moves. A bit carrying a stored copy of Bob's value, made when the preparation had the bits interact, would answer the local description while reproducing $\mathbf{p}_A$ exactly: the law would read the copy, which sits locally, and would never reach across. That is the hidden memory of Section 5.3. The escape reaches past Alice's own step. When Bob perturbs second, his step has a local law of its own: found at 0 his bit flips to 1 for certain, found at 1 it takes a fair coin, and the matrix $\begin{pmatrix} 0 & 1/2 \\ 1 & 1/2 \end{pmatrix}$ carries $\mathbf{p}_A$ to $\mathbf{p}_C$ exactly. With the copy on Alice's side and this law on Bob's, the whole both-perturb run has a local account, trajectories and all.

What no local account covers is the four cases together, and an explanation of the experiment must cover all four with one law. Run the copy account on Case B, where Bob perturbs and Alice reads plainly. A law that knows only Bob's bit must treat every run at $B = 0$ alike, and the state demands opposite things of those runs: the runs at $A = 1, B = 0$ must all end at $\beta = 1$, since $A\beta = 10$ is impossible, while the runs at $A = 0, B = 0$ must split evenly between $A\beta = 00$ and 01 , each at $1/6$ of the ensemble. The law above, applied to Case B, sends the runs from 00 to 01 for certain and splits the runs from 11 , placing $1/6$ of the ensemble on the forbidden outcome. A copy stored on Bob's side covers Case B and fails Case A the same way. Each case on its own has a local account; the four together have none.

Pre-assignment is where Chapter 10's argument bites. To explain all four outcomes with one set of hidden memories is to give each run values for both of Alice's readings, A and α , and both of Bob's, B and β , each reading being what the experiment finds in two of the four cases. The two readings are measurements in different bases. The plain reading is a measurement on the Z basis. The perturbed reading is a measurement on the X basis, being the gate H followed by a measurement on the Z basis, as Section 10.2 recast it. Pre-assigning values to both is giving every bit a definite value on all bases at once, waiting context-independently to be read off. That is the doctrine Section 10.3 named *non-contextuality* and the GHZ experiment refuted: its four XOR lines are values assigned across bases that cannot all hold at once. The algebra here is the same conviction reached with different lines. The copy account fails Case B precisely because its pre-assigned values cannot be made consistent across the four cases, and the footnote's CHSH combination is the arithmetic that certifies the failure, as the four XOR lines are the arithmetic that certifies it there.

One thing did not go into the argument at all: the transition law. No step above used the construction of Section 6.6, or any other choice of law for the interval. The four equations were built from $\mathbf{p}_0$ and $\mathbf{p}_A$, and those are Born distributions, fixed by quantum mechanics before any question arises about what happens between them. The argument therefore runs unchanged for every member of the equivalence class of Section 6.7. Whatever law fills the interval, no local stochastic perturbation of Alice's bit reproduces $\mathbf{p}_A$, and non-locality here is not something a different choice of law could undo.

11.4 The time-ordering of events

Since the perturbations consult each other's bits, it matters *in which order they fire*: this is where the contradiction unravels. Chapter 6's iterative proportional fitting gives the unique single-step transition matrices for the whole experiment, so we can simply compute the two orderings and watch.

The preparation is common to both. From the base case $\mathbf{e}_{00}$, the three preparation gates contribute their fitted steps: $|\mathbf{R}_Y(\theta)|^2 \otimes \mathbf{I}$, the CX permutation, and the CH's fitted step, carrying the ensemble to $\mathbf{p}_0 = (1/3, 0, 1/3, 1/3)$, with each individual pair really in configuration 00, 10, or 11.

Version 1: Alice perturbs first. Fitting Alice's H (transport $\mathbf{p}_0 \rightarrow \mathbf{p}_A$, default $|\mathbf{H} \otimes \mathbf{I}|^2$) and then Bob's H (transport $\mathbf{p}_A \rightarrow \mathbf{p}_C$, default $|\mathbf{I} \otimes \mathbf{H}|^2$) gives

$$\Gamma_{\text{Alice}} = \begin{pmatrix} 1 & 0 & 1 & 0 \\ 0 & \frac{1}{2} & 0 & \frac{1}{2} \\ 0 & 0 & 0 & 0 \\ 0 & \frac{1}{2} & 0 & \frac{1}{2} \end{pmatrix} \quad \Gamma_{\text{Bob}} = \begin{pmatrix} \frac{1}{10} & \frac{1}{10} & 0 & 0 \\ \frac{9}{10} & \frac{9}{10} & 0 & 0 \\ 0 & 0 & \frac{1}{2} & \frac{1}{2} \\ 0 & 0 & \frac{1}{2} & \frac{1}{2} \end{pmatrix}.$$

Read the columns of Γ_{Alice} : a pair in configuration 10 goes to 00 *with certainty* (Alice's bit lands 0) while a pair in 11 sends her bit to a fair coin. Her perturbation's outcome depends on Bob's bit, the non-local dependence we just proved must exist. And read Γ_{Bob} : by the time Bob fires, his coin's bias depends on Alice's bit, which has already been perturbed.

Version 2: Bob perturbs first. The same fitting in the other order gives the mirror image:

$$\Gamma'_{\text{Bob}} = \begin{pmatrix} \frac{1}{2} & \frac{1}{2} & 0 & 0 \\ \frac{1}{2} & \frac{1}{2} & 0 & 0 \\ 0 & 0 & 0 & 0 \\ 0 & 0 & 1 & 1 \end{pmatrix} \quad \Gamma'_{\text{Alice}} = \begin{pmatrix} \frac{1}{2} & 0 & \frac{1}{2} & 0 \\ 0 & \frac{9}{10} & 0 & \frac{9}{10} \\ \frac{1}{2} & 0 & \frac{1}{2} & 0 \\ 0 & \frac{1}{10} & 0 & \frac{1}{10} \end{pmatrix}.$$

Both versions transport the ensemble to exactly the same final distribution $\mathbf{p}_C$, so the observable statistics cannot tell the orderings apart. But now follow the definite values through each chain and score the two inferences. Conditioning on the trajectories (the chapter’s code listing does the arithmetic):

	$P(AB = 11 \mid \alpha = 1)$	$P(AB = 00 \mid \beta = 0)$
Version 1 (Alice first)	1	1/5
Version 2 (Bob first)	1/5	1

Whoever perturbs *first* draws an inference that is correct with certainty. Whoever perturbs second draws an inference that is usually wrong. The mechanism is visible in the matrices: Alice’s inference from $\alpha = 1$ leans on her perturbation consulting Bob’s bit *while it still holds its original value B*. If she fires first, that is exactly what happens, and her conclusion $AB = 11$ is guaranteed. By firing, however, she has perturbed the very configuration that Bob’s perturbation will consult, so when Bob fires second, his coin consults not the original pair but the pair as Alice left it, and his textbook reasoning, which was derived for the case where Alice’s bit sat unperturbed, no longer applies to the situation he is actually in. (If the two perturbations are treated as one simultaneous joint step, the fitted matrix turns out to have identical columns. Conditioning on α or β then teaches nothing at all about the prior configuration, and neither inference is licensed.)

The contradiction is thereby resolved by noticing that the two inferences are never valid *in the same run*, and the pair’s definite values persist throughout. Only the first mover’s conclusion is sound; the second mover’s is void. The resolution demands one admission: *the dynamics of quantum mechanics, read statistically, carry an implicit time-ordering of events*. Which of two mutually non-local perturbations fired first is a fact, and physical conclusions depend on it.

Frauchiger and Renner themselves meet this structure in their paper. Analyzing their protocol within Bohmian mechanics (the realist, definite-values model of Chapter 10), they note that “the time order of the measurements carried out by agents W and $\overline{W}$ is relevant,” with one agent’s conclusion valid and the other’s invalid according to who measured first, and they call this “a clear departure from standard quantum mechanics,” where the order of operations on separate systems is irrelevant.³ From our standpoint that is no departure at all: the order-dependence is not an artifact of Bohm’s particular model but a structural feature of *any* reading that keeps definite values. We have just derived it from transition matrices alone, with no pilot wave in sight.

³Frauchiger and Renner, “Quantum theory cannot consistently describe the use of itself,” *Methods*, “Analysis within Bohmian mechanics.”

11.5 The EPR paradox

Section 2.6 built a correlated pair out of two ordinary bits: COIN on the first, then CX to record its value onto the second, leaving

$$\mathbf{P}_{\text{correlated}} = \begin{pmatrix} 0.5 \\ 0 \\ 0 \\ 0.5 \end{pmatrix}.$$

Only two configurations carry any probability, 00 and 11. Separate the bits by any distance you like and read one of them. If it reads 0 the pair must be 00, so the far bit is 0; if it reads 1 the pair must be 11, so the far bit is 1. That is Bayes' theorem of Section 2.9 applied to a distribution with two entries in it: reading one bit locally settles the value of a bit that may be arbitrarily far away.

Nothing non-local has happened, because the far bit held that value all along. The CX gave the two bits matching values while they were in the same place, and each has held its value since. What the reading changed is the observer's knowledge of the pair, and a change in the subjective description expresses no physical change.

Now take the other view, on which the bits hold no values until they are read: the value is genuinely undetermined beforehand and manifests at the moment of measurement. This is the collapse picture. Alice reads her bit and a value manifests for it; Bob reads his and a value manifests for his; the two events happen far apart. If the two events are independent of one another, they agree only half the time, as independent fair draws do. They agree every time. So the two events are not independent: one outcome depends on the other across whatever distance separates them, which is what non-locality means. Nothing need travel between them for that to be so.

If non-locality is unacceptable, the natural conclusion is the one the classical pair already illustrated: the bits acquire definite values when they interact locally, and carry them afterward. Quantum mechanics stands in the way. A bit has a value on the basis it is read in, and no bit has definite values on every basis at once. Chapter 10 made that a theorem: no assignment of values to all the observables exists at all. Zeilinger has gone as far as proposing this as the foundational principle of quantum information: an elementary system carries one bit, the answer to a single question, rather than a bit for each basis one might have chosen.⁴

Albert Einstein, Boris Podolsky and Nathan Rosen took the difficulty up in 1935, and their argument is commonly called the EPR paradox.⁵ Their own construction used the positions and momenta of two particles. What follows recasts it onto two-valued quantities, as Bohm did, and runs it on bits.⁶ Run the program of Section 3.5, which is the one above

⁴A. Zeilinger, "A Foundational Principle for Quantum Mechanics," *Foundations of Physics* 29, no. 4 (1999), 631–643.

⁵A. Einstein, B. Podolsky and N. Rosen, "Can Quantum-Mechanical Description of Physical Reality Be Considered Complete?" *Physical Review* 47 (1935), 777–780.

⁶D. Bohm, *Quantum Theory* (New York: Prentice-Hall, 1951), §§22.15–22.18.

with the COIN replaced by a Hadamard, and the pair comes out in the Bell state

$$\psi = \frac{1}{\sqrt{2}} \begin{pmatrix} 1 \\ 0 \\ 0 \\ 1 \end{pmatrix}.$$

Separate the bits and have Alice read hers on the Z basis. The Born distribution is $\mathbf{p}_{\text{correlated}}$ again, so reading her own bit lets her predict Bob's with certainty, and she has done nothing to Bob's bit in the process. Their criterion makes that enough:

If, without in any way disturbing a system, we can predict with certainty (i.e., with probability equal to unity) the value of a physical quantity, then there exists an element of physical reality corresponding to this physical quantity.

So Bob's Z value is real. Now try the other two bases. Section 10.2 recast a rotated reading as a perturbation followed by an ordinary one: the X basis means applying H to each bit first, the Y basis $S^\dagger$ and then H. Applied to the Bell state, the three distributions are

$$|\psi|^2 = \begin{pmatrix} 0.5 \\ 0 \\ 0 \\ 0.5 \end{pmatrix}, \quad |(H \otimes H)\psi|^2 = \begin{pmatrix} 0.5 \\ 0 \\ 0 \\ 0.5 \end{pmatrix}, \quad |(\text{HS}^\dagger \otimes \text{HS}^\dagger)\psi|^2 = \begin{pmatrix} 0 \\ 0.5 \\ 0.5 \\ 0 \end{pmatrix}.$$

On Z and X the bits agree; on Y they disagree. Whichever of the three bases Alice picks, reading her own bit on that basis lets her predict, with certainty and without disturbing it, what Bob's bit would give if it were read on the same basis. The choice is hers alone, made after the bits have separated, and by the paper's other premise it reaches Bob not at all: "since at the time of measurement the two systems no longer interact, no real change can take place in the second system in consequence of anything that may be done to the first system." Whatever is real at Bob therefore cannot depend on which basis Alice picks. All three predictions are available to her, so all three values must have been real already, while ψ assigns a definite value on no basis at all.

A *complete* ψ describes everything there is to a system between measurements, so that the system holds no values of its own and its bits acquire values only on being read. An *incomplete* ψ has to be supplemented with values the system carries alongside it. The distinction says nothing about determinism. An incomplete ψ requires no deterministic theory underneath it, only that ψ leave something out. Einstein, Podolsky and Rosen conclude that ψ is incomplete, and the argument above makes it incomplete on every basis at once. The two classical bits restored locality by carrying their values from the moment they interacted, and nothing less than all three values carried the same way will do it here.

The obvious reply is that only the value Alice actually predicted was ever real. They raised it themselves, in the paper's last full paragraph, and refused it:

This makes the reality of P and Q depend upon the process of measurement carried out on the first system, which does not disturb the second system in any way. No reasonable definition of reality could be expected to permit this.

Read that with Chapter 10 in hand and it says more than its authors meant. They meant it as a refutation: the consequence could not stand, so the premise behind it had to go, and the premise they gave up was the completeness of ψ . Kochen and Specker closed that route afterward. The values cannot all be definite together, so the escape Einstein wanted is unavailable, and what survives of his dilemma is the choice he thought absurd. What is real at Bob does depend on what Alice does, far away and after they have parted. Einstein expressed disbelief at this, referring to it as “spooky action at a distance.”⁷

The contrast with the classical pair is exact. There, reading one bit changed only the observer’s knowledge of the other, and a change in the subjective description expresses no physical change. Here the reading is a stochastic perturbation of what was there, which is how Section 10.4 describes the measurement of an emergeable, and Section 11.3 already showed that such a perturbation cannot be confined to Alice’s bit.

Staggering the two perturbations locates what cannot be done locally. Section 10.2 recast a rotated reading as a disturbance followed by an ordinary one, so measuring both bits on X means applying H to each and then reading the ontic basis. Let the two H gates act on the Bell state ψ at different times, Alice’s first and Bob’s after it, with both readings deferred until each gate has acted. The pair’s distribution runs

$$\begin{pmatrix} 0.5 \\ 0 \\ 0 \\ 0.5 \end{pmatrix} \xrightarrow{H \otimes I} \begin{pmatrix} 0.25 \\ 0.25 \\ 0.25 \\ 0.25 \end{pmatrix} \xrightarrow{I \otimes H} \begin{pmatrix} 0.5 \\ 0 \\ 0 \\ 0.5 \end{pmatrix},$$

Alice’s gate spreading the pair over all four configurations and Bob’s restoring the correlation, since $H \otimes H$ returns the Bell state to itself.

Only one of the two steps is in any doubt. Alice’s gate carries a perfectly correlated pair to a uniform one, and a fair coin on her own bit does that much: $\Gamma = \begin{pmatrix} 0.5 & 0.5 \\ 0.5 & 0.5 \end{pmatrix}$ satisfies $(\Gamma \otimes I) \mathbf{p}_{\text{correlated}} = \mathbf{p}_{\text{uniform}}$ exactly, with both columns summing to one and every entry a probability. Nothing non-local is asked of her. Bob’s step is the one to put to the test of Section 11.3, and it is the step that restores the correlation.

Write the candidate out entry by entry as before, $\Gamma = \begin{pmatrix} \gamma_{00} & \gamma_{01} \\ \gamma_{10} & \gamma_{11} \end{pmatrix}$, acting now on Bob’s bit, with nothing assumed about the four numbers. Working out $(I \otimes \Gamma) \mathbf{p}_{\text{uniform}}$ and setting it equal to $\mathbf{p}_{\text{correlated}}$ gives

$$(I \otimes \Gamma) \mathbf{p}_{\text{uniform}} = \frac{1}{4} \begin{pmatrix} \gamma_{00} + \gamma_{01} \\ \gamma_{10} + \gamma_{11} \\ \gamma_{00} + \gamma_{01} \\ \gamma_{10} + \gamma_{11} \end{pmatrix} = \begin{pmatrix} 0.5 \\ 0 \\ 0 \\ 0.5 \end{pmatrix}.$$

The first component asks that $\gamma_{00} + \gamma_{01}$ equal 2 and the third asks that the same sum equal 0. The second and fourth do the same to $\gamma_{10} + \gamma_{11}$, in the other order. There is no solution, and this time no appeal to the bounds on a probability is needed: the equations contradict one another before any question of range arises.

The reason they do is the uniformity Alice’s gate left behind. Bob’s bit is as likely to face a 0 on her side as a 1, so a matrix on his bit alone returns the same distribution either way,

⁷A. Einstein to M. Born, 3 March 1947, in *The Born–Einstein Letters*, trans. I. Born (London: Macmillan, 1971).

while the correlation it is asked to produce requires his bit to come out 0 when hers is 0 and 1 when hers is 1. A law reading only his bit cannot tell the two cases apart. Nor can one reading his bit together with a stored copy of hers, for the reason Section 11.3 gave: the copy was made before her gate acted and no longer reports what her bit holds. Section 4.9 named the state Alice's gate leaves behind: its classical mutual information is zero while its quantum mutual information is maximal, the correlation held in the phases where measurements of values cannot see it. Bob's gate is the local rotation that brings it back, and no matrix over the values can do that work, since the values are where the correlation is not.

