Endgame

Bitcoin, Sovereign Debt, and the Return of Sound Capital

KEVIN BELL, CFA

Founder, Cadena Bitcoin

ADVANCE MANUSCRIPT · 2026

Endgame: Bitcoin, Sovereign Debt, and the Return of Sound Capital

Kevin Bell, CFA — Founder, Cadena Bitcoin

cadenabitcoin.com

CONTENTS

PART IV COMPLETING THE SYSTEM

Author's Note

In 2025 I founded a company called Cadena Bitcoin. Cadena builds a non-custodial, DLC-enforced credit market in which Bitcoin holders can lend and borrow against their Bitcoin without ever surrendering their keys. The thesis of this book — that the existing monetary system is epistemically broken, that Bitcoin is the response, and that a market-discovered yield curve denominated against Bitcoin collateral is what completes that response — is also the professional bet I am making. If the argument in this book is right, the work I do becomes meaningful. If the argument is wrong, my company is misallocated effort and will still be one of the most rewarding experiences of my life.

I would rather you know that going in than discover it on the back cover. The book stands or falls on the quality of its arguments and the durability of its evidence, not on my conviction. I have tried to write the version of these ideas that survives the most skeptical reader.

I have also deliberately limited the book's scope. I do not write about how to set up a hardware wallet. I do not write about which jurisdictions are favorable, what the tax treatment is in any country, or what the regulatory landscape looks like in any given month. Those subjects matter, but they are moments in time. They will be obsolete before this book is in your hands. The frame I am trying to give you is durable; the products and rules that operate within that frame will keep evolving. Custody mechanics, payment rails, jurisdictional arbitrage, and the specific shape of any one credit product are downstream of the frame.

What I am trying to do is supply the frame.

I will not repeat this disclosure throughout the book. The reader should assume it applies to every chapter, and particularly to Parts III and IV, where the recommendations bear most directly on the kind of work my company does. I have tried to write arguments that survive the disclosure rather than depend on the reader forgetting it. Where the productive-layer infrastructure that Cadena builds is relevant to the structural argument, I describe it as a category of products being built by a number of operators, rather than as my company specifically, because the structural claim is what matters and the specific operator who serves any given reader is downstream of considerations — jurisdiction, fee structure, custody architecture, counterparty relationships — that this book is the wrong place to litigate. The reader who finds the framework persuasive will have his own decisions to make about specific operators, and I trust him to make them on the merits rather than on the strength of any name-recognition I might attempt to leverage in these pages.

Preface

This book is for Bitcoiners. It is not an attempt to convert anyone, and there is no Bitcoin 101 chapter. If you do not already understand what Bitcoin is, why it matters, and roughly how it works, this is not the book for you, and I recommend you start with Saifedean Ammous's *The Bitcoin Standard* or Lyn Alden's *Broken Money*. Read those, then come back. I will be here.

If you are already convinced of Bitcoin's importance, this book is trying to do three things you may not yet have seen done together in the existing literature.

First, it makes the Hayekian epistemic critique of central banking explicit and complete. There is a great deal of moral and distributional argument in the Bitcoin canon — debasement is theft, money printing benefits insiders, every fiat regime ends — and all of it is correct. But there is a deeper argument, hiding in plain sight. The argument is that central banking is not bad because the central bankers are corrupt or captured. It is bad because the project itself is, on the terms of mainstream economics, epistemically incoherent. The knowledge required to set the right monetary aggregate or interest rate cannot in principle be possessed by any committee, because it only emerges from the market process the committee is interrupting. This argument works against the best possible central banks, staffed by saints. And it slots into a Hayekian intellectual tradition that almost every serious economist already accepts in other domains. Part I of this book develops this argument carefully.

Second, it reframes the Bitcoin allocation question through the lens of portfolio construction rather than through the lens of conviction. Bitcoiners have a tendency to argue from inside the worldview, which produces unfalsifiable claims about inevitability and price targets. I want to argue from outside, in the language a CFA charterholder uses with a sophisticated client: asymmetric payoff distributions, position sizing under uncertainty, the form-of-holding question (ETF versus self-custody as different risk profiles), and — most importantly — the case that sovereign debt is not actually risk-free and should be replaced in serious portfolios by a different fixed-income allocation. Part III of this book is that argument.