Reversing the order of the two gates exchanges the two roles and leaves the final distribution untouched, so which perturbation is the one no local matrix accounts for is settled by which gate acts first. Where the two are space-like separated, Section 11.4 has already made that a question about the slicing.

Quantum theory as written is therefore not local. What reaches across is the law Bob's gate follows, which consults a value his side does not hold; no disturbance travels, and neither gate moves the bit it does not touch, which is the pseudo-locality of Section 6.7.

In a case this simple, locality can be recovered by adding to the theory from outside, and the addition Einstein, Podolsky and Rosen were after is a hidden memory the two bits carry away from the preparation. A model along those lines is easy to write down, and seeing how far it gets is the best way to see what the no-go theorems are for. We will call it the *strong value model*.

Let a bit hold +1 and -1 in place of 0 and 1, and extend its memory so that it carries three values at once, x , y and z , one for each basis it might be read in. Call the result a *strong object* and write it as a vector,

$$\mathbf{s} = \begin{pmatrix} 1 \\ x \\ y \\ z \end{pmatrix},$$

the leading entry standing for the identity observable, whose value is always 1. A reading on a given basis returns the entry belonging to that basis and nothing else is consulted. The three values cannot be drawn independently, since Y is not independent of X and Z : the entries must satisfy $y = \pm xz$, so fixing z at +1 and drawing x leaves a strong object one free coin rather than two. Every unitary operator then acts on $\mathbf{s}$ through its Pauli transfer matrix, the real matrix defined in Section 7.5, and for H that matrix is the signed permutation displayed there: it exchanges the x and z entries and flips the sign of y . Where a gate acts on more than one strong object, the objects are combined with the tensor product, the gate's matrix is applied to the combination, and the result is factored back into separate objects.

That much reproduces the EPR experiment. Running the preparation leaves the two objects agreeing on their x and z entries and disagreeing on y , which is the agreement on Z , the agreement on X , and the disagreement on Y that the three distributions above record, and the six mismatched pairs of settings come out uniform as they should. The staggered circuit is the same model in motion: both objects begin showing z and agree, Alice's gate exchanges her entries so that she shows x while Bob still shows z , and a reading then compares two independent values and finds no correlation. Bob's gate exchanges his, the reading compares x against x , and the correlation is back. It was never destroyed, only moved into the entry

nobody was reading, which is the part the phases played in the quantum account. Nothing crosses the distance between them, and every value Alice can predict with certainty was written into Bob's object before the two parted, which is the element of physical reality their criterion asked for.

What defeats the model is the factoring step, which is also the whole of its locality. Run the GHZ circuit of Chapter 10 through it and three of the four runs come out right while the fourth does not. Skip the factoring, keep the combined vector, and all four come out right, which sounds like a repair until one asks what the combined vector is. For n objects it carries 4^n entries, and it is the quantum state itself written in another basis, the Pauli basis of Section 3.4. The entries that hold the correlations belong to no individual object, so keeping them means keeping a fact about the group somewhere that is local to none of its members. Factoring is what makes the model local and is what makes it wrong about GHZ, and the cure gives up the thing the model was built for.

The model is trivial enough to write down on a napkin, and it goes further than one expects: it settles Bohm's version of the EPR experiment completely, across all nine pairings of the three bases, it runs Deutsch's algorithm (but not Deutsch-Jozsa), and it runs superdense coding, recovering all four messages on every run and as many pairs in parallel as one cares to prepare. The GHZ experiment that defeats it uses three objects and no gate more exotic than the ones already named.

How easily a model of this kind reproduces part of quantum mechanics is what makes it tempting to think the rest will come with effort. Einstein thought so and tried. In May 1927 he presented a paper to the Prussian Academy assigning definite trajectories to the solutions of Schrödinger's equation, a completion of just the sort the EPR argument would later say was needed, and he withdrew it before publication.⁸ A postscript he added to the proofs gives the reason: for two parts that do not interact, his scheme made the motion of each depend on the coordinates of the other. What defeated his own attempt is the separability the EPR paper would go on to take as a premise. A completion that did work came from Bohm a quarter of a century later, and Einstein would not have that either. Bohm's theory is deterministic and reproduces the predictions in full, and all Einstein wrote to Born about it was that "that way seems too cheap to me."⁹ He gave no reason there, though Bohm's particles are guided by a wave that answers to the whole arrangement however far apart its pieces are, which is what his own attempt had foundered on.

The strong value model shows how misleading a local model can be. It gets enough right to make a fully local account of quantum mechanics look like a matter of further work. Yet the GHZ argument of Chapter 10, the two-bit distillation of Section 11.1, and Bell's theorem of the next section rule out this model and every way of patching it.

The same reading settles what Bell's theorem shows. The result is usually reported as ruling out "local realism," where denying the realism means denying that the bits hold definite values between measurements. That phrasing invites the reading that locality survives once the definite values are dropped, and that realism is what brought the non-locality in. The two halves of the argument say otherwise. Einstein, Podolsky and Rosen showed that dropping the

⁸A. Einstein, "Bestimmt Schrödingers Wellenmechanik die Bewegung eines Systems vollständig oder nur im Sinne der Statistik?" (Berlin, 5 May 1927), Einstein Archive 2-100, Doc. 516 in *The Collected Papers of Albert Einstein*, Vol. 15 (Princeton: Princeton University Press, 2018).

⁹A. Einstein, letter to M. Born, 12 May 1952, in *The Born-Einstein Letters*, trans. I. Born (London: Macmillan, 1971), letter 99.

values gives non-locality, since nothing is then left to coordinate the two outcomes. Section 11.6 shows that keeping them gives non-locality as well. Neither choice is local, and the non-locality belongs to the correlations rather than to any view about what exists between measurements.

What survives is the contextuality of Chapter 10. The bits do hold definite values before anyone looks, but on one basis at a time, never on all three taken separately. Once correlated bits are read on different bases, what is measured cannot be accounted for locally.

The claim that ψ is complete goes further than denying values on the other two bases. The bit then holds no value at all before it is read, not even on one basis, and manifests one at every reading. That is the case already treated above: two manifestations far apart, with nothing making them agree, give uncorrelated outcomes. The only way around non-locality in this picture is to accept that the outcomes really are uncorrelated and to hold that reality fragments, so that no observer is ever placed to notice. Section 11.12 takes that reading up, and concerns with it. Short of that, the choice is between non-locality and retrocausality, which Section 11.10 examines, and the rest of this chapter is written on the first.

11.6 Bell's theorem

This section is an aside, and a reader who would rather skip it can do so without losing the thread: the argument resumes in Section 11.7 from where Section 11.4 left it. The mathematics here is heavier than the rest of the book, and it is written with integrals, though everywhere else we keep deliberately to discrete quantities.

The argument just completed is a single-run argument: one pair, one outcome $\alpha\beta = 10$, two irreconcilable stories about its past. Bell's 1964 theorem reaches the same conclusion from ensembles and an inequality instead.¹⁰ It is the argument the literature cites and the experiments test, and it convicts locality in a different place than Section 11.3 did, which the end of this section explains.

Bell works with two spin-half particles in the *singlet* state, prepared together and then carried apart. It is one bit-flip away from the Bell state of Chapter 3: the same circuit, run on $\mathbf{e}_{11}$ instead of $\mathbf{e}_{00}$:

$$\psi_- = CX(H \otimes I_2)(\mathbf{e}_1 \otimes \mathbf{e}_1) = \frac{1}{\sqrt{2}} \begin{pmatrix} 0 \\ 1 \\ -1 \\ 0 \end{pmatrix}.$$

Alice measures her bit along an axis at angle a , Bob his along an axis at angle b . Chapter 10 already told us what a rotated measurement is: a perturbation, followed by an ordinary reading in the ontic basis, with the perturbation given by the recasting unitary. For an axis at angle θ that unitary is the rotation gate of Chapter 3,

$$R_Y(\theta) Z R_Y(\theta)^\dagger = \cos \theta Z + \sin \theta X,$$

so measuring at angle θ means applying $R_Y(-\theta)$ to the bit and then reading it.

¹⁰J. S. Bell, "On the Einstein Podolsky Rosen paradox," *Physics* 1 (1964), 195–200; reprinted in *Speakable and Unspeakable in Quantum Mechanics*, pp. 14–21.

Write $E(a, b)$ for the correlation between the two readings: the probability that the bits agree minus the probability that they disagree,

$$E(a, b) = P(00) + P(11) - P(01) - P(10),$$

with the four probabilities given by the Born rule of $(R_Y(-a) \otimes R_Y(-b)) \psi_-$. The chapter's code evaluates it at any pair of angles; the answer, every time, is

$$E(a, b) = -\cos(a - b).$$

Two features of that formula matter for what follows. It depends only on the *difference* of the two angles. And at $a = b$ it equals -1 : measured along the same axis, the two bits *never* agree, with certainty, however far apart they have been carried.

The assumption on trial

Now Bell's question: what would it take for the two outcomes to be produced *locally*? Let λ stand for whatever the pair really is as it leaves the source. On our reading λ is the pair's ontic configuration, distributed over runs by some $\rho(\lambda)$, which is nothing but a base-case distribution in the sense of Chapter 6. Locality of the law is then the demand that Alice's outcome be a function of her own setting and λ alone, and Bob's of his,

$$A(a, \lambda) = \pm 1, \quad B(b, \lambda) = \pm 1,$$

with no b appearing in the first and no a in the second, so that the observed correlation is the average of their product,

$$E(a, b) = \int d\lambda \rho(\lambda) A(a, \lambda) B(b, \lambda).$$

We must be precise about what has just been assumed, because it is often stated incorrectly. Non-contextuality is not assumed, by Bell or by us. What is taken to be definite is λ , the beables of Section 10.4, and the outcome of a measurement along a is $A(a, \lambda)$: a lawful function of the real state and the setting, which is exactly the criterion of Section 10.5 and exactly what an emergeable is. What is assumed on top of that is that the law producing Alice's value never consults Bob's setting. That is the same assumption Section 11.3 caught failing at a single transport, and Bell is about to show it failing again, by a different method.

The inequality

Perfect anticorrelation is used first. Since $E(a, a) = -1$ while A and B take only the values ± 1 , their product must equal -1 for almost every λ , which forces $B(b, \lambda) = -A(b, \lambda)$: Bob's function is Alice's, negated. Substituting it and subtracting two correlations,

$$E(a, b) - E(a, c) = - \int d\lambda \rho(\lambda) [A(a, \lambda)A(b, \lambda) - A(a, \lambda)A(c, \lambda)].$$

Since $A(b, \lambda)^2 = 1$, the bracket factors:

$$E(a, b) - E(a, c) = - \int d\lambda \rho(\lambda) A(a, \lambda)A(b, \lambda) [1 - A(b, \lambda)A(c, \lambda)].$$

The factor $A(a, \lambda)A(b, \lambda)$ has magnitude 1, and the bracket is never negative, so taking absolute values can only increase the right-hand side:

$$|E(a, b) - E(a, c)| \leq \int d\lambda \rho(\lambda) [1 - A(b, \lambda)A(c, \lambda)] = 1 + E(b, c),$$

the last equality because $E(b, c) = -\int \rho A(b)A(c)$ by the same substitution. That is Bell's inequality in its original form. Every step used only that the values exist and that each is produced locally; no quantum mechanics entered anywhere.

The violation

Now feed it the quantum correlations. Put the three axes at 0° , 60° , and 120° :

$$E(a, b) = -\cos 60^\circ = -\frac{1}{2}, \quad E(a, c) = -\cos 120^\circ = +\frac{1}{2}, \quad E(b, c) = -\cos 60^\circ = -\frac{1}{2},$$

and the inequality asserts

$$|-\frac{1}{2} - \frac{1}{2}| = 1 \leq 1 - \frac{1}{2} = \frac{1}{2},$$

which is false by a factor of two. Quantum mechanics predicts the left-hand side, and the laboratory has reported the same in every Bell test ever run. So the premise fails: *no law that produces Alice's value from her own setting and the pair's configuration alone can reproduce the correlations of the singlet*. Note again what the theorem leaves untouched. Nothing in the derivation questioned whether λ was definite, and the theorem runs word for word if λ is as definite as we insist. What is refuted is the locality of the *law*, exactly as Chapter 10's theorem constrained the laws and left the values alone, and as Chapter 12's will.

Why three axes

This argument and Section 11.3's fail locality in different places, and the difference explains why Bell needed three axes where we needed one gate.

In Section 11.3, locality failed at a *single* transport: one perturbation, one pair of distributions, and the entries the local matrix would have required came out at $3/2$ and $-1/2$. Ask the same question here, of the singlet with Alice alone perturbing, and the answer is the opposite. A perfectly legal local matrix exists, at every angle:

$$\Gamma_A(a) = \begin{pmatrix} \cos^2(a/2) & \sin^2(a/2) \\ \sin^2(a/2) & \cos^2(a/2) \end{pmatrix},$$

non-negative, columns summing to one, transporting the pair's distribution correctly (the code checks it). Any *one* setting-pair of the Bell experiment can be explained locally. What cannot be done is to explain all three at once, with a single λ standing behind them, and that is precisely what the inequality tests and what the third axis is for. Bell's theorem is a constraint on a whole family of experiments rather than on any one member of it, which is why it had to be stated as an inequality, and why the argument went undiscovered for the thirty years after Einstein, Podolsky and Rosen.

The transition matrices of Section 11.4 are non-local in exactly the way Bell's theorem demands. Their columns are indexed by the configuration of *both* bits, so Alice's hop consults Bob's value as a matter of law: Γ_{Alice} sent the configuration 10 to 00 with certainty while sending 11 to a fair coin. Bell proves that some such consultation is mandatory for anyone who keeps the values. The next section works out what that costs.

11.7 The collision with relativity

We now return to the transition matrices of Section 11.4. Alice's perturbation is an operation on her bit; Bob's is an operation on his. Nothing prevents the two bits from being carried far apart before the perturbations fire, and nothing prevents the two firings from happening in such quick succession that light cannot cross the gap in between: the two events can be *space-like separated*.

In special relativity, the time-ordering of space-like separated events is relative: observers in different frames of reference disagree about which came first, and the geometry of Minkowski space declares the disagreement unresolvable, because no signal can connect the events. Ordinarily the relativity of that ordering is harmless: if nothing traveling between the events could depend on the order, the order carries no physics. But we have just computed that something *does* depend on it. If Alice fired first, the pair was 11 and her account of its history is the true one; if Bob fired first, the pair was 00. The configuration the pair actually occupied (in every run, per the first requirement, a definite fact) differs between the two orderings. Let the ordering be relative, and the fact becomes relative: in one frame the pair's past is 11, in another it is 00. The values of the bits fail to be *Lorentz invariant*.

Replace the bits with brain-bits and we meet the same failure as before. If the configurations of the particles composing an observer's memory vary with the reference frame of whoever describes them, then other minds hold definite states only relative to a frame, and Chapter 9's third requirement falls, by the now-familiar route, into solipsism. This is the same verdict the previous section extracted from Bell's inequality: the laws governing the values must be non-local, and non-local laws cannot have Lorentz-invariant content. The Frauchiger–Renner argument, however, needs no inequalities and no ensembles: one run, one $\alpha\beta = 10$, two irreconcilable pasts.

Three constraints now stand together. The values must be invariant (requirement three). The values cannot be *Lorentz invariant* (just shown). The only invariance left standing is the *Galilean* kind: there must be an *absolute* fact about which perturbation fired first.

11.8 Apparent relativity

That sounds like a declaration of war on special relativity. It is not, and the person who explained why was, once again, Bell, in an essay titled “How to teach special relativity.”¹¹

Bell's essay retells special relativity in the manner of Lorentz rather than Einstein: fix a single reference frame, assume the dynamical laws take their usual form in it, and derive what happens to *equipment* that moves. Rods physically contract and clocks physically slow, as consequences of the same laws that hold the equipment together. A moving observer, her instruments altered in exactly compensating ways, notices nothing; misled in synchronizing her clocks, she adopts a shifted time variable and, in Bell's narration, “is led to think” that it “is the real time.” When two such observers compare notes, each attributes the other's variables to physical effects in the other's equipment, and, as Bell says of the disagreement, “Her view will be logically consistent and in perfect accord with the observable facts. He will have no way of persuading her that she is wrong.”

¹¹J. S. Bell, “How to teach special relativity,” *Progress in Scientific Culture* 1, no. 2 (1976); reprinted in *Speakable and Unspeakable in Quantum Mechanics*, pp. 67–80.

On the philosophy, Bell is explicit: Einstein declares the notions “really resting” and “really moving” meaningless, since experiment cannot separate them; “Lorentz, on the other hand, preferred the view that there is indeed a state of real rest, defined by the ‘aether’, even though the laws of physics conspire to prevent us identifying it experimentally.” And the verdict: “The facts of physics do not oblige us to accept one philosophy rather than the other.” Every formula of special relativity survives untouched on either philosophy; the choice between them is not an experimental question.

Lorentz’s philosophy is exactly what our argument needs, and its ingredients need names (the names are our framing, not Bell’s). Designate one rest frame as privileged. Rods and clocks at rest in that frame measure *real* space and *real* time; rods and clocks in motion relative to it are physically altered (genuinely contracted, genuinely slowed) and so measure quantities that deviate from the real ones: *apparent* space and *apparent* time. Observers still disagree about simultaneity, but the disagreement is now a disagreement among distorted instruments, with a fact of the matter underneath. The technical name for the privileged frame’s family of constant-time surfaces is a *preferred slicing*, or preferred foliation, of spacetime. Call the resulting picture *apparent relativity*: Minkowski geometry as the phenomenology of moving instruments, Galilean order underneath. The view has a name in the literature, where a reading of this shape, Newtonian structure underneath with a preferred frame, is called the *neo-Lorentzian* interpretation of special relativity.¹²

Apparent suggests an illusion, and if the picture is right then special relativity is one. But to call something an illusion is to say it is unreal, and that says nothing until one has answered what is real. Answer it and the illusion dissolves. What was being called an illusion turns out to be a real feature of the world, emergent rather than fundamental, which is why *emergent relativity* would serve as a name just as well.