Third, and this is the part I think no other Bitcoin book has yet attempted, it argues that what is currently emerging in Bitcoin-native credit markets is not just another product category. It is the operationalization of the Hayekian critique into working infrastructure. A market-discovered yield curve denominated in Bitcoin collateral, built from many loans of varying tenors and structures, would be the first true yield curve in any major economy in fifty years. The yield curve is not a financial product. It is the price signal that coordinates capital across time. When central banks suppress and shape the curve, capital is misallocated across the entire time structure of the economy — long-duration projects get funded that shouldn't, savers cannot price the future, and the productive capital structure gets warped as does every price signal in the economy. Restoring a market-discovered curve, on a sound-money base, completes the system that fixed supply alone cannot complete. Part IV of this book is that argument.

I am building on Ammous, Alden, and Booth. They are the three thinkers who have most shaped the contemporary Bitcoin macro position, and you will hear echoes of all of them throughout these pages. I owe each of them a substantial intellectual debt. Where I diverge, I try to

mark the divergence; where I am only restating, I try to credit them. None of them are responsible for any errors in this book.

A final note on what you will not find here. I will not predict the price of Bitcoin in any year. I will not tell you Bitcoin is inevitable. I will not tell you the dollar will collapse on a particular timeline. The argument I am making does not depend on any of those things, and I am suspicious of writers who have to depend on them. The case for Bitcoin in a serious portfolio is the case for an asymmetric payoff under a probability-weighted distribution of monetary regimes. That case stands whether or not the most dramatic outcomes occur. It is the case I will try to make to you across the rest of this book.

PART I

The Diagnosis

CHAPTER ONE

The Knowledge Problem

✦

In September of 1945, in the *American Economic Review*, Friedrich Hayek published a paper called "The Use of Knowledge in Society." It is fourteen pages long. It is one of the most cited papers in the history of economics.

Hayek's central claim is that the economic problem of society is not what most economists at the time thought it was. The problem is not the allocation of given resources, with given preferences, against given technical possibilities. If we knew all of those things — if a planner could simply be handed the relevant data — then economics would reduce to a calculation problem, solvable by sufficiently powerful computers and sufficiently smart bureaucrats. Hayek's claim is that the economic problem is not this problem at all. The economic problem is the use of dispersed knowledge that exists nowhere in concentrated form. The relevant facts about supply and demand, about local conditions, about transient shortages and surpluses, about the small adjustments by which an economy continually re-coordinates itself, are not facts that any single mind possesses or can possess. They exist in the heads of billions of actors, distributed across time and place, much of it tacit, none of it ever fully articulated.

This is why Hayek so admired the price system. The price system, he argued, is a mechanism that aggregates this dispersed local knowledge into a single, actionable signal. When tin becomes scarce, for whatever reason — a strike in Bolivia, a new demand for tin somewhere in

industry, a flood at a smelter — the price of tin rises. Nobody who responds to that price needs to know why tin became scarce. They just need to see the price rise and adjust. A user of tin economizes; a producer ramps up; a substitute material enters the calculation. The economy reorganizes around the new condition without anyone in particular having to understand the condition in detail. The price did the work. The price was the knowledge.

This insight applies, with particular force, to Bitcoin. Bitcoin's price is the only continuous, real-time aggregation of human judgment about the value of a sound-money alternative to the existing fiat regime. The price moves not because anyone has told it to but because millions of participants, each with their own information about their own situations and time preferences and tolerance for monetary risk, are making continuous decisions about how much purchasing power to commit to the asset. The price is the knowledge, and the knowledge is too dispersed for any analyst or institution to reproduce. I will return to the importance of listening to the price in the Afterword. For now, the point is that Hayek's insight about the price system is not a theoretical proposition you can take or leave. It is an operational fact about the only mechanism humans have ever devised for aggregating dispersed knowledge about value at scale.

This is what Hayek called the "marvel" of the market. He did not mean it figuratively. He meant that no other system humans have constructed, or could construct, accomplishes this knowledge problem at the scale and speed the price system accomplishes it.

It is worth pausing on how broadly accepted this argument is. Hayek won the Nobel Prize in 1974 in part for this work. The collapse of the Soviet Union is widely understood by serious economists as the empirical vindication of Hayek's argument: a planner with apparently unlimited authority and apparently unlimited computational resources

could not, in the end, replicate what the price system does because the planner did not have access to the knowledge that the price system aggregates. Even economists who do not call themselves Austrians accept Hayek's basic point. It is one of the rare propositions in economics that crosses ideological lines.