The illusion holds only for someone who takes space and time to be whatever rods and clocks report. On this reading those reports are emergent, produced by length contraction and time dilation acting on real space and real time, so what an instrument measures is fixed by the real quantities together with the physical effects the instrument has undergone. That is the relation Section 10.5 drew between what is really there and what a measurement returns: an emergeable is lawfully determined by the beable and the arrangement doing the measuring, and it is no less real for being derived.

There is a middle ground between special relativity and the neo-Lorentzian reading. One can hold that spacetime is not Galilean, with velocities adding as relativity has them add and light keeping one speed in every frame, and still single out one frame as preferred. That is all the preceding sections require: an absolute fact about which perturbation fired first. Whether length contraction and time dilation are on top of that explained as physical effects on moving equipment, the way Bell explains them, is a further question the argument leaves open.

Einstein’s objection to privileging a frame is the one we have met before: the mathematics is symmetric under change of frame, so no experiment can identify which frame is the real one. That is the underdetermination objection, and Chapter 10 taught us its exact weight. It cannot rule the preferred frame out; it establishes only that *if* one exists, experiment will not locate it. And our argument stands against the temptation to wield Occam’s razor (“undetectable, therefore dispensable”): dispensing with the preferred ordering is *inconsistent*

¹²The characterization, and an argument against it, are in Y. Balashov and M. Janssen, “Presentism and Relativity,” *British Journal for the Philosophy of Science* 54 (2003), 327–346.

with keeping invariant values. The GHZ and Wigner’s-friend arguments were, in effect, a proof by contradiction that some basis must be ontologically privileged even though the mathematics treats all bases alike. The Bell and Frauchiger–Renner arguments are a proof by contradiction of precisely parallel shape: some frame must be ontologically privileged even though the mathematics treats all frames alike. In both cases the symmetry of the formalism is real, and in both cases the realist learns that the symmetry cannot extend all the way down into the ontology.

As with the ontic basis, two postures are available, and we regard both as legitimate. One may treat the choice as a convention fixed per context. Or one may argue for a particular frame on grounds beyond present experiment. General relativity, in the Friedmann solutions that describe our actual expanding universe, spontaneously singles out a cosmic rest frame: the frame in which the expansion is isotropic and cosmic time is defined. It is measurable. The cosmic microwave background carries a dipole anisotropy from our motion through it, and from that dipole the solar system’s velocity relative to the universe’s rest frame has been measured at roughly 370 kilometers per second.¹³ If a preferred slicing must exist, the slicing of cosmic time is the obvious candidate, a frame the universe itself supplies rather than one any observer must invent. On this reading, which perturbation *really* fired first is decided by clocks at rest with respect to the universe; every frame-dependent disagreement about the order is apparent, an artifact of distorted instruments; and the pair of bits has exactly one past.

11.9 Non-locality without signals

One question about non-local laws remains to be answered. A reader who has accepted them is entitled to ask why laboratories cannot exploit them: if the outcome of Alice’s perturbation depends on Bob’s distant bit, can Bob not wiggle his bit and send Alice a message faster than light? The answer is no. The proof is given outright, in the same element-by-element style as everything else here, because it shows *exactly* how the theory manages both at once.

Take a completely arbitrary two-bit quantum state and a completely arbitrary gate applied to the *first* bit only:

$$\psi = \begin{pmatrix} a \\ b \\ c \\ d \end{pmatrix} \quad U = \begin{pmatrix} u_{00} & u_{01} \\ u_{10} & u_{11} \end{pmatrix}, \quad (U \otimes \mathbf{I}) \psi = \begin{pmatrix} u_{00}a + u_{01}c \\ u_{00}b + u_{01}d \\ u_{10}a + u_{11}c \\ u_{10}b + u_{11}d \end{pmatrix}.$$

Compute the marginal Born probability that the *second* bit reads 0, after the gate. Two configurations have the second bit 0 (the first and third entries), so:

$$P(B=0) = |u_{00}a + u_{01}c|^2 + |u_{10}a + u_{11}c|^2.$$

Both square magnitudes expand by $|z|^2 = z\bar{z}$ of Section 3.1, the overbar being the conjugate defined there. Taking the first of them,

$$|u_{00}a + u_{01}c|^2 = (u_{00}a + u_{01}c)(\overline{u_{00}a + u_{01}c}) = |u_{00}|^2|a|^2 + |u_{01}|^2|c|^2 + u_{00}\overline{u_{01}}a\bar{c} + \overline{u_{00}}u_{01}\bar{a}c,$$

¹³Planck Collaboration, “Planck 2018 results. I. Overview and the cosmological legacy of *Planck*,” *Astronomy & Astrophysics* 641 (2020), A1.

where the last two terms are conjugates of one another, so their sum is twice the real part of either. The second square magnitude expands the same way with u_{10} and u_{11} in place of u_{00} and u_{01} . Adding the two and collecting like terms:

$$P(B=0) = (|u_{00}|^2 + |u_{10}|^2) |a|^2 + (|u_{01}|^2 + |u_{11}|^2) |c|^2 + 2 \operatorname{Re}[(u_{00}\overline{u_{01}} + u_{10}\overline{u_{11}}) a \bar{c}].$$

Now recall what it is to be a valid gate. Chapter 3's condition was $UU^\dagger = \mathbf{I}$; and since quantum information is time-symmetric (it is a matter of convention which of U and $U^\dagger$ is the time-forward operator), the product in the other order is the identity as well, $U^\dagger U = \mathbf{I}$. Written out, the dagger transposes and conjugates, so that condition is

$$U^\dagger U = \begin{pmatrix} \overline{u_{00}} & \overline{u_{10}} \\ \overline{u_{01}} & \overline{u_{11}} \end{pmatrix} \begin{pmatrix} u_{00} & u_{01} \\ u_{10} & u_{11} \end{pmatrix} = \begin{pmatrix} |u_{00}|^2 + |u_{10}|^2 & \overline{u_{00}}u_{01} + \overline{u_{10}}u_{11} \\ \overline{u_{01}}u_{00} + \overline{u_{11}}u_{10} & |u_{01}|^2 + |u_{11}|^2 \end{pmatrix} = \begin{pmatrix} 1 & 0 \\ 0 & 1 \end{pmatrix}.$$

Its four entries are the very quantities standing in the expansion, so there is nothing to do but substitute. The diagonal entries set $|u_{00}|^2 + |u_{10}|^2$ and $|u_{01}|^2 + |u_{11}|^2$ to 1, and the upper-right entry sets $\overline{u_{00}}u_{01} + \overline{u_{10}}u_{11}$ to 0, whose conjugate is the bracket multiplying $a\bar{c}$. The interference term therefore vanishes identically, and what remains is

$$P(B=0) = |a|^2 + |c|^2,$$

which is the second bit's marginal *before* the gate. The same cancellation gives $P(B=1) = |b|^2 + |d|^2$. No unitary applied to the first bit (for *any* state, entangled or not, and *any* gate) can move the second bit's marginal statistics by any amount whatsoever. What we have verified entry-by-entry for one bit acting on another, the physicist Philippe Eberhard proved in full generality: local operations on one subsystem cannot alter the marginal probabilities of another.¹⁴

So the non-locality runs in exactly one direction. The outcome of a perturbation on Alice's bit can *depend on* Bob's bit; nothing Alice does *moves* the statistics of Bob's bit. A dependence is not a lever. And here is the apparent loophole: Alice perturbs her bit, which consults Bob's, so surely the consultation leaves a trace Bob can read. It fails for a reason the reader can now diagnose in one line: to find out how her own stochastic perturbation came out, Alice would have to measure her bit, and measuring it creates a base case, maximal QMI between her and the bit, which severs the very entanglement the scheme meant to exploit.

The theory is non-local in its laws and no-signaling in its statistics, and both properties are theorems of the same formalism. What we have checked here for one gate on one bit of a pair is the smallest case of a general result, known in the literature as the *no-communication theorem*: no operation carried out on one part of a composite system can alter the statistics of any other part, so no arrangement of shared entanglement can be made to carry a message.

That theorem is also what keeps our own reading consistent. Section 11.8 left the preferred slicing undetectable, and an undetectable posit invites suspicion. Non-locality, by the theorem, never appears directly. It is always inferred, from statistics gathered and compared afterwards, and never witnessed in the moment as a disturbance in anybody's bit.

Return to the contradiction of Section 11.2 with that in hand. Alice retrodicts that the pair was $AB = 11$ and Bob that it was $AB = 00$, and they cannot both be right. Neither ever

¹⁴P. H. Eberhard, "Bell's theorem and the different concepts of locality," *Il Nuovo Cimento B* 46 (1978), 392–419.

saw the configuration they are disagreeing about. Both perturbed the pair, and between them the perturbations destroyed the very values in dispute, so the disagreement is between two inferences about a past of which no record survives. It follows that the disagreement has no physical consequence: no measurement anyone can perform depends upon which retrodiction is the correct one, and the no-communication theorem is exactly what forbids the difference from surfacing anywhere.

The undetectability of the slicing is therefore harmless rather than fatal. *Any* preferred slicing resolves the contradiction, and any two of them resolve it in ways no experiment distinguishes. To show the reading consistent we therefore do not have to identify the true slicing. We have only to state that one exists, since the existence alone supplies a fact of the matter about which perturbation came first, and every candidate fact discharges the contradiction equally well. The point is not that we can identify the slicing, only that one must be assumed to exist at all for the account to be logically consistent.

The familiar picture of entanglement does not survive this. In that picture one member of a pair is twisted and the other, however far away, instantly twists to match: an action here, an answering action there. Nothing of the sort happens. Alice's perturbation changes Alice's bit and nothing else. What Bob's bit supplies is not a response but an input, because how Alice's bit comes out depends on the configuration Bob's is in; his bit is consulted, not disturbed. Its marginal statistics are precisely what they were, for every state and every gate, as the calculation above establishes, and an experimenter sitting with Bob's bit would find nothing whatever different about it at the moment Alice acted. Nothing propagated, because nothing was sent.

What replaces the image is the reading Chapter 8 reached by another route: the law consults the arrangement. Alice's bit hops under a law that takes the configuration of the pair as its input, Bob's half included. That is a claim about what the law reads, not about what Alice's gate touches.

That distinction also settles a claim frequently made on behalf of Many Worlds: that it is the one interpretation that keeps locality, and is to be preferred on that ground. The physicist Aurélien Drezet has examined the claim in detail using the GHZ correlations of Chapter 10.¹⁵ Within any single branch, every observer reports one definite outcome, and in every branch where the three results have been brought together, they satisfy the GHZ constraints. Those correlations are then subject to the same theorem as everyone else's, and no locally causal account of them exists; the standard reply, that the theorem tacitly assumes outcomes are unique, fails because the theorem needs only the outcome each observer reports within a branch. What Many Worlds does retain is no-signaling. But no interpretation includes signaling: no-signaling is a theorem of the shared formalism, and we have just proved it for a theory whose laws are overtly non-local. To redefine locality as no-signaling is therefore not to show that Many Worlds is local; it is to make every interpretation local by fiat. Locality in the ordinary sense, that an object *here* cannot influence an object *there* without something crossing the space between them, is violated by every reading that keeps the past fixed. Judged by the locality Bell defined, Many Worlds fails with everyone else; judged by the locality it advocates substitute, everyone else passes with it. Drezet's verdict is that Everett's theory is "either wrong or nonlocally-causal."

¹⁵A. Drezet, "An Elementary Proof That Everett's Quantum Multiverse Is Nonlocal," arXiv:2306.07794 (2023).

The non-locality has a simple picture behind it. When two particles interact locally and then separate, each carries its history with it, and the histories are not independent: they extend backward in time like two tails that still cross at the point of the old interaction. In the hypergraph language of Chapter 7, the pair edge written during the interaction persists after the particles part. A law that consults history (and Part II established that quantum laws do) thereby consults, through the crossing point, the other particle's present. Non-locality is what the non-Markovian picture looks like when a history is shared.

11.10 The other horn: retrocausality

The chapter has drawn a firm conclusion: the laws are non-local. But this particular conclusion hangs on a choice, and a different choice leads somewhere else. The paradox of the two bits was a collision between two tacit premises, and we resolved it by sacrificing one of them. The other is equally available. Whether the preferred frame survives that second choice is the harder question, and the answer turns on which form the alternative takes: one form appears to need the frame exactly as the non-local horn does, while another may escape it, by a route we come to at the end of the section.

The two premises are these. The first is that “ $AB = 01$ is impossible” is a settled fact about the pair, holding of every run and independent of what anyone will later choose to measure. Notice how a claim of that shape is established: one prepares the pair many times, measures on the ontic basis, never once sees $AB = 01$, and generalizes from those runs to a law covering all runs. Call it the premise of *induction*. The second is that each perturbing measurement reveals a reliable fact about the other bit: the premise of *inference*. When both parties measure and obtain $\alpha\beta = 10$, the two inferences combine to say the pair was in the configuration $AB = 01$, which the first premise declared impossible. One premise must go.

We sacrificed inference: the second mover's reading is unreliable because the perturbation consults a bit the first mover has already disturbed. That is the non-local horn, and it is the route Sections 11.6 and 11.7 took to the preferred frame. But one may instead sacrifice the induction. On this reading the impossibility of $AB = 01$ was never an unconditional fact. The runs that established it were runs in which the pair was going to be measured on the ontic basis, and nothing in them licenses the conclusion for runs in which it is not; the joint future choice to measure both bits on the perturbed basis *opens* $AB = 01$ as a genuine possibility for the pair's past. The future setting reaches back and edits what the past could have been. This is *retrocausality*, and it dissolves the paradox not through any non-local influence across space, but through influence that runs along time instead.

The two-state vector formalism of Yakir Aharonov and Lev Vaidman makes the retrocausal reading precise, and precise enough to compute.¹⁶ Its premise is the one this section is entertaining: that the laws carry no preferred direction in time. A boundary condition imposed at the end of an experiment is then as legitimate as one imposed at the beginning, and neither has a prior claim to be the cause of the other. Ordinary quantum mechanics uses one of the two, a state prepared at the start and evolved forward. This formalism uses both.

A state ψ_i is carried forward from the preparation and a second state ψ_f is carried backward from the outcome the run ends on, and a system between the two measurements is

¹⁶Y. Aharonov and L. Vaidman, “The Two-State Vector Formalism: An Updated Review,” arXiv:quant-ph/0105101 (2002).

described by the pair together. Neither alone settles what happens in between. Their overlap does, and it constrains the intermediate values more tightly than the forward state can on its own, since a history the preparation allows is struck out if it cannot also arrive at the outcome that was recorded.

Both states have to be brought to whatever moment t is being asked about. Write $U(t \leftarrow t_i)$ for the run's gates from the preparation up to t , and $U(t_f \leftarrow t)$ for the gates from t onward to the final measurement. Then

$$\psi_i(t) = U(t \leftarrow t_i) \psi_i, \quad \psi_f(t) = U(t_f \leftarrow t)^\dagger \psi_f.$$

The dagger is what turns the second one around. Section 3.3 showed $U^\dagger$ to be the *time-reverse* of U , and a time-reverse is precisely what the backward state calls for: carrying ψ_f back through a stretch of circuit means running that stretch's gates in reverse. So $\psi_i(t)$ is the preparation pushed forward to t and $\psi_f(t)$ is the recorded outcome pulled back to t , with the same gates used both ways. The weak value of an observable A at that moment is built from the two together,

$$A_w = \frac{\psi_f(t)^\dagger A \psi_i(t)}{\psi_f(t)^\dagger \psi_i(t)},$$

a definition due to Aharonov, Albert and Vaidman.¹⁷ Its denominator is the overlap of the two boundaries met at t : how much of what was prepared arrives at what was recorded. Everything below is evaluated at one such moment, so the argument is dropped and ψ_i and ψ_f stand for the states already carried there.

For a single ontic bit the weak value is converted into an ordinary probability by the Aharonov–Bergmann–Lebowitz rule.¹⁸ That rule is usually written with projectors. For an observable A with the two outcomes ± 1 , write Π_+ and Π_- for the projectors onto the eigenstates that return $+1$ and -1 . The rule gives

$$P(\pm) = \frac{|\psi_f^\dagger \Pi_\pm \psi_i|^2}{|\psi_f^\dagger \Pi_+ \psi_i|^2 + |\psi_f^\dagger \Pi_- \psi_i|^2}.$$

We want it in terms of the weak value instead, since that is the quantity the formalism hands us. Abbreviate $p = \psi_f^\dagger \Pi_+ \psi_i$ and $m = \psi_f^\dagger \Pi_- \psi_i$. Since $A = \Pi_+ - \Pi_-$ and $\Pi_+ + \Pi_- = I$, the weak value is the ratio

$$A_w = \frac{\psi_f^\dagger A \psi_i}{\psi_f^\dagger \psi_i} = \frac{p - m}{p + m}, \quad \text{so} \quad A_w + 1 = \frac{2p}{p + m}, \quad A_w - 1 = \frac{-2m}{p + m}.$$

Taking square magnitudes, the common factor $4/|p + m|^2$ cancels between numerator and denominator of any ratio built from them, and

$$\frac{|A_w + 1|^2 - |A_w - 1|^2}{|A_w + 1|^2 + |A_w - 1|^2} = \frac{|p|^2 - |m|^2}{|p|^2 + |m|^2} = P(+)-P(-) = \langle A \rangle.$$

¹⁷Y. Aharonov, D. Z. Albert and L. Vaidman, "How the result of a measurement of a component of the spin of a spin-1/2 particle can turn out to be 100," *Physical Review Letters* 60 (1988), 1351–1354.

¹⁸Y. Aharonov, P. G. Bergmann and J. L. Lebowitz, "Time Symmetry in the Quantum Process of Measurement," *Physical Review* 134 (1964), B1410–B1416.