And then those same economists walk into the newly renovated Eccles Building and the neighboring 1951 Constitution Avenue Building, where the Federal Reserve employs more than 500 researchers, including more than 400 PhD economists and then sets, by committee, the price of money.

If it weren't true, you couldn't make it up.

This is the asymmetry that I want to put in front of you, because I think it is the central intellectual fact of the modern monetary order. We accept Hayek for tin. We accept Hayek for labor. We accept Hayek for capital goods, for raw materials, for consumer products, for almost every market in the economy. We do not accept Hayek for the unit in which all of those things are denominated. For the price of money — that is, for interest rates, for the size of the monetary aggregate, for the rate at which credit expands or contracts — we — everyone, everywhere, without voting for it, in every major economy in the world — have decided that a committee of twelve to nineteen people, meeting eight times a year, should set the value by deliberate judgment.

The asymmetry is not defended. It is just inherited. It would not survive if it were defended on Hayekian terms, because there is no Hayekian defense of it. The knowledge problem applies at least as much to the price of money as it applies to the price of tin. Probably more. The right level of the money supply at any given moment is a function of countless dispersed facts about the demand for cash balances, the velocity of circulation, the productive capacity of the economy, the time preferences of millions of actors, the structure of credit markets, the state

of expectations, and a dozen other variables that no committee possesses in real time. The right interest rate is the price at which the supply of savings clears the demand for borrowable funds across time, and that price is the most knowledge-intensive price in the economy because it coordinates the entire intertemporal structure of capital. If Hayek is right that the price of tin cannot be set by committee, then the price of money cannot be set by committee, *a fortiori*.

This is not a moral argument. I am not claiming that the people on the Federal Open Market Committee are corrupt or captured. I have no reason to believe any of those things, and the argument I am making does not require me to believe them. The members of the FOMC could be saints. They could have superhuman judgment. They would still be doing something epistemically incoherent, because the knowledge required to do their job correctly is not knowledge that any committee can possess. The market process they are interrupting is the only mechanism by which that knowledge is produced.

This is the strongest version of the case against central banking, and it is also the version least often made by Bitcoiners. Most of us learned this argument in its populist form: the Fed is corrupt, money printer go brrr, bankers are insiders, debasement is theft. All of those statements have something to them, and I do not want to dismiss them. But each of them invites a counterargument that depends on the moral character of the people involved or on the political economy of the institution. The Hayekian argument does not depend on any of that. It says: even if the people are saints, even if the institution is honest, the project is incoherent and is leading to suboptimal economic outcomes that are, by now, obvious to anyone who has been paying attention.

I want to plant a flag here, because I am going to come back to this throughout the book. The case against central banking that I will be making is not that central bankers are bad people. The case is that they

are attempting, by deliberation, to set a variable that cannot in principle be set by deliberation. That variable is the price of money. The market discovery of that price is not optional. It is the only mechanism that produces it.

The rest of Part I will develop this claim in three further directions. In the next chapter I will show why money is special — why corrupting the unit of account is not like corrupting any other price, but is in fact a distortion of every price in the economy simultaneously. In the chapter after that I will show what the consequences look like when this distortion runs for fifty years, using the Cantillon framework as the lens. Then I will show what the consequence has been for productivity and wages — the part of the story that mainstream economists are most reluctant to acknowledge. And I will close Part I with the international evidence, which makes clear that this is not an American story but a global one.

If you find the Hayekian argument congenial, the rest of the diagnosis will follow naturally. If you do not, I would ask you to sit with it for a few chapters before deciding. It took Hayek his entire career to develop, and it has taken me longer than I would like to admit to see how completely it applies to money itself.

CHAPTER TWO

Money as Information

✦

Money is not a commodity priced by the market. Money is the substrate in which all other prices are denominated.

This sentence does most of the work of the book, so it is worth sitting with for a moment.

When we say the price of oil is seventy-five dollars a barrel, we are making two claims. The first claim is about oil — about its scarcity, its production cost, its substitutes, its logistical situation, the strategic and political conditions affecting its supply. The second claim is about dollars — about how many of them it takes today, in this place, to acquire one barrel. Both halves of the price are doing work. If oil conditions change, the number changes. If dollar conditions change, the number changes. A price quoted in dollars contains information about both the numerator and the denominator.