So the weak value determines the ABL probabilities outright. For a single ontic bit, with $A = Z$ and the outcome -1 read as the bit value 1,

$$A_w = \frac{\psi_f^\dagger A \psi_i}{\psi_f^\dagger \psi_i}, \quad \langle A \rangle = \frac{|A_w + 1|^2 - |A_w - 1|^2}{|A_w + 1|^2 + |A_w - 1|^2}, \quad P(x=1) = \frac{1 - \langle A \rangle}{2}.$$

Apply this at the moment after the preparation but before the measurement bases are reoriented, taking $\psi_i = \psi_0$ and ψ_f the state the future measurement post-selects. The observables are the two ontic-basis values: $Z \otimes I$ for the first bit, $I \otimes Z$ for the second, so $P(x=1)$ is the probability that the bit's pre-measurement ontic value was 1.

Taking Z here is the choice of Section 10.4 rather than anything the rule dictates: the ontic basis is the computational basis, conventionally identified with the Pauli- Z basis, so these probabilities report beables, the values the bits actually carry between measurements. The derivation is tied to that basis for exactly that reason. Run it with A taken on some other basis and what comes back is an emergeable, a value the measurement would manufacture, and Section 10.1 refuted holding both sets of answers of one bit at the same time. Nothing here recovers all-bases-at-once realism by a retrocausal route.

Take the branch that matters, in which both bits are measured on the perturbed basis, and post-select each of the four outcomes in turn:

$\alpha\beta$	$P(A = 1)$	$P(B = 1)$
00	0	1/5
01	4/5	1/5
10	0	1
11	4/5	1

Read the row for the outcome $\alpha\beta = 10$. Both marginals are degenerate: the first bit's ontic value was 0 with certainty, the second's was 1 with certainty, so the pre-measurement configuration was $AB = 01$ (exactly the configuration the preparation alone declares impossible). The future choice to measure both bits on the perturbed basis, together with this outcome, has reached back and made the forbidden past actual.

Notice how different this is from the resolution of Section 11.4. There, the fact of the matter was fixed and the question was only *which* observer could read it; here the fact of the matter about the past is itself contingent on a choice made in the future. It is not that whoever measured first is right. Rather, the past holds a new possibility, licensed retroactively by what is done later.

The same method extends to the full four-bit circuit of Frauchiger and Renner, and reaches the same kind of verdict. Their setup uses four registers: a coin, the memory in which the first friend records the coin, a letter that friend writes to the second, and the memory in which the second friend records the letter, so that the string 0011 ought to be impossible. Read it left to right: if the coin is 0, the first friend's memory records 0; the letter is written by a controlled-Hadamard gate whose control is that memory, and a controlled-Hadamard leaves its target at 0 when its control is 0; so the letter stays 0, the second friend reads 0, and the string is 0000. A letter of 1 atop a coin of 0 would require the gate to flip its target from 0 to 1 with the control off, which it cannot do. Indeed, one verifies directly that the forward circuit gives only 0000, 1100, and 1111, each with probability 1/3, and never 0011.

Yet the same weak-value method, post-selected now on the future Bell-basis measurements of the two external observers, assigns the letter the ontic value 1 in the rare branch that produces Frauchiger and Renner’s contradiction (the forbidden 0011, with the letter carrying a 1 though its control was 0).

There is no inconsistency here, and the reason is a feature of their construction that is easy to miss: the first friend never actually *observes* the letter she writes. To call the controlled-Hadamard “writing a note” is misleading, because a note can be read back and checked; the gate is better pictured as a button on a sealed box that the friend presses and trusts to produce the letter, without ever looking at what came out. A retrocausal reading can then say that in this rare branch the box produces the letter in error (a 1 where a 0 was intended) and that this is not a violation of the gate’s law but an instance of it, because the gate’s behavior is contextual. We have already seen operators change their effect with context; the retrocausal reading simply allows the relevant context to lie in the future, so that the external observers’ later choice to measure on the Bell basis is what makes the controlled-Hadamard write the unexpected 1.

That calculation uses the ABL rule in a non-standard way, and the departure matters. The rule is applied to the two bits separately, through the degenerate projectors $Z \otimes I$ and $I \otimes Z$, so what comes back is a marginal probability for each bit on its own. Applied instead to the four one-dimensional projectors, which is its standard form, it returns the forward distribution conditioned on the recorded outcome by Bayes’ rule, and a conditioning of that kind cannot move probability onto a configuration the preparation assigns zero. The forbidden past is opened by the coarse-grained reading and closed by the fine-grained one.

Chapter 6 needs more than marginals. A transition matrix is fitted to a distribution over configurations, and marginals do not determine one. In this case they happen to, because when the paradox occurs they come out degenerate, and a degenerate marginal settles the joint: a first bit certainly 0 together with a second certainly 1 leaves only the configuration $AB = 01$. That is a large exception rather than the general rule. Most circuits, most of the time, return marginals nothing like degenerate, and for those the joint is simply missing. So this use of the rule does not amount to a complete retrocausal formulation of quantum information, though it may yet be usable to build one.

What such a formulation would have to supply can be said precisely. Section 5.1 showed the statistics to be non-Markovian: the law consults the past, and not only the past gates but the past distributions. Two Hadamards in a row return a bit prepared at 0 to 0, and return a bit prepared at 1 to 1, so the second gate carries Γ_{zero} in the first run and Γ_{one} in the second, though the distribution reaching it is uniform either way. A retrocausal construction needs that dependence running in both directions, with the gate’s behavior settled partly by the circuit’s final output. On two bits our formulation already allows one gate four behaviors, one for each configuration it might receive; a retrocausal formulation would allow four more, run by run, since a pair prepared at 00 can end on any of 00, 01, 10 or 11, and each ending would call for the gate to have behaved differently.

Weak values are already quantities of that kind, computed from the preparation and the recorded outcome together, so they belong to a single run rather than to an ensemble. The paradox then has a retrocausal reading if some rule holds that the controlled-Hadamard, in those runs that reach it holding 00 and end on $\alpha\beta = 10$, behaves so as to leave the pair at $AB = 01$ after all, violating the premise of induction.

The rule to reach for first is the one Chapter 6 already uses. There the two distributions are known and the matrix is not, and a few reasonable demands settle it: start from $|U|^2$, deform it no further than carrying one distribution to the other requires, and fit by iterative proportional fitting, which leaves the zeros alone. Here only the marginals are in hand, which is a weaker footing, though marginals can still serve as the constraints that do the deforming.

The harder obstruction is that deforming is not enough. For the retrocausal reading the controlled-Hadamard must take an input of 00 to an output of 01, which the gate cannot ordinarily do, and scaling $|U|^2$ preserves its zeros, so no scaling produces that transition. The zeros themselves have to move, and in three of this circuit's four branches something of the kind is unavoidable. One tentative repair, which may be no better than a hack, adds a small error term so the fitting can move them, weighted heavily enough that it stays with $|U|^2$ unless forced off.

Set against that weaker footing is a stronger demand. The purpose of going retrocausal at all is to be rid of non-locality entirely, reaching past Section 6.7's pseudo-locality to genuine locality, and a demand of that strength is a substantial constraint in its own right, enough to make up some of what the missing joint withholds. Expressing it as cleanly as the non-local law is expressed is the difficulty. Solutions exist, and they look contrived. We leave the question there.

It is tempting to think this horn escapes the preferred frame that the non-local horn forced on us: if the influence runs backward along time rather than sideways across space, perhaps a Lorentz-invariant ontology is back within reach. Whether it is depends on what sort of retrocausality is meant, and the two sorts on offer come apart exactly here.

Take first the version that keeps a dynamics: a state that still evolves from one moment to the next, with a contribution now propagating backward from a future boundary condition. That version looks to inherit the problem rather than escape it. The backward-evolving state is defined on a hypersurface of constant time, propagating from the future boundary back toward the present across a *slice* of spacetime, and space-like separated frames slice spacetime differently. Which future measurements have already been made, and therefore what past possibilities they have opened, depends on how the slicing is drawn. The past ontic state the formalism assigns is then, once again, not a frame-invariant fact: two observers foliating spacetime differently assign the pair different pasts, exactly as in Section 11.4 they assigned it different presents. On this reading retrocausality may do no more than move the non-locality from space to time, leaving the need to privilege a slicing where it stood.

The second version may avoid the demand altogether. Emily Adlam has argued that the dynamical picture, with influences traveling backward along a time axis, is the less coherent of the two options available, and that the better one is an *all-at-once* reading.¹⁹ On that reading the theory evolves nothing at all. Rather than carrying a state from one slice to the next in either direction, it assigns probabilities directly to entire histories: the whole four-dimensional pattern of events is constrained in one stroke, and no hypersurface does any work. The argument of the previous paragraph runs entirely on how a state gets carried from one hypersurface to another, so against a theory in which nothing is carried it never starts. There is no slicing to privilege, because there is no evolution across slices to make single-valued.

What that gives up is dynamics itself. A theory of this kind does not say how the world

¹⁹E. Adlam, "Two Roads to Retrocausality," arXiv:2201.12934 (2022).

develops from one moment to the next, because it does not describe development at all; it states a global condition that whole histories either meet or fail to meet. Everything we have built is a dynamics in precisely the sense being surrendered: Chapter 6's transition matrix carries a distribution from one time interval to the next, and Chapter 5's laws consult a history in order to produce the next step. We are not willing to give that up, so we do not take this road. But we cannot claim it is closed. A reader willing to give up time-evolution as a fundamental notion may have a real route to keeping Lorentz invariance, and we flag it as an open question rather than a settled one.

The same trade reappears, in a different form, in the doctrine of *superdeterminism*: the proposal that the correlations of quantum mechanics reflect correlations laid down in the initial conditions of the universe, between the systems measured and the settings that will measure them, so that the assumption of statistical independence between them fails.²⁰ Superdeterminism is often presented as a way to keep both locality and definite values: once the initial data are fixed, everything evolves forward, locally, with no influence reaching across space and no influence reaching back in time. We read it, sympathetically, as a form of retrocausality rather than an alternative to it.

For the initial conditions to produce the observed correlations, they must already be correlated with measurement settings that lie in their future. That dependence on the future is pushed all the way back to the beginning and frozen into the starting configuration, but it remains, and one advocate, Gerard 't Hooft, is admirably candid about it.²¹ The particles, he writes, “already carry information about the settings that will be used by Alice and Bob . . . non-local information about the future is already present in the ‘hidden’ ontological data.” Information about future measurement choices, present from the beginning of time in the ontological data: that is a dependence on the future written into the initial conditions, and 't Hooft's own phrase “non-local information about the future” says as much.

So superdeterminism inherits the choice rather than escaping it. Carrying a dependence on the future, it must take one of the two forms just distinguished, and of those only the all-at-once form looks promising for avoiding a preferred frame altogether.

Retrocausality is a logically consistent option: it satisfies our criterion for a realist model as fully as the non-local reading does, assigning definite values grounded in an ontic state. Where the non-local reading gives up locality in space, the retrocausal reading gives up the fixity of the past. What it settles on the question of the frame depends on which version is taken. Kept dynamical, it appears to buy nothing: its values are no more Lorentz invariant than the non-local reading's, and it looks to need a preferred foliation for the same reason. Taken all-at-once, it may escape the demand, but only by giving up dynamics, and with it the whole machinery of laws that move a state from one moment to the next.

We do not endorse the retrocausal horn (it adopts the non-local reading throughout), but the horn is real, and a reader who finds non-locality harder to accept than backward-in-time influence is entitled to take it. What no reader is entitled to do is imagine that either intuition can be kept in full. Every road here leaves the comfortable, classical, local-and-forward-in-time picture of the world behind. What we do claim is narrower than it may have seemed: the values are real, and for any theory that still moves them from one moment to the next, they

²⁰For a sympathetic overview, see S. Hossenfelder, “Superdeterminism: A Guide for the Perplexed,” arXiv:2010.01324 (2020).

²¹G. 't Hooft, *The Cellular Automaton Interpretation of Quantum Mechanics* (Cham: Springer, 2016), §3.8 and §5.2.

are not Lorentz invariant and a privileged slicing is required to make them single-valued. The one escape we know of is to stop moving them at all.

11.11 Both conclusions in one branch

Section 11.2 left something hanging, and now that the chapter’s machinery is built it can be settled. The contradiction there was a contradiction between two people. Alice, reading $\alpha = 1$, retrodicts that the pair was $AB = 11$; Bob, reading $\beta = 0$, retrodicts that it was $AB = 00$; both cannot be right. The Many Worlds interpretation, which Section 9.1 declined on other grounds, has a reply ready here. On that reading a measurement does not select one outcome from several. It branches the world, and every outcome occurs, each in a branch of its own. Two records contradict one another only if they are held together; file them in separate branches and nobody is ever inconsistent with themselves, so no contradiction was ever derived.

Many Worlds is usually presented, and almost always popularized, as the claim that every outcome happens: a photon meets a beam splitter, the world divides, and the same observer sees it reflected in one branch and transmitted in the other. On that telling each branch is an ordinary classical world with a settled past, and the only novelty is that there are many of them. This is how the reply is ordinarily made to Bell’s theorem, and to the paradox presented by Frauchiger and Renner as well. That paradox requires a single-outcome assumption, which forbids an agent from being certain of two incompatible things at once, and its authors note that many-worlds interpretations “manifestly violate” it.²² The measurement problem is then said to cost little to solve: drop the assumption, allow that more than one outcome occurs, and nothing further is required.

One extra bit disposes of the reply, by putting both records in the same place. What follows adds no new charge to the one Section 9.1 already made; it closes a repair. The picture of a branch as an ordinary world is what makes the interpretation attractive, since it appears to hand back everything Sections 9.1 and 9.2 said was lost, and what this section shows is that no branch can be such a world, because the construction below puts two incompatible records of one run inside one of them. The charge is *inconsistency*, leveled on the interpretation’s own terms and using only steps it already accepts.

Let Alice tell Bob what she read. Give Bob a second bit, a notepad, starting at 0, and copy Alice’s reading onto it with a CX in which her bit is the control and the notepad is the target. Sending a message is nothing more than that here: the value sitting on her bit is written onto a bit in his keeping. Appended to the preparation of Section 11.1 and the two perturbations of Case C, the experiment becomes a three-bit circuit:

Alice	$\mathbf{e}_0$	$R_Y(\theta)$	$\overline{CX}$	$\overline{CH}$	H	$\overline{CX}$	$\perp$
Bob	$\mathbf{e}_0$		CX	CH	H		$\perp$
Notepad	$\mathbf{e}_0$					CX	$\perp$

Nothing in it is new. The first three gate columns are the preparation that produced ψ_0 , the fourth is Case C (a perturbation on each of the two bits), and the fifth is the message. The notepad sits idle until that last column, and no gate ever acts on it and Bob’s bit together.

²²Frauchiger and Renner, “Quantum theory cannot consistently describe the use of itself,” Box 4 and “Theories that violate Assumption (S).”

We will write a three-bit configuration as $\alpha\beta r$, where r represents what is written to the notepad. Running the circuit puts probability $1/12$ on each of 000, 101 and 111, and $3/4$ on 010. The remaining four configurations are impossible.

Two features of that distribution have to be checked before anything is concluded from it, and each is a line of arithmetic.

The first is that the message was only a message. Marginalize the notepad away, adding together the configurations that agree on $\alpha\beta$ and differ on r , and what returns is $\mathbf{p}_C$, entry for entry, unaltered. The copy moved no probability anywhere. Alice told Bob something; she did not do anything to him, and in particular she did not disturb the statistics the rest of the chapter has been computing. This matters, because the contradiction being chased is a contradiction about those very statistics: had the copy shifted them, the exercise would prove nothing.

The second is that the message is faithful. The configuration 100 carries probability exactly zero, and so does 001, and so does every other configuration in which r differs from α . There is no branch anywhere in which Bob's notepad misreports what Alice read. Whatever stands written on it is what she actually had.

Now take the configuration $\alpha\beta r = 101$, which comes up in one run out of twelve, and read it one bit at a time. Bob's own bit gives $\beta = 0$, and his Case-B reasoning from Section 11.2 converts that into the claim that the pair was $AB = 00$. His notepad gives $r = 1$, which by the previous paragraph is exactly Alice's $\alpha = 1$, and her Case-A reasoning converts that into the claim that the pair was $AB = 11$. Both claims are now in Bob's possession. They concern one run of one experiment; they are carried by two bits of a single configuration; and they cannot both be true.

The branching reply has nothing left to do here. It works by holding incompatible records apart, and these two are not apart: they are two bits of the same configuration, occurring together with probability $1/12$. Nor can they be pushed apart by objecting that some step was illegitimate, because every step was a unitary acting on the joint system, the preparation and the perturbations and the copy alike, and Many Worlds has granted all of them in advance: granting them is exactly what it means to say that measurement is unitary and nothing collapses. So the interpretation is committed to a branch of its own universal wavefunction in which one agent holds two incompatible records of a single run. Frauchiger and Renner diagnosed Many Worlds as an interpretation that simply denies their single-outcome assumption, which sounds like something its advocates were always content to concede. The damage is worse than that. The inconsistency is not between branches, where branching could absorb it; it is inside one. The sharpening is the philosopher of physics Jeffrey Bub's.²³

Our own resolution, made back in Section 11.3, was the other one available: Alice's and Bob's retrodictions are not both sound, because the second mover's inference consults a bit the first mover has already disturbed. That resolution gives up locality. Many Worlds offers branching in its place, as a way of keeping both locality and the soundness of the inferences. Branching works by separating incompatible records into different branches, and the notepad puts both of them inside one. The reply therefore fails, and the inconsistency stands.

Bub takes the result further, concluding that quantum mechanics cannot be read as a representational theory at all, and he defines an interpretation broadly enough to cover ours:

²³J. Bub, "Two Dogmas' Redux," arXiv:1907.06240 (2019).

a proposal to reformulate the theory in representational terms, “either by introducing hidden variables, or by proposing that every possible outcome occurs in a measurement, or in some other way.” The argument he gives is against the Everett interpretation, and the rest are swept in by that definition. It is not applicable to an interpretation that takes the first resolution, since the contradiction needs both retrodictions to be sound and Section 11.3 denies that they are.

For textbook quantum mechanics we agree with him, by another route. Section 11.9 established two things at once: the laws are non-local, and the statistics carry no signal. So the non-locality never surfaces in what a measurement reports. Each outcome is a local record that every frame agrees on, and nothing done elsewhere moves its statistics, while Section 11.7 showed that the configurations underneath cannot be Lorentz invariant at all. A theory confined to measurement outcomes is therefore Lorentz invariant, while one that says what those measurements are measurements of is not. That is why the measurement postulate has to be taken as primitive: in a theory of measurement outcomes alone there is nothing underneath for it to follow from. The Frauchiger–Renner setup demonstrates the same point from the other side, since avoiding the contradiction requires an absolute ordering of the perturbations, and an ordering is exactly what a theory of measurement outcomes alone declines to supply.

This result and Drezet’s go together. Section 11.9 reported his finding that the non-locality survives inside a single branch, because the GHZ argument needs nothing beyond the outcome each observer reports there. What this section adds is that the contradiction survives there too, and for the same reason, since the contradiction is what a non-local law produces once two retrodictions about one past are collected in one place. Branching separates neither.

An Everettian may say that no configuration of the pair was ever real inside the branch, so the two retrodictions have nothing to disagree about. That reply is available, and with branching closed off it is the only one left. Its cost is the whole of the popular picture. A branch in which no configuration of the pair was ever real is not a world in which the photon was reflected, since nothing in it did anything at all. The consequences were set out in Section 9.2: a theory whose records imply no past is the theory Bell called a radical solipsism, in which every memory is an account of a history that never actually occurred. That is a kind of *Last Thursdayism*, a world arriving already furnished with records of a past it never had.

So the popular picture is unavailable on either reading. Grant the branch a definite configuration for the pair and it holds a contradiction; deny it one and it holds no past at all. Neither is an ordinary world, and neither lets the interpretation say that the photon was reflected here and transmitted there, since no branch holds a settled account of even one run. Whatever Many Worlds delivers, it is something other than what its popular telling promises.

What that telling promised was something anyone would want. A multiverse is extravagant, but what the extravagance appeared to yield was a reading anyone could picture: ordinary worlds, one outcome in each, related only by having branched from a common origin. Nothing in it requires a collapse, a wave that steers the particle, or facts that hold only relative to an observer. That is why it is the interpretation most often reached for by those who want quantum mechanics to make ordinary sense, and the contradiction assembled above removes precisely that. What survives is no easier to picture than the readings it was meant to improve upon.

Section 8.4 is where this was set up. A wave that will not fit into three-dimensional space

leaves one route open to anyone who wants it real: stop putting the wave in the world and put the world inside the wave, as a pattern in a space of M^N amplitudes. The move has a requirement attached. If the world we experience sits inside the wave, there has to be a consistent way of picking it out again, and the attempt is what this section has just made. Giving a branch a definite history is what produced the contradiction.

What is left is the other possibility: that the world as we know it, furnished with a definite past, never existed to begin with. That is what Bell called a radical solipsism, reached here from the opposite direction: the whole history of the world, our own lives included, an illusion owed to impressions left on our memories. The wave could not be placed in the world, so the world was placed in the wave, and the world did not survive the move.

11.12 Relational quantum mechanics

The argument bears on one further position. Relational quantum mechanics, developed by Carlo Rovelli, holds that the values of physical quantities belong to a system only relative to another system that has interacted with it.

The relational reading keeps locality. Smerlak and Rovelli argue that the correlations between distant bits can be accounted for without giving it up, that the apparent non-locality is an illusion, and that locality is saved by weakening realism instead.²⁴ A state of the pair that is definite for everyone at once would be a description given from no particular place, and giving it would require an observer standing outside every interaction in the experiment. No such standpoint exists: two observers can compare what they hold only by interacting with each other, so there is nowhere that Alice's record and Bob's record are both definite and set side by side.²⁵ *Juxtapose* the two perspectives, in Rovelli's word, and the correlations look non-local. Decline to juxtapose them, and each observer's own account is local from beginning to end.

They state explicitly that they are abandoning what they call *Einstein's realism*, the assumption that "there exists a physical reality independent of substantiation and perception"; what survives is a realism with no single account of events that every observer shares.²⁶ The picture that results is a fragmented one. Rovelli speaks of the *fragmentation of points of view* that follows from properties being only relative: the facts do not assemble into one world, but into as many partial accounts as there are systems to hold them. Call each such account a *fragment*. Consistency is then required within a fragment and nowhere else. Restricting it that way is what makes the locality claim work, because two fragments that disagree are never in contradiction: a contradiction needs both claims to be facts in one place, and no fragment holds both.

The previous section assembles a contradiction inside a single fragment. Bob holds two bits, his own $\beta = 0$ and his notepad's $r = 1$, in one configuration of one run. From the first he infers that the pair was $AB = 00$; from the second, which is Alice's reading, that it was $AB = 11$. No second observer is required and no two perspectives are juxtaposed. One observer, one fragment, two records that cannot both be right.

²⁴M. Smerlak and C. Rovelli, "Relational EPR," *Foundations of Physics* 37 (2007), 427–445, §1.

²⁵C. Rovelli, "Relational quantum mechanics," *International Journal of Theoretical Physics* 35 (1996), 1637–1678, §3.3.

²⁶Smerlak and Rovelli, "Relational EPR," §2.

Relational quantum mechanics answers this by denying that the two claims are about anything real. Physical reality, on that view, is formed by the events through which interacting systems affect one another: quantum events “exist only in interactions,” and which properties a system has is only relative to another system that interacts with it.²⁷ Bob never interacted with the pair at the time his two inferences are about. The perturbations of Case C destroyed those values before he read anything, so relative to Bob neither $AB = 00$ nor $AB = 11$ was ever a fact at all, and a contradiction between two claims that have no relational existence for the one holding them is a contradiction about nothing.

Consider what that reply commits one to. Replace the bits with brain-bits, as in Section 11.7, and the records stop being about particles. They are about what an observer experienced, and the sharpest case is the one Section 9.3 set out. The friend sealed in the laboratory sees a definite outcome. Wigner, waiting outside it, describes her as having no definite outcome at all. On the relational reading both descriptions stand, each relative to the system holding it, so the friend as she experiences herself and the friend as Wigner describes her are not the same thing and are allowed to differ.

A simpler case shows where the difference comes from. Prepare a Bell state, whose two bits always match, and suppose neither carries a value before it is measured. Again, write A for the outcome on Alice’s bit and B for the outcome on Bob’s. Alice measures hers and it comes out $A = 0$. Since the two always match, she infers at once that the pair is $AB = 00$, so Bob’s bit must read 0 as well. Either it acquired that value at the moment she measured, which is the non-locality at issue, or it acquires one later and on its own, in which case nothing ties it to hers and the two will sometimes fail to match. Plain statistics leaves no third option, since events that settle independently do not come out correlated every time.

Relational quantum mechanics accepts that independence. Relative to Bob, his own reading is the only one there is, so nothing obliges it to agree with what Alice holds. Smerlak and Rovelli grant the point in those terms: Alice and Bob are space-like separated, so “there cannot exist an observer with respect to which both of these outcomes are actual,” comparing them is “meaningless,” and each outcome is “fully independent” of the other.²⁸

Fragmentation is what follows. Alice reading $A = 0$ makes the pair $AB = 00$ relative to her. Bob’s outcome is independent of hers, so in half the runs he finds $B = 1$, and in those the pair is $AB = 11$ relative to him. Both fragments have the bits matching, just as the preparation requires, so neither observer finds anything out of order in the records available to them. The two accounts nonetheless disagree about both bits.

In such a run, nothing Bob can do will bring the disagreement into view. His fragment holds that $A = 1$, while Alice’s own reading of it is $A = 0$. Should he ask her what she saw, what answers him is Alice relative to Bob, and she reports $A = 1$, the value his fragment already carried. Each observer hears the others confirm the account they already hold, which Smerlak and Rovelli offer as a sound definition of objectivity: everyone “hears everybody else stating that they see the same elephant they see.”²⁹ The two accounts, they say, “cannot and should not be juxtaposed.”³⁰

One might object that the argument has already broken the rule. If the two accounts cannot be juxtaposed, then setting them beside each other as the last few paragraphs did

²⁷Smerlak and Rovelli, “Relational EPR,” §2.

²⁸Smerlak and Rovelli, “Relational EPR,” §4.1.

²⁹Smerlak and Rovelli, “Relational EPR,” §4.3.

³⁰Smerlak and Rovelli, “Relational EPR,” §4.4.

is the very move the view forbids, and the fragmentation was assembled from a standpoint nobody occupies.

Their own statement of the rule does the same thing. Reality, they write, “admits one description per (observing) system, each being internally consistent.”³¹ That sentence ranges over every observing system at once and says what holds of each: one description apiece, internally consistent. It is a claim about the whole collection, made from no system’s vantage in particular. The rule bars requiring the members of that collection to agree with one another; it does not bar speaking of them together. Consistency is demanded within a description and nowhere else, and the argument above asks for nothing more.

What the argument requires is weaker than what the prohibition rules out. It requires only that something exist relative to Bob. It does not require that anyone hold Bob’s value and Alice’s at once, and the fragmentation above was never derived by anyone holding both. That weaker thing the relational view has to grant. Suppose it refused, and Alice were permitted to admit nothing beyond what stands in relation to Alice. She would then be denying that Bob has a perspective at all: no experience of the world belonging to him, and none belonging to anyone else either, with her own taken as the only true one. That is solipsism, and calling the facts relative does not change it. The interest of the view lies in denying exactly that, and its authors say so in the same passage: the scheme is offered as sufficient for describing nature and for our exchanging information about it, “hence circumventing solipsism.”³² So it must allow Bob’s perspective contents of its own.

Granting that much does not force those contents to agree with hers. The relational view can say, and does say, that the two may differ and are not to be set side by side, and that is the stance it takes in order to keep clear of solipsism while holding perspectives apart. The Bell pair shows what the stance costs. Once both accounts are granted to exist, and nothing obliges them to match, they differ; and the fragmentation follows from those two admissions alone, without anyone ever standing where both values are held together.

Quantum mechanics still predicts that the two values match, so Alice can travel to Bob and measure his bit herself, or ask Bob what he measured and so measure it indirectly. Either way the two come out the same. The journey from Alice to Bob is a local one, and it carries the record of her own measurement with it, whether in a note she wrote or in the memory she holds of it. When the interaction with Bob’s bit takes place relative to Alice, everything needed to guarantee the match is present at that one place. The correlation is therefore fixed by a local process, with nothing acting across the gap.

The consequence of that locality is the freedom of one fragment to disagree with another, applied to a person rather than a bit, as in the case of Wigner and his friend. What he finds when he looks at her and what she experiences of herself are therefore allowed to differ. Once they can differ, each observer is sealed inside an account of the world in which the other observers are not the ones having experiences of their own. The ones who do are elsewhere, in fragments of their own, and no interaction ever reaches them, since every interaction yields one more fact relative to whoever performed it. Wigner can learn any number of facts about the friend relative to Wigner, and never one fact about the friend relative to herself, since the two are not the same thing. However much he learns about his friend, he can never truly know her.

³¹Smerlak and Rovelli, “Relational EPR,” §4.4.

³²Smerlak and Rovelli, “Relational EPR,” §4.4.

Adlam has argued that a science is confirmed by results that more than one person can hold, and that an observer who can never obtain anyone else's result has no evidence for quantum mechanics beyond his own present records, which leaves the view without the support that recommended it.³³ Science is a social process, and a social process requires the exchange of information between individuals. The relational reading leaves each of them sealed off, and thus undermines the very foundation by which it came to be justified.

Adlam and Rovelli have since proposed a repair. They replace one of the relational postulates with a principle called *cross-perspective links*: if Alice measures a variable, and nothing subsequently destroys the record of it held in her physical variables, her memory of the measurement included, then when Bob measures that record he obtains what she obtained.³⁴ Their own treatment of Frauchiger and Renner depends upon that proviso. The measurement that would expose the disagreement is also the one that erases the record disagreeing with it, so the mistaken observer is left with no memory of the mistake.

The concession is significant. Quantum events, the authors write, “must now be regarded as observer-independent in some sense, although quantum states remain relational.”³⁵ That is the one thing the relational view had done without. Blake Stacey has put the objection more sharply, that the revision backs away from the relationalism which had characterized the interpretation: under the new postulate an observer's choice of measurement, her outcome and her later interactions are all, in principle, facts for everyone, which leaves the position hard to distinguish from the orthodoxies it was meant to improve on.³⁶

Additionally, the proviso in the postulate is itself a condition on the future. Alice measures and obtains a value; whether Bob will find the same value when he looks, and so whether a cross-perspective link obtains between them at all, depends on whether anything afterwards destroys her record. Sharing does not settle it either. Once Alice has told Bob, a later measurement on the systems holding both records can destroy them together, which is exactly what the external observers do in Frauchiger and Renner's argument. What decides between a shared fact and a fact for Alice alone is something that has not happened yet.

The authors do not resist this reading. They answer that the postulate requires only that the two values match, with neither causing the other, so the events “can simply be generated ‘all-at-once’ in a time-symmetric and atemporal fashion,” and their conclusion adopts an ontology in which the distribution of events is “determined all-at-once in an external, time-symmetric way.”³⁷ Nothing is sent backward along a time axis. What settles the matter is a consistency condition laid across the whole of history at once, with no regard for the order in which its parts occur.

That is the all-at-once picture of Section 11.10, which Adlam counts as one of the two roads to retrocausality and the more coherent of them. So the repaired theory arrives at retrocausality after all. But Section 11.10 showed retrocausality resolving the paradox of the two bits by itself, by giving up induction rather than inference, with no appeal to relative

³³E. Adlam, “Does Science need Intersubjectivity? The Problem of Confirmation in Orthodox Interpretations of Quantum Mechanics,” *Synthese* 200 (2022), 522.

³⁴E. Adlam and C. Rovelli, “Information is Physical: Cross-Perspective Links in Relational Quantum Mechanics,” arXiv:2203.13342 (2022), Definition 4.1; the Frauchiger–Renner discussion is §7.

³⁵Adlam and Rovelli, “Information is Physical: Cross-Perspective Links in Relational Quantum Mechanics,” §8.

³⁶B. C. Stacey, “The De-Relationalizing of Relational Quantum Mechanics,” arXiv:2211.03230 (2022).

³⁷Adlam and Rovelli, “Information is Physical: Cross-Perspective Links in Relational Quantum Mechanics,” §§6 and 8.

facts anywhere in the argument. Whoever is willing to go that far has no further use for the relational machinery: the paradox is already answered, and relativizing facts to observers adds a second mechanism where one had sufficed.

None of this shows the relational interpretation to be inconsistent. What the objections bear on is what adopting it commits one to. Weakening realism asks little while the systems are individual particles: saying that a bit has a value only with respect to whatever interacted with it takes away nothing anyone wanted, and a good deal of the measurement problem dissolves along the way. The difficulty arrives when the theory is applied to observers as well, which is what treating it as universal requires. What gets relativized is then somebody's experience, the fragments are people, and the account grows more convoluted rather than simpler. The initial elegance of its supposed simplicity then seems to dissolve away.

Code: this chapter in Octave

```

% --- the gates, and the preparation of Section 11.1 ---
I2 = eye(2); H = (1/sqrt(2))*[1 1; 1 -1]; sq = 1/sqrt(2);
th = 2*acos(1/sqrt(3));
RY = [cos(th/2) -sin(th/2); sin(th/2) cos(th/2)];
CX = [1 0 0 0; 0 1 0 0; 0 0 0 1; 0 0 1 0];
CH = [1 0 0 0; 0 1 0 0; 0 0 sq sq; 0 0 sq -sq];

% Preparation: RY on Alice's bit, CX, CH
psi0 = CH * CX * kron(RY,I2) * [1;0;0;0]; % (1/sqrt(3))[1 0 1 -1]
p0 = abs(psi0).^2 % [1/3 0 1/3 1/3]: 01 impossible
pA = abs(kron(H,I2)*psi0).^2 % [2/3 1/6 0 1/6]: aB=10 impossible
pB = abs(kron(I2,H)*psi0).^2 % [1/6 1/6 0 2/3]: Ab=10 impossible
pC = abs(kron(H,H)*psi0).^2 % [1/12 3/4 1/12 1/12]: ab=10 at 1/12

% --- No local matrix on Alice's bit alone can effect p0 -> pA ---
% Four equations, four unknowns. Nothing is assumed about Gamma: the
% system alone fixes it, and the entries it fixes are not probabilities.
g01 = 3*pA(2); g11 = 3*pA(4); g00 = 3*pA(1) - g01; g10 = 3*pA(3) - g11;
G = [g00 g01; g10 g11] % [1.5 0.5; -0.5 0.5]
norm(kron(G,I2)*p0 - pA) % 0: it really is the solution
sum(G) % [1 1]: normalization is not the failure
[min(G(:)) max(G(:))] % -0.5 and 1.5: outside [0,1]

% --- IPF single-step matrices for the two orderings ---
function G = ipf(A, pt, pt1)
    v = ones(length(pt),1);
    for k = 1:20000
        u = pt1 ./ (A*(v.*pt)); u(~isfinite(u)) = 0;
        v = 1 ./ (A'*u); v(~isfinite(v)) = 0;
    end
    G = diag(u)*A*diag(v);
end
dA = abs(kron(H,I2)).^2; dB = abs(kron(I2,H)).^2;

GA1 = ipf(dA, p0, pA), GB1 = ipf(dB, pA, pC) % Version 1: Alice first
GB2 = ipf(dB, p0, pB), GA2 = ipf(dA, pB, pC) % Version 2: Bob first

% joint over (final config, initial config): J = G2*G1*diag(p0)
J1 = GB1*GA1*diag(p0); J2 = GA2*GB2*diag(p0);
% alpha=1: final first bit 1 -> rows 3:4
% beta=0: final second bit 0 -> rows [1 3]
V1_alice = sum(J1(3:4,4)) / sum(sum(J1(3:4,:))) % Pr(AB=11|a=1) -> 1
V1_bob = sum(J1([1 3],1)) / sum(sum(J1([1 3],:))) % Pr(AB=00|b=0) -> 0.2
V2_alice = sum(J2(3:4,4)) / sum(sum(J2(3:4,:))) % Pr(AB=11|a=1) -> 0.2
V2_bob = sum(J2([1 3],1)) / sum(sum(J2([1 3],:))) % Pr(AB=00|b=0) -> 1

```

```

% both orderings give the same observable statistics:
[GB1*GA1*p0, GA2*GB2*p0, pC] % three equal columns

% --- Section 11.5: the EPR pair with two staggered H gates ---
bell = CX * kron(H,I2) * [1;0;0;0]; % (1/sqrt(2))[1 0 0 1]
q0 = abs(bell).^2 % [.5 0 0 .5]
q1 = abs(kron(H,I2)*bell).^2 % [.25 .25 .25 .25]: Alice's gate
q2 = abs(kron(I2,H)*kron(H,I2)*bell).^2 % [.5 0 0 .5]: |H(x)H| restores it
norm(abs(kron(H,I2)*kron(I2,H)*bell).^2 - q2) % 0: order-independent

% Alice's step needs nothing non-local: a fair coin on her own bit
Gc = [.5 .5; .5 .5]; norm(kron(Gc,I2)*q0 - q1) % 0: exact

% Bob's step, by Section 11.3's method. (I kron G) q1 = q2 asks the
% same sum of unknowns to take two different values, so no G exists.
[4*q2(1), 4*q2(3)] % gamma00+gamma01 must be 2 and must be 0
[4*q2(2), 4*q2(4)] % gamma10+gamma11 must be 0 and must be 2

% --- Section 11.5: the strong value model ---
% A strong object is s = [1; x; y; z] with entries +-1 and y = -x*z,
% so one coin apiece. Gates act through their Pauli transfer matrix,
% R(j,k) = Tr(P_j U P_k U')/2^n, built here over the Pauli strings.
function P = paulis(n)
    B = {eye(2), [0 1;1 0], [0 -1i;1i 0], [1 0;0 -1]};
    if n == 1, P = B; return; end
    Q = paulis(n-1); P = {};
    for j = 1:4, for k = 1:length(Q), P{end+1} = kron(B{j},Q{k}); end, end
end
function R = ptm(U, n)
    P = paulis(n); d = 2^n; m = 4^n; V = zeros(m, d*d);
    for j = 1:m, V(j,:) = reshape(P{j}.', 1, []); end
    R = zeros(m);
    for k = 1:m, R(:,k) = real(V * reshape(U*P{k}*U', [], 1)) / d; end
    R(abs(R) < 1e-9) = 0;
end
coin = @() 2*(rand < 0.5) - 1;
obj = @(x) [1; x; -x; 1]; % z fixed at +1, y = -x*z
RH = ptm(H,1) % exchanges x and z, flips y
all(abs(sum(abs(RH)) - 1) < 1e-9) % 1: a signed permutation

% Three objects through the GHZ circuit, factored after every CX.
RCX = ptm(CX,2); mar = @(w) {w([1 5 9 13]), w([1 2 3 4])};
runs = {[2 3 3],[3 2 3],[3 3 2],[2 2 2]}; % X is entry 2, Y is 3
acc = zeros(1,4); N = 2000;
for t = 1:N

```

```

q = {obj(coin()), obj(coin()), obj(coin())};
q{1} = RH*q{1};
d = mar(RCX*kron(q{1},q{2})); q{1} = d{1}; q{2} = d{2};
d = mar(RCX*kron(q{2},q{3})); q{2} = d{1}; q{3} = d{2};
for r = 1:4
    k = runs{r}; acc(r) += q{1}(k(1))*q{2}(k(2))*q{3}(k(3));
end
end
acc/N          % [+1 -1 -1 +1]; quantum is [-1 -1 -1 +1], so XYY is wrong

% Without the factoring the combined vector IS the quantum state,
% written in the Pauli basis: no longer a model of anything local.
G3 = kron(I2,CX)*kron(CX,I2)*kron(H,kron(I2,I2));
ghz = G3*[1;0;0;0;0;0;0;0]; P3 = paulis(3); qv = zeros(64,1);
for j = 1:64, qv(j) = real(trace(P3{j}*(ghz*ghz'))); end
s0 = [1;0;0;1]; % |0> as a Pauli vector
norm(ptm(G3,3)*kron(s0,kron(s0,s0)) - qv) % 3e-16

% The simplest local hidden memory for this experiment: each bit
% carries one entry per basis, all drawn at the preparation.
R = {I2, H, H*[1 0; 0 -1i]}; % Z, X, Y recast onto the ontic basis
err = zeros(3);
for i=1:3, for j=1:3
    P = zeros(4,1);
    for z=0:1, for x=0:1, for y=0:1
        A = [z x y](i); B = [z x 1-y](j); % Alice's row, Bob's row
        P(A*2+B+1) += 0.125;
    end, end, end
    err(i,j) = norm(P - abs(kron(R{i},R{j})*bell).^2);
end, end
max(max(err)) % 3e-16: all nine setting pairs reproduced

% --- Section 11.6: Bell's 1964 theorem, on the singlet ---
e1 = [0;1];
Rot = @(t) [cos(t/2) -sin(t/2); sin(t/2) cos(t/2)]; % R_Y(t)
sing = CX * kron(H,I2) * kron(e1,e1) % (1/sqrt(2))[0 1 -1 0]
% measuring at angle t = recast by R_Y(-t), then read the ontic basis
E = @(a,b) [1 -1 -1 1] * abs(kron(Rot(-a),Rot(-b))*sing).^2;
[E(0.7,0.7), E(0,pi/3), -cos(0-pi/3)] % -1, and E(a,b) = -cos(a-b)

% Bell's inequality |E(a,b)-E(a,c)| <= 1 + E(b,c), axes 0, 60, 120 deg
a = 0; b = pi/3; c = 2*pi/3;
lhs = abs(E(a,b) - E(a,c)), rhs = 1 + E(b,c) % 1 and 0.5: broken

% One setting-pair alone IS locally explicable (unlike Section 11.3):
pA1 = abs(kron(Rot(-b),I2)*sing).^2;

```

```

Gloc = [pA1(2) pA1(1); pA1(4) pA1(3)] / 0.5 % legal stochastic matrix
norm(kron(Gloc,I2)*abs(sing).^2 - pA1) % -> 0: transports exactly

% --- Eberhard: a gate on bit 1 cannot move bit 2's marginals ---
psi = randn(4,1)+1i*randn(4,1); psi = psi/norm(psi);
t = randn(); ph = randn();
U = [cos(t)*exp(1i*ph) -sin(t); sin(t) cos(t)*exp(-1i*ph)];
m = @(v) [v(1)+v(3); v(2)+v(4)];
[m(abs(psi).^2), m(abs(kron(U,I2)*psi).^2)] % identical columns

% --- Retrocausality: TSVF weak values + the ABL rule ---
% psi0 = prepared state (post-CH, pre-final-H), the forward-evolving state
Z = [1 0; 0 -1]; ZI = kron(Z,I2); IZ = kron(I2,Z);
wv = @(A,pf,pi) (pf'*A*pi) / (pf'*pi); % TSVF weak value
abl = @(a) (abs(a+1)^2 - abs(a-1)^2)/(abs(a+1)^2 + abs(a-1)^2);
pr1 = @(A,pf,pi) (1 - abl(wv(A,pf,pi)))/2; % Pr(bit = 1)
for b = 0:3
    pf = kron(H,H) * ((1:4 == b+1)'); % post-selected outcome
    disp([pr1(ZI,pf,psi0), pr1(IZ,pf,psi0)]) % [first bit, second bit]
end
% outcomes 00,01,10,11 -> [0 .2],[.8 .2],[0 1],[.8 1]
% outcome 10 -> [0, 1]: the past ontic pair was 01, the forbidden one

% --- Section 11.11: both conclusions land in one branch ---
psi3 = kron(kron(H,H)*psi0, [1;0]); % Case C, then Bob's notepad bit = 0
CXam = eye(8); % CX: Alice's bit -> notepad
CXam([5 6],[5 6]) = [0 1;1 0]; CXam([7 8],[7 8]) = [0 1;1 0];
p3 = abs(CXam*psi3).^2 % 1/12 on 000,101,111; 3/4 on 010
marg = [p3(1)+p3(2); p3(3)+p3(4); p3(5)+p3(6); p3(7)+p3(8)];
norm(marg - pC) % 0: the copy moved no probability

```

Chapter 12

The PBR Theorem

One theorem remains, and it is widely reputed to be the strongest. In 2012 the physicists Matthew Pusey, Jonathan Barrett, and Terry Rudolph published a theorem, now known by their initials, under the title “On the reality of the quantum state.”¹ Its advertised conclusion is aimed directly at us: that the quantum state cannot be a mere statistical description. More precisely, it holds that “any model in which a quantum state represents mere information about an underlying physical state of the system, and in which systems that are prepared independently have independent physical states, must make predictions which contradict those of quantum theory.” If that sentence ended at the first comma, Part II would be refuted. Our work here is to see that everything depends upon what comes after the second comma, and even more so on an assumption buried deeper still.

The original proof is compact but its assumptions are easy to miss. We will work instead with a simplified two-bit version,² which needs nothing beyond the machinery already in hand.

12.1 Four preparations

Prepare two bits, independently of one another, in one of four ways: leave both bits alone, apply a Hadamard to the first, apply a Hadamard to the second, or apply a Hadamard to both. Statistically: each bit is either left in a *degenerate* distribution (D), which is certainly 0, or placed in a *uniform* distribution (U). The four preparations and their Born distributions:

$$\begin{array}{cc}
 \begin{array}{c} \mathbf{e}_0 \\ \mathbf{e}_0 \end{array} & \mathbf{p}_{DD} = \begin{pmatrix} 1 \\ 0 \\ 0 \\ 0 \end{pmatrix} & \begin{array}{c} \mathbf{e}_0 \text{ H} \\ \mathbf{e}_0 \end{array} & \mathbf{p}_{UD} = \begin{pmatrix} 0.5 \\ 0 \\ 0.5 \\ 0 \end{pmatrix} \\
 \begin{array}{c} \mathbf{e}_0 \\ \mathbf{e}_0 \text{ H} \end{array} & \mathbf{p}_{DU} = \begin{pmatrix} 0.5 \\ 0.5 \\ 0 \\ 0 \end{pmatrix} & \begin{array}{c} \mathbf{e}_0 \text{ H} \\ \mathbf{e}_0 \text{ H} \end{array} & \mathbf{p}_{UU} = \begin{pmatrix} 0.25 \\ 0.25 \\ 0.25 \\ 0.25 \end{pmatrix}
 \end{array}$$

¹M. F. Pusey, J. Barrett, and T. Rudolph, “On the reality of the quantum state,” *Nature Physics* 8 (2012), 475–478.

²The author first encountered this simplified presentation in a lecture by the physicist Robert Spekkens: R. W. Spekkens, “Why I Am Not a Psi-ontologist,” Perimeter Institute, PIRSA:12050021 (2012), <https://pirsa.org/12050021>.

We have devoted eleven chapters to defending the position that quantum information really is statistical. If so, then in each individual run the pair of bits really occupies one definite configuration, drawn from the preparation's distribution. And the distributions *overlap*. The configuration 00 can occur under all four preparations; 01 under DU and UU; 10 under UD and UU; only 11 is unique to UU. A pair of bits sitting in the configuration 00 carries no marking that says which preparation produced it.

It follows that no device receiving the bits can reliably identify the preparation. That is a limitation quantum mechanics itself respects, since the four quantum states are not mutually orthogonal. Nor can a device learn the preparation from long-run frequencies if we choose the preparation afresh, at random, on every run: each run hands the device one pair of bits, and the statistics across runs are a mixture from which the individual run's label cannot be recovered.

12.2 The black box

The surprise is what a single fixed gate can do with these four preparations. Send the two bits, whichever way they were prepared, into a black box that applies one fixed unitary (the same gate every time, with no knowledge of the preparation) and then measures both bits:

$$U = \begin{pmatrix} 0 & \frac{1}{\sqrt{2}} & \frac{1}{\sqrt{2}} & 0 \\ \frac{1}{2} & -\frac{1}{2} & \frac{1}{2} & \frac{1}{2} \\ \frac{1}{2} & \frac{1}{2} & -\frac{1}{2} & \frac{1}{2} \\ \frac{1}{\sqrt{2}} & 0 & 0 & -\frac{1}{\sqrt{2}} \end{pmatrix}.$$

(The reader can check $UU^\dagger = I$; in fact U is real, so it is an orthogonal matrix, with $\det U = -1$.) Nothing exotic is being smuggled in: the box can be built from the gates already in hand, as

$$\begin{array}{cccccccc} \mathbf{e}_0 & R_Y(-\pi/2) & \overline{CX} & R_Y(-\pi/2) & SW & G(\pi/4) & \overline{CX} & \perp \\ \mathbf{e}_0 & & CX & & SW & \overline{G}(\pi/4) & YXYH & CX \quad HZ \quad \perp \end{array}$$

where $G(\pi/4)$ is Chapter 8's Givens operator, itself built there from SWAP, CX and R_Y . The circuit is long and none of its detail matters below; only the total operator does. The chapter's code checks that it multiplies out to U exactly. Apply U to each of the four preparations and take the Born rule of each result:

$$\mathbf{p}'_{DD} = \begin{pmatrix} 0 \\ 0.25 \\ 0.25 \\ 0.5 \end{pmatrix} \quad \mathbf{p}'_{UD} = \begin{pmatrix} 0.25 \\ 0.5 \\ 0 \\ 0.25 \end{pmatrix} \quad \mathbf{p}'_{DU} = \begin{pmatrix} 0.25 \\ 0 \\ 0.5 \\ 0.25 \end{pmatrix} \quad \mathbf{p}'_{UU} = \begin{pmatrix} 0.5 \\ 0.25 \\ 0.25 \\ 0 \end{pmatrix}.$$

Look at where the zeros fall. Under the preparation DD, the outcome 00 *never* occurs. Under UD, the outcome 10 never occurs; under DU, never 01; under UU, never 11. Encode each preparation by which bits were uniform ($D \rightarrow 0$, $U \rightarrow 1$) and the pattern becomes crisp: *outcome k never occurs under preparation k* . The box cannot say which preparation it received. Yet every single outcome it produces *veto*es one of the four candidates, with certainty, on every run.

12.3 Why no preparation-blind box can do this

Suppose the box were what a Markovian statistical theory says a box must be: a fixed stochastic matrix Γ , its response determined by the configuration of the two bits it receives, so that the same column is applied to the configuration 00 no matter how the 00 came about. The veto behavior is then impossible, and the proof takes four lines.

Consider the column of Γ for the input configuration 00, and chase it through the four preparations, each of which can deliver a 00:

- Under DD the input is always 00, and outcome 00 never occurs: $\Gamma_{00,00} = 0$.
- Under DU the inputs are 00 and 01, and outcome 01 never occurs; since both inputs arise with positive probability, both columns must give 01 nothing: $\Gamma_{01,00} = 0$.
- Under UD, by the same reasoning with outcome 10: $\Gamma_{10,00} = 0$.
- Under UU all four inputs arise, and outcome 11 never occurs: $\Gamma_{11,00} = 0$.

Every entry of the 00 column is zero, and a column of zeros cannot sum to one. There is no such stochastic matrix. A box whose response is a function of the incoming bit values alone cannot produce the veto behavior that quantum mechanics predicts and the laboratory delivers.

This is the PBR theorem, in miniature and in full logical strength. The standard conclusion is drawn immediately: the bits must carry something *beyond* their configuration: a hidden memory of their preparation. This memory is none other than the quantum state, which is conveniently distinct in the four cases. The quantum state must therefore be physically real, “ontic” in the vocabulary of Harrigan and Spekkens, and every merely statistical, “epistemic” reading is refuted.³ So the advertisement runs.

12.4 The second premise

The proof is correct; we endorse every line of it. What the advertisement omits is that the proof consumed more than the premise on trial. That premise is that the bits have definite configurations distributed statistically. On top of it sit two independence assumptions. One is stated in the theorem’s own abstract: “systems that are prepared independently have independent physical states.” The other is buried inside the definition of the box as a fixed stochastic matrix: the box’s response depends on the present physical state of the bits *and on nothing else*.

Pusey, Barrett, and Rudolph are careful about this, more careful than the advertisement: they note explicitly that “models where the quantum state is not a physical property can be constructed by dropping our assumption of preparation independence,” and the literature has confirmed the point with explicit constructions.⁴ Their verdict on such models is that

³N. Harrigan and R. W. Spekkens, “Einstein, incompleteness, and the epistemic view of quantum states,” *Foundations of Physics* 40 (2010), 125–157.

⁴P. G. Lewis, D. Jennings, J. Barrett, and T. Rudolph, “Distinct quantum states can be compatible with a single state of reality,” *Physical Review Letters* 109 (2012), 150404. For a thorough survey of the theorem and its escape routes, see M. S. Leifer, “Is the quantum state real? An extended review of ψ -ontology theorems,” *Quanta* 3 (2014), 67–155.

they “appear extremely contrived.” The model that violates the assumption, however, was not invented to dodge a theorem. It is the model Part II derived, for entirely independent reasons, from a puzzle about two Hadamard gates.

12.5 The law consults the history

Recall what Chapter 5 established and Chapters 6 through 8 built out: the dynamical laws of quantum mechanics, read statistically, are non-Markovian. The transition matrix governing an interaction is a function of the interacting systems’ statistical histories, computed (uniquely, by Chapter 6’s iterative proportional fitting) from everything that has happened since the last base case. Apply that machinery to the black box, and the paradox dissolves into arithmetic.

A bit prepared D sits at a base case: measured, definite, no history to speak of. A bit prepared U is one Hadamard past its base case: uniformly distributed, with that Hadamard in its history. The four preparations therefore hand the box four different *histories*, and the fitted single-step matrix at the box differs accordingly. For the DD preparation the incoming pair is itself a base case, so the box’s matrix is the unistochastic default $|U|^2$; for the others, fitting transports each preparation’s distribution to its output distribution. The response columns for the configurations that actually occur (the code listing computes them all):

	$f(00)$	$f(01)$	$f(10)$	$f(11)$
DD	$(0, \frac{1}{4}, \frac{1}{4}, \frac{1}{2})$			
UD	$(0, \frac{1}{2}, 0, \frac{1}{2})$		$(\frac{1}{2}, \frac{1}{2}, 0, 0)$	
DU	$(0, 0, \frac{1}{2}, \frac{1}{2})$	$(\frac{1}{2}, 0, \frac{1}{2}, 0)$		
UU	$(0, \frac{1}{2}, \frac{1}{2}, 0)$	$(1, 0, 0, 0)$	$(1, 0, 0, 0)$	$(0, \frac{1}{2}, \frac{1}{2}, 0)$

Each row is a valid stochastic response, each transports its preparation’s distribution to exactly the quantum output, and each has zeros in precisely the vetoed outcome’s position. Now read down the $f(00)$ column. The *same* incoming configuration receives four different responses under the four preparations, even though it is just two bits reading 00, physically indistinguishable in the moment. That is the violation of the theorem’s assumption, exhibited entry by entry: the box’s behavior is not a function of the present configuration alone. It is a function of the configuration *and the history that delivered it*. The box is, in the only sense we will ever need, *aware* of how the bits were prepared. This is not because the bits carry a memory, but because the law that fires when systems meet consults their pasts. The bits hold nothing beyond their values; the history-dependence belongs to the law. That is nomological history-dependence, Chapter 5’s distinction, doing here exactly the work it was introduced to do.

Is that contrived? It would be, if the model had been reverse-engineered from the theorem. It was not. The non-Markovian law was forced on us in Chapter 5, by the bare fact that one Hadamard makes a coin and two Hadamards make a certainty. PBR then arrives and proves, rigorously, that a *Markovian* statistical reading cannot survive: a reading in which the transition matrix at the box is fixed once and for all, blind to history. We agree, and had already retired that reading seven chapters ago, for its own reasons. What PBR actually establishes is not that the quantum state exists as a physical object; it is that the statistical laws must reach into the past: if the values are real and statistical, the laws are non-Markovian.

The theorem, read correctly, is not an argument against us. It proves precisely what we have been arguing: that real statistical values require non-Markovian laws.

Code: this chapter in Octave

```

% --- the gates, and the box of Section 12.2 ---
I2 = eye(2); H = (1/sqrt(2))*[1 1; 1 -1]; s = 1/sqrt(2);
U = [0 s s 0; .5 -.5 .5 .5; .5 .5 -.5 .5; s 0 0 -s];
norm(U*U' - eye(4)) % -> 0: unitary
e00 = [1;0;0;0];

% Four independent preparations: DD, UD, DU, UU
prep = {e00, kron(H,I2)*e00, kron(I2,H)*e00, kron(H,H)*e00};
labels = {'DD','UD','DU','UU'};
for k = 1:4
    disp([labels{k} ':']);
    disp(abs(U*prep{k}).^2) % zeros at outcome 00, 10, 01, 11 resp.
end
% outcome k never occurs under preparation k (D->0, U->1 encoding)

% No preparation-blind stochastic matrix exists: its column for the
% input configuration 00 would need zeros at outputs 00 (from DD),
% 01 (from DU), 10 (from UD), and 11 (from UU) - and cannot sum to 1.

% --- IPF: the box's fitted response per preparation ---
function G = ipf(A, pt, pt1)
    v = ones(length(pt),1);
    for k = 1:20000
        u = pt1 ./ (A*(v.*pt)); u(~isfinite(u)) = 0;
        v = 1 ./ (A'*u); v(~isfinite(v)) = 0;
    end
    G = diag(u)*A*diag(v);
end
dU = abs(U).^2;
for k = 1:4
    pt = abs(prepare{k}).^2; G = ipf(dU, pt, abs(U*prep{k}).^2);
    disp([labels{k} ' response to configuration 00:']);
    disp(G(:,1)) % four different responses to the same input
end
% DD's response equals the unistochastic default dU(:,1)' (base case).
% The law consults the history, not just the present configuration.

% --- the black box built from elementary gates (Chapter 3's RY and CX,
% Chapter 1's SWAP, Chapter 8's Givens) -----
RY = @(t) [cos(t/2) -sin(t/2); sin(t/2) cos(t/2)];
Xp = [0 1; 1 0]; Yp = [0 -1i; 1i 0]; Zp = [1 0; 0 -1];
CX12 = [1 0 0 0; 0 1 0 0; 0 0 0 1; 0 0 1 0];
SWg = [1 0 0 0; 0 0 1 0; 0 1 0 0; 0 0 0 1];
Gv = @(t) [1 0 0 0; 0 cos(t) -sin(t) 0; 0 sin(t) cos(t) 0; 0 0 0 1];
Ucirc = kron(I2, Zp*H) * CX12 * kron(I2, H*Yp*Xp*Yp) * Gv(pi/4) * SWg ...

```

```
    * kron(RY(-pi/2), I2) * CX12 * kron(RY(-pi/2), I2);  
norm(Ucirc - U)           % -> 0: the circuit is exactly U
```

Conclusion: Grappling with Underdetermination

Return, one final time, to the magician. You copied a single bit onto a thumb drive and handed it over; he named the category of program it came from, run after run, without fail. The introduction called this impossible, and as a statement about the bit it still is: a lone 0 carries no information about the ensemble it was drawn from. But the reader can now say what the introduction could not, which is where the information actually lives. The magician never read the bit. He *knew how the bit came to be*: he knew which program you had run to produce it, perhaps because he could see your computer’s screen the whole time.

Nature’s version of the trick, the second Hadamard gate that undoes the first, has the same explanation, with one refinement: there is no record. Nothing rides along with the bit. No hidden memory, no pilot, no physically existing distribution; whenever the memory of a quantum computer is read, out come bits with definite values and nothing else. The gate’s statistical law is simply a function of the circuit that preceded it, reaching back to the last base case. The trick was never in the bit, and never in an accompanying wave. It was in the law. *The law consults the history*. By now the act has also been performed under laboratory conditions: the Hadamard pair of Chapter 5 was the warm-up, and Chapter 12’s box was the full routine, two bits at a time, with a theorem certifying that nothing reading only the present could pull it off.

Now the summing up.

We have made one sustained claim: quantum mechanics is mathematically equivalent to a statistical theory of bits with definite values, governed by stochastic laws that are non-Markovian (they consult the history), non-local (they consult the arrangement), and contextual (they consult the interaction taking place).

The claim was defended twice over. Part II established it constructively: Chapter 5 proved the dynamics non-Markovian and traded the quantum state for a probability vector plus a sufficient statistic of its history; Chapter 6 rebuilt the theory as unistochastic dynamics grounded in base cases and supplied the unique single-step transition matrices; Chapters 7 and 8 then turned the machinery on the showpieces, tracing what a quantum algorithm works with to information stored in the history and computing interference through without a particle ever being in two places. Part III defended the claim against the no-go theorems: Kochen–Specker, Bell, Frauchiger–Renner, and Pusey–Barrett–Rudolph, each translated into an explicit circuit and computed in full, and each found to constrain the *laws* while leaving the *values* untouched.

The celebrated proofs that “reality is impossible” turned out, one after another, to be a

systematic characterization of what reality's laws are like. The statistical reading does not make quantum mechanics unmysterious; it relocates the mystery from the existence of the world to the laws of the world, where a mystery can be studied.

The dominant interpretations did not survive the same walk. Many Worlds takes ψ to be everything and never supplies a single observable with a value; it cannot explain what is observed, and by Bell's own reckoning it dissolves the past into a solipsism of the present moment. Collapse that varies from observer to observer makes other minds perspectival, which is a form of solipsism. Collapse that is physical and invariant is a different theory (GRW, Diósi–Penrose), deviating from quantum mechanics in principle and being hunted in exactly the regimes where it must show itself. These are not aesthetic complaints. Each reading, at the point where it meets observation, is inconsistent with the joint demands of realism and the quantum predictions, and Part III's purpose was to establish that strong claim one calculation at a time.

What survives is the statistical reading. But surviving is not the same as being finished, and the right way to close is to say exactly what remains open.

Three questions that must have answers

A realist who keeps definite values must supply three answers: which law moves the values, in which basis they live, and by which clock they are ordered.

The law first. If bits have values at every moment, then between any two moments there is a fact about how the values were transported: a single-step transition matrix. Chapter 6 supplied one: the unique column-stochastic matrix that reproduces the Born statistics at both ends while deforming the unistochastic default as little as possible, in the precise sense of minimizing Kullback–Leibler divergence, computed by iterative proportional fitting.

The basis second. Chapter 10 showed that definite values cannot be assigned to every basis at once; contextuality forbids it. A realist must therefore crown one basis as ontic, the basis whose values are the real configuration, and demote the rest to emergeables, manufactured at measurement.

The clock third. Chapter 11 showed that the values cannot be Lorentz invariant: the two-bit Frauchiger–Renner distillation makes the pre-measurement configuration depend on which of two space-like separated perturbations fired first. A realist who wants invariant values, and must want them on pain of solipsism, has to privilege one time-ordering: a preferred slicing of spacetime. A reader willing to give up dynamics altogether may have an escape here, since the all-at-once retrocausal reading of Section 11.10 carries nothing from one moment to the next and so has no slicing to privilege; we are not, and take the question as it stands.

Each question, taken alone, has an answer. The difficulty is that each has too many.

Three symmetries that hide the answers

For each question, a symmetry of quantum theory makes a whole family of distinct answers empirically identical, matching in every prediction and indistinguishable by any experiment.

The law is hidden by *indivisibility*. Quantum mechanics fixes the statistics at special moments, the base cases; between them, as Barandes's analysis makes precise, an entire

equivalence class of stochastic processes reproduces everything the theory fixes. The KL-minimizing matrix is a principled member of that class, and we have argued it is the right member to model, but the mathematics tolerates its rivals.

The basis is hidden by the *symmetry of configuration space*. Rotate every state and every operator by a common unitary and no statistic changes. Nothing in the predictive content of the theory crowns one basis over another.

The clock is hidden by the *symmetry of Minkowski space*. As Bell's Lorentzian retelling of relativity makes plain, a theory with a single real rest frame is empirically identical to a theory with none; the facts of physics, in his words, do not oblige us to choose. No measurement locates the preferred slicing.

Three questions that must have answers; three symmetries that keep the answers from ever being read off the mathematics. This is not a defect of the statistical reading. The symmetries belong to quantum theory itself, so these freedoms are the freedoms of *any* realist account, and the situation is the same underdetermination first discussed in Chapter 6 and met again in Chapters 10 and 11. As Chapter 10 put it: the theorems do not eliminate the realist reading; they inventory its freedoms.

Committing to answers

Commit to an answer for each of the three questions and what results is a fully specified, logically consistent ontology. This is no idle promise, because the literature contains a standing example. Commit to the position basis and to the guiding equation and one has non-relativistic Bohmian mechanics: particles with definite positions at all times, moving on deterministic trajectories, the wave function demoted to a nomological object that steers them. It reproduces every prediction of non-relativistic quantum mechanics and falls to none of the inconsistencies that sank the readings above.

Making it relativistic requires the third commitment, a preferred frame, and as Hrvoje Nikolić has shown this is not the obstacle it is usually assumed to be: a Bohmian theory that reproduces the measurable predictions of relativistic quantum field theory in *one* Lorentz frame thereby reproduces them, by the invariance of those predictions themselves, in all frames.⁵ His model revises the first commitment as well, taking the configuration of the fields as the beable in place of particle position, as Section 10.6 described.

The Lorentz-violating ontology hides behind Lorentz-invariant statistics. Which frame is the real one is then a further commitment; the cosmic rest frame, measurable through the microwave background dipole, is a plausible candidate.

We made three choices of its own, suited to its laboratory. For the law, the KL-minimizing transition matrix of Chapter 6. For the basis, the computational basis of the quantum computer running the program. For the clock, the rest frame of that computer. Together they deliver the picture we have drawn throughout: bits inside the machine, carrying well-defined values at all times, hopping stochastically from configuration to configuration under laws that consult their history. Everything is definite; nothing is weird; the sorcery is gone.

And yet all three are conventions rather than findings. Imagine a second machine, built along different physical axes, in motion relative to ours. Its computational basis is not ours; its rest frame is not ours; its fitted law transports its own steps. Its owner's picture is as

⁵H. Nikolić, "Relativistic QFT from a Bohmian perspective: A proof of concept," arXiv:2205.05986 (2022).

internally coherent as ours, and it disagrees with ours: Chapter 11 computed how sharply, with two slicings of time assigning the very same bits different, mutually incompatible pasts. Both accounts cannot be correct, so some fact favors one of them. And the three symmetries guarantee that no experiment will ever say which.

What Occam's razor keeps

It is tempting to read the symmetries' guarantee, that no experiment can ever say which machine's frame, basis, and clock are the real ones, as proof that there is no fact: that the values are simply indefinite until observed. The temptation should be refused, because the entire century-long dispute lives at this exact spot.

That the answers cannot be confirmed does not show that they do not exist. Denying the values, the doctrine of value indefiniteness, is itself a substantive metaphysical position, and Part III found it untenable in every form it can take: it makes other minds perspectival, which is solipsism; it offers a world of pure statistics with nothing left to point at, which cannot explain what observation reports; and its physical implementations either deviate from quantum mechanics in principle or never say when collapse occurs. The no-go theorems, read correctly, never established it. They constrain the laws; the values they leave alone.

The persistent confusion has its root in a misreading of Occam's razor. The razor does not say: discard every assumption that cannot be empirically verified. It says: assume no more than is necessary for a logically consistent account of what *can* be verified. These are very different rules. The assumption that the bits have values regardless of observation cannot be verified directly; but drop it, and the outcomes of the Bell tests, the contextuality experiments, and the interference and Wigner's-friend circuits we have computed become inexplicable. An assumption whose removal makes the verifiable inexplicable is a necessity the razor keeps.

Two paths

Two paths remain open to a realist who keeps a dynamics, and they differ in one respect only: whether the three answers are held as conventions or as claims.

On the first, one never commits. The three questions have answers, because realism requires it; the answers cannot be known, because the symmetries forbid it; so one adopts whatever triple the bench suggests (the computational basis of the machine at hand, its rest frame, its fitted transition law), uses it as a working model of the underdetermined fact, and exchanges it without regret when the context changes. The choice is a convention, held the way we have held all its conventions, while keeping clear which is which. This is the path we have taken.

On the second, one commits: a particular basis, a particular frame, a particular law, held not as convention but as the actual ontology (perhaps the position basis or the configuration of the fields, the guiding equation, and the cosmic rest frame, which assemble into a relativistic Bohmian mechanics). Experiment cannot justify the commitment, because the symmetries have silenced experiment on precisely this question; the justification must be explanatory, metaphysical, a case that one basis or one frame is more fundamental on grounds outside the reach of measurement. It is a respectable path, and it is where the realist tradition from de Broglie to Bell has generally settled.

Choosing between the two is a question of philosophical temperament, and we will not pretend to force it. What it does insist on is what the two have in common. Both keep the values real. Both accept laws that are non-Markovian, non-local, and contextual. Both acknowledge the three symmetries and the humility they impose. What neither permits is the move so much of the last century made: abandoning the reality of the values altogether. The option is closed by contradiction, not by preference.

What has been shown, in sum: the measurement problem dissolves into Bayesian updating; interference is computed without a particle ever being in two places; the friend, the spins, and the black boxes of the no-go theorems are each explained with definite values intact; what a quantum algorithm works with is traced to information in the history. What has been conceded: laws that consult the history and the whole arrangement, and three answers that no experiment can supply. What was never conceded is the definiteness of the world. The bits hold values between the gates, whether or not anyone reads them, and the second Hadamard gate behaves as it does because of the circuit that came before it, not because anything traveled alongside the bit.

References

- Aaronson, S., “Quantum Computing and Hidden Variables,” *Physical Review A* 71 (2005), 032325.
- Aaronson, S., Y. Atia, and L. Susskind, “On the hardness of detecting macroscopic superpositions,” arXiv:2009.07450 (2020).
- Aaronson, S., and D. Gottesman, “Improved Simulation of Stabilizer Circuits,” *Physical Review A* 70 (2004), 052328.
- Adlam, E., “Does Science need Intersubjectivity? The Problem of Confirmation in Orthodox Interpretations of Quantum Mechanics,” *Synthese* 200 (2022), 522.
- Adlam, E., “Two Roads to Retrocausality,” arXiv:2201.12934 (2022).
- Adlam, E., and C. Rovelli, “Information is Physical: Cross-Perspective Links in Relational Quantum Mechanics,” arXiv:2203.13342 (2022).
- Aharonov, Y., D. Z. Albert, and L. Vaidman, “How the result of a measurement of a component of the spin of a spin-1/2 particle can turn out to be 100,” *Physical Review Letters* 60 (1988), 1351–1354.
- Aharonov, Y., P. G. Bergmann, and J. L. Lebowitz, “Time Symmetry in the Quantum Process of Measurement,” *Physical Review* 134 (1964), B1410–B1416.
- Aharonov, Y., and L. Vaidman, “The Two-State Vector Formalism: An Updated Review,” arXiv:quant-ph/0105101 (2002).
- Albert, D. Z., “Quantum mechanics without a wave-function,” *Synthese* 208 (2026), article 141.
- Althusser, L., “Contradiction and Overdetermination,” in *For Marx*, trans. B. Brewster (London: Verso, 2010).
- Aaronson, S., in discussion with J. Barandes, hosted by C. Jaimungal, “Harvard Scientist Rewrites the Rules of Quantum Mechanics | Scott Aaronson & Jacob Barandes,” video (3 March 2025).
- Balashov, Y., and M. Janssen, “Presentism and Relativity,” *British Journal for the Philosophy of Science* 54 (2003), 327–346.
- Ballentine, L. E., “The Statistical Interpretation of Quantum Mechanics,” *Reviews of Modern Physics* 42 (1970), 358–381.
- Barandes, J. A., “The Stochastic-Quantum Correspondence,” arXiv:2302.10778 (2023).
- Bravyi, S., and D. Gosset, “Improved Classical Simulation of Quantum Circuits Dominated by Clifford Gates,” *Physical Review Letters* 116 (2016), 250501.

- Bell, J. S., "On the Einstein Podolsky Rosen paradox," *Physics* 1 (1964), 195–200.
- Bell, J. S., "Subject and object," in *The Physicist's Conception of Nature* (Dordrecht: Reidel, 1973).
- Bell, J. S., "How to teach special relativity," *Progress in Scientific Culture* 1, no. 2 (1976).
- Bell, J. S., "Quantum mechanics for cosmologists," in C. J. Isham, R. Penrose, and D. W. Sciama (eds.), *Quantum Gravity 2* (Oxford: Clarendon Press, 1981).
- Bell, J. S., "On the Problem of Hidden Variables in Quantum Mechanics," *Reviews of Modern Physics* 38 (1966), 447–452.
- Bell, J. S., "On the impossible pilot wave," *Foundations of Physics* 12 (1982), 989–999.
- Bell, J. S., "Six possible worlds of quantum mechanics" (1986), in *Speakable and Unsayable in Quantum Mechanics*, second edition (Cambridge: Cambridge University Press, 2004), p. 191.
- Bell, J. S., "Are there quantum jumps?" in C. W. Kilmister (ed.), *Schrödinger: Centenary Celebration of a Polymath* (Cambridge: Cambridge University Press, 1987), pp. 41–52.
- Bell, J. S., *Speakable and Unsayable in Quantum Mechanics* (Cambridge: Cambridge University Press, 1987).
- Bell, J. S., "Against 'measurement'," *Physics World* 3, no. 8 (August 1990), 33–40.
- Bennett, C. H., and G. Brassard, "Quantum cryptography: Public key distribution and coin tossing," in *Proceedings of IEEE International Conference on Computers, Systems and Signal Processing* (Bangalore, 1984), 175–179.
- Benoist, J., *Toward a Contextual Realism* (Cambridge, MA: Harvard University Press, 2021).
- Bloch, F., "Nuclear Induction," *Physical Review* 70 (1946), 460–474.
- Blokhintsev, D. I., "Kritika idealisticheskogo ponimaniya kvantovoy teorii" [Critique of the idealist understanding of quantum theory], *Uspekhi Fizicheskikh Nauk* 45, no. 2 (1951), 195–228.
- Blokhintsev, D. I., *The Philosophy of Quantum Mechanics* (Dordrecht: D. Reidel, 1968).
- Bohm, D., *Wholeness and the Implicate Order* (London: Routledge & Kegan Paul, 1980).
- Bohm, D., *Quantum Theory* (New York: Prentice-Hall, 1951).
- Bohm, D., "A Suggested Interpretation of the Quantum Theory in Terms of 'Hidden' Variables," I and II, *Physical Review* 85 (1952), 166–179 and 180–193.
- Born, M., "Zur Quantenmechanik der Stoßvorgänge," *Zeitschrift für Physik* 37 (1926), 863–867.
- Born, M., "Quantenmechanik der Stoßvorgänge," *Zeitschrift für Physik* 38 (1926), 803–827.
- Bub, J., "'Two Dogmas' Redux," arXiv:1907.06240 (2019).
- Buniy, R. V., and S. D. H. Hsu, "Everything is Entangled," *Physics Letters B* 718 (2012), 233–236.
- Catani, L., M. Leifer, D. Schmid, and R. W. Spekkens, "Why interference phenomena do not capture the essence of quantum theory," *Quantum* 7 (2023), 1119.
- Chalmers, D. J., *The Conscious Mind: In Search of a Fundamental Theory* (New York: Oxford University Press, 1996).
- Clauser, J. F., M. A. Horne, A. Shimony and R. A. Holt, "Proposed Experiment to Test Local

- Hidden-Variable Theories,” *Physical Review Letters* 23 (1969), 880–884.
- Deming, W. E., and F. F. Stephan, “On a Least Squares Adjustment of a Sampled Frequency Table When the Expected Marginal Totals are Known,” *Annals of Mathematical Statistics* 11 (1940), 427–444.
- Deutsch, D., “Quantum theory, the Church–Turing principle and the universal quantum computer,” *Proceedings of the Royal Society of London A* 400 (1985), 97–117.
- Deutsch, D., “Quantum theory as a universal physical theory,” *International Journal of Theoretical Physics* 24, no. 1 (1985), 1–41.
- Deutsch, D., *The Fabric of Reality* (London: Allen Lane, 1997).
- Deutsch, D. and R. Jozsa, “Rapid solution of problems by quantum computation,” *Proceedings of the Royal Society of London A* 439 (1992), 553–558.
- Diósi, L., “Models for universal reduction of macroscopic quantum fluctuations,” *Physical Review A* 40 (1989), 1165–1174.
- Drezet, A., “An Elementary Proof That Everett’s Quantum Multiverse Is Nonlocal,” arXiv:2306.07794 (2023).
- Eberhard, P. H., “Bell’s theorem and the different concepts of locality,” *Il Nuovo Cimento B* 46 (1978), 392–419.
- Einstein, A., “Autobiographical Notes,” in P. A. Schilpp (ed.), *Albert Einstein: Philosopher-Scientist* (New York: MJF Books, 2001).
- Einstein, A., “Bestimmt Schrödingers Wellenmechanik die Bewegung eines Systems vollständig oder nur im Sinne der Statistik?” (Berlin, 5 May 1927), Einstein Archive 2-100, Doc. 516 in *The Collected Papers of Albert Einstein*, Vol. 15 (Princeton: Princeton University Press, 2018).
- Einstein, A., “Clerk Maxwell’s Influence on the Evolution of the Idea of Physical Reality” (1931), reprinted in *The World As I See It* (New York: Covici-Friede, 1934).
- Einstein, A., “Quanten-Mechanik und Wirklichkeit,” *Dialectica* 2, no. 3–4 (1948), 320–324.
- Einstein, A., “Remarks Concerning the Essays Brought together in this Co-operative Volume,” in P. A. Schilpp (ed.), *Albert Einstein: Philosopher-Scientist* (New York: MJF Books, 2001).
- Einstein, A., letter to M. Born, 3 March 1947, in *The Born–Einstein Letters*, trans. I. Born (London: Macmillan, 1971).
- Einstein, A., letter to M. Born, 12 May 1952, in *The Born–Einstein Letters*, trans. I. Born (London: Macmillan, 1971), letter 99.
- Einstein, A., B. Podolsky, and N. Rosen, “Can Quantum-Mechanical Description of Physical Reality Be Considered Complete?” *Physical Review* 47 (1935), 777–780.
- Elitzur, A. C., and L. Vaidman, “Quantum mechanical interaction-free measurements,” *Foundations of Physics* 23 (1993), 987–997.
- Engels, F., *Dialectics of Nature*, trans. C. Dutt (Moscow: Progress Publishers, 1954).
- Feynman, R. P., “Space-Time Approach to Non-Relativistic Quantum Mechanics,” *Reviews of Modern Physics* 20 (1948), 367–387.
- Feynman, R. P., R. B. Leighton, and M. Sands, *The Feynman Lectures on Physics*, vol. 3

- (Reading, MA: Addison–Wesley, 1965).
- Frauchiger, D., and R. Renner, “Quantum theory cannot consistently describe the use of itself,” *Nature Communications* 9, 3711 (2018).
- Galvan, B., “On the preferred-basis problem and its possible solutions,” arXiv:1008.3708 (2010).
- Ghirardi, G. C., A. Rimini, and T. Weber, “Unified dynamics for microscopic and macroscopic systems,” *Physical Review D* 34 (1986), 470–491.
- Gottesman, D., and I. Chuang, “Quantum Digital Signatures,” arXiv:quant-ph/0105032 (2001).
- Greenberger, D. M., M. A. Horne, and A. Zeilinger, “Going beyond Bell’s theorem,” in M. Kafatos (ed.), *Bell’s Theorem, Quantum Theory and Conceptions of the Universe* (Dordrecht: Kluwer Academic, 1989), pp. 69–72.
- Gottesman, D., “The Heisenberg Representation of Quantum Computers,” arXiv:quant-ph/9807006 (1998).
- Gu, M., K. Wiesner, E. Rieper, and V. Vedral, “Quantum mechanics can reduce the complexity of classical models,” *Nature Communications* 3, 762 (2012).
- Harrigan, N., and R. W. Spekkens, “Einstein, incompleteness, and the epistemic view of quantum states,” *Foundations of Physics* 40 (2010), 125–157.
- ’t Hooft, G., “Explicit construction of Local Hidden Variables for any quantum theory up to any desired accuracy,” arXiv:2103.04335 (2021).
- ’t Hooft, G., *The Cellular Automaton Interpretation of Quantum Mechanics* (Cham: Springer, 2016).
- Hossenfelder, S., “Superdeterminism: A Guide for the Perplexed,” arXiv:2010.01324 (2020).
- Ilyenkov, E. V., *Dialectical Logic: Essays on its History and Theory*, trans. H. Campbell Creighton (Moscow: Progress Publishers, 1977).
- Kastner, R. E., “The ‘Delayed Choice Quantum Eraser’ Neither Erases Nor Delays,” *Foundations of Physics* 49 (2019), 717–727.
- Kent, A., “Does it Make Sense to Speak of Self-Locating Uncertainty in the Universal Wave Function? Remarks on Sebens and Carroll,” *Foundations of Physics* 45 (2015), 211–217.
- Kent, A., “Against Many-Worlds Interpretations,” *International Journal of Modern Physics A* 5 (1990), 1745–1762.
- Kim, Y.-H., R. Yu, S. P. Kulik, Y. Shih, and M. O. Scully, “Delayed ‘Choice’ Quantum Eraser,” *Physical Review Letters* 84 (2000), 1–5.
- Kochen, S., and E. P. Specker, “The problem of hidden variables in quantum mechanics,” *Journal of Mathematics and Mechanics* 17 (1967), 59–87.
- Korzekwa, K., and M. Lostaglio, “Quantum Advantage in Simulating Stochastic Processes,” *Physical Review X* 11 (2021), 021019.
- Kripke, S. A., *Wittgenstein on Rules and Private Language* (Cambridge, MA: Harvard University Press, 1982).
- Kumar, M., *Quantum: Einstein, Bohr and the Great Debate about the Nature of Reality* (London: Icon Books, 2008).

- Kruithof, J., “Telefoonverkeersrekening,” *De Ingenieur* 52 (1937), E15–E25.
- Leifer, M. S., “Is the quantum state real? An extended review of ψ -ontology theorems,” *Quanta* 3 (2014), 67–155.
- Lewis, P. G., D. Jennings, J. Barrett, and T. Rudolph, “Distinct quantum states can be compatible with a single state of reality,” *Physical Review Letters* 109 (2012), 150404.
- Malavolta, G., H. Nieuwboer, A. Ramachandran, and M. Walter, “Robustness of Hidden-Variable Theories and Matrix Scaling,” arXiv:2609.14519 (2026).
- Maudlin, T., “Can the World be Only Wavefunction?” in S. Saunders, J. Barrett, A. Kent, and D. Wallace (eds.), *Many Worlds? Everett, Quantum Theory, and Reality* (Oxford: Oxford University Press, 2010), pp. 121–143.
- Maudlin, T., *Philosophy of Physics: Quantum Theory* (Princeton: Princeton University Press, 2019).
- Mermin, N. D., “Quantum mysteries revisited,” *American Journal of Physics* 58 (1990), 731–734.
- Nikolić, H., “Relativistic QFT from a Bohmian perspective: A proof of concept,” arXiv:2205.05986 (2022).
- Pais, A., “*Subtle is the Lord...: The Science and the Life of Albert Einstein* (Oxford: Oxford University Press, 1982).
- Penrose, R., *The Emperor’s New Mind: Concerning Computers, Minds, and the Laws of Physics* (Oxford: Oxford University Press, 1989).
- Penrose, R., “On gravity’s role in quantum state reduction,” *General Relativity and Gravitation* 28 (1996), 581–600.
- Planck Collaboration, “Planck 2018 results. I. Overview and the cosmological legacy of *Planck*,” *Astronomy & Astrophysics* 641 (2020), A1.
- Pris, I., *Contextual Realism and Quantum Mechanics* (Chisinau: Scholars’ Press, 2020).
- Pris, F.-I., “On Quantum Fundamentalism,” arXiv:2307.12992 (2023).
- Pris, F.-I., “The Real Meaning of Quantum Mechanics,” arXiv:2107.10666 (2021).
- Pusey, M. F., J. Barrett, and T. Rudolph, “On the reality of the quantum state,” *Nature Physics* 8 (2012), 475–478.
- Qureshi, T., “The Enigma of Delayed Choice Quantum Eraser,” *Quanta* 14 (2025), 66–74.
- Resnick, S., and R. Wolff, “Everythingism, or Better Still, Overdetermination,” *New Left Review* I/195 (September–October 1992), 124–126.
- Rovelli, C., “Relational quantum mechanics,” *International Journal of Theoretical Physics* 35 (1996), 1637–1678.
- Rovelli, C., *Helgoland: Making Sense of the Quantum Revolution*, trans. E. Segre and S. Carnell (New York: Riverhead Books, 2021).
- Sanchioni, M., “Exploring the Interplay Between Wave Function Realism and Gauge Symmetry Interpretations in Quantum Mechanics,” *Foundations of Physics* 54 (2024), article 59.
- Schlosshauer, M., “Decoherence, the Measurement Problem, and Interpretations of Quantum Mechanics,” *Reviews of Modern Physics* 76 (2004), 1267–1305.
- Schrödinger, E., “Über die Umkehrung der Naturgesetze,” *Sitzungsberichte der Preussischen*

- Akademie der Wissenschaften, Physikalisch-mathematische Klasse* 8 N9 (1931), 144–153. Translated by R. Chetrite, P. Muratore-Ginanneschi and K. Schwieger as “On the Reversal of the Laws of Nature,” *European Physical Journal H* 46 (2021), article 28.
- Schrödinger, E., “Die gegenwärtige Situation in der Quantenmechanik,” *Naturwissenschaften* 23 (1935), 807–812, 823–828, 844–849; translated by J. D. Trimmer as “The Present Situation in Quantum Mechanics,” *Proceedings of the American Philosophical Society* 124 (1980), 323–338.
- Schrödinger, E., *Science and Humanism* (Cambridge: Cambridge University Press, 1951).
- Scully, M. O., and K. Drühl, “Quantum eraser: A proposed photon correlation experiment concerning observation and ‘delayed choice’ in quantum mechanics,” *Physical Review A* 25 (1982), 2208–2213.
- Sebens, C. T., and S. M. Carroll, “Self-Locating Uncertainty and the Origin of Probability in Everettian Quantum Mechanics,” *British Journal for the Philosophy of Science* 69 (2018), 24–75.
- Sinkhorn, R., and P. Knopp, “Concerning nonnegative matrices and doubly stochastic matrices,” *Pacific Journal of Mathematics* 21 (1967), 343–348.
- Smerlak, M., and C. Rovelli, “Relational EPR,” *Foundations of Physics* 37 (2007), 427–445.
- Spekkens, R. W., “Why I Am Not a Psi-ontologist,” lecture, Perimeter Institute, PIRSA:12050021 (2012), <https://pirsa.org/12050021>.
- Stacey, B. C., “The De-Relationalizing of Relational Quantum Mechanics,” arXiv:2211.03230 (2022).
- Stoica, O. C., “3D-Space and the preferred basis cannot uniquely emerge from the quantum structure,” *Advances in Theoretical and Mathematical Physics* 26 (2022), 3895–3962.
- von Neumann, J., *Mathematische Grundlagen der Quantenmechanik* (Berlin: Springer, 1932).
- Wheeler, J. A., “The ‘Past’ and the ‘Delayed-Choice’ Double-Slit Experiment,” in A. R. Marlow (ed.), *Mathematical Foundations of Quantum Theory* (New York: Academic Press, 1978), 9–48.
- Wheeler, J. A., “Law Without Law,” in J. A. Wheeler and W. H. Zurek (eds.), *Quantum Theory and Measurement* (Princeton: Princeton University Press, 1983).
- Wigner, E. P., “Remarks on the mind–body question,” in I. J. Good (ed.), *The Scientist Speculates* (London: Heinemann, 1961); reprinted in E. P. Wigner, *Symmetries and Reflections* (Bloomington: Indiana University Press, 1967), pp. 171–184.
- Wittgenstein, L., *Philosophical Investigations*, trans. G. E. M. Anscombe (Oxford: Blackwell, 1953).
- Wootters, W. K., and W. H. Zurek, “A single quantum cannot be cloned,” *Nature* 299 (1982), 802–803.
- Zeilinger, A., “A Foundational Principle for Quantum Mechanics,” *Foundations of Physics* 29, no. 4 (1999), 631–643.
- Zurek, W. H., “Relative States and the Environment: Einselection, Envariance, Quantum Darwinism, and the Existential Interpretation,” arXiv:0707.2832 (2007).
- Zurek, W. H., “Decoherence, einselection and the existential interpretation (the rough guide),” *Philosophical Transactions of the Royal Society of London A* 356 (1998), 1793–1821.