This is true of every price in the economy. Every one of them. The price of labor, the price of housing, the price of a coffee, the price of a share of stock, the price of a bond, the price of a barrel of oil. Every one of those prices is a ratio between the thing being priced and the unit in which it is being priced. If the unit is unstable, every price in the economy contains a confounded signal — partly information about the priced thing, partly information about the unit. A market participant cannot, from a single price observation, distinguish how much of the change is real and how much is monetary. The information content of every individual price is degraded.

This is a different problem than the one we have when the price of tin moves. When the price of tin moves, the rest of the economy adjusts. When the price of money moves, the rest of the economy is repriced. There is no equivalent adjustment, because there is no other unit to flee to. Every other price in the economy is denominated in the same unit that has just moved. To repeat the formulation from the previous chapter: corrupting the unit corrupts every signal in the system, simultaneously.

Hayek understood this, but he did not fully apply his own insight to money in his early work. It took him decades to come around to the position that money itself ought to be exposed to market discipline rather than monopolized by the state, and the book in which he made the argument most clearly — *Denationalisation of Money* (1976) — was published when he was seventy-seven years old. He titled the central proposal of that book the "concurrent currencies" idea: let private issuers compete to provide money, let users choose the unit they preferred, let the market select the best unit through use. He understood that as long as a single state-issued money was the only unit available, no market discipline could discover the right one, because there was no alternative against which to measure it.

What Hayek imagined as competing private fiat issuers, Bitcoin instantiates as a single non-state issuer with credibly fixed supply. The structural achievement is the same: a money whose properties are not determined by deliberation. Whether this counts as one of Hayek's competing currencies, or as something stronger — a single money whose rules are constitutional rather than legislative — is a question I will come back to in Chapter 8. For now, the point is that Bitcoin is an instantiation, in working code, of a Hayekian critique that Hayek himself only got to articulate near the end of his life.

There is a second thread to pull on, which makes the case sharper. The most important price in any economy is not the price of any

particular good. It is the price of money over time — that is, the interest rate. The interest rate is the price at which the demand for borrowable funds today is equilibrated against the supply of savings deferred from today. It is the price that determines which long-duration projects get funded and which do not. It is the price that coordinates the time structure of production with the time preferences of savers. Every other price in the economy depends on it, because every other price is set by actors who are weighing present and future, and the rate at which they trade off present against future is the interest rate.

In the Austrian tradition, this is called intertemporal coordination, and the argument that running an economy without market-discovered interest rates produces malinvestment is the core of the Austrian theory of the business cycle. When central banks set the price of credit below the rate at which savers would voluntarily provide it, capital gets allocated to projects that are profitable only at the artificially low rate. When the rate eventually has to rise — because the underlying scarcity of real savings reasserts itself — those projects fail. The recession is not the cause of the malinvestment; it is the discovery of malinvestment that already happened. This is what Mises and Hayek were saying about the 1920s in real time, well before 1929 — and it might be something to consider today, with negative real interest rates for much of the past decade, record stock market valuations, and soaring real estate values.

The mechanism does not require subscribing to the entire Austrian apparatus. It only requires accepting that interest rates are prices, that prices coordinate behavior, and that prices set by deliberation rather than discovery contain less information than prices set by discovery. If you accept Hayek on tin, you should accept this on rates.

Here is where I want to make the connection that the rest of the book will turn on. Once you have accepted that the unit of account ought to be exogenous to political deliberation, and once you have accepted that

interest rates ought to be discovered rather than set, the question becomes: how would you ever get a discovered interest rate again? You cannot get one inside the existing system, because the existing system has, by deliberate design, made the central bank the marginal setter of the rate. You cannot get one in fiat-denominated lending markets, because those markets are downstream of the central bank's chosen rate; the LIBOR successor rates, the Treasury curve, every observable benchmark traces back to a policy rate that was deliberately set.

You can only get a discovered interest rate in a market where neither side of the transaction is subject to a central bank. That is, you can only get it in a credit market denominated in a money that is exogenous to any central bank. Bitcoin-denominated credit, or credit collateralized in Bitcoin and discovered in a market that does not look to any monetary authority for guidance, is the only place where the question "what is the natural rate of interest" can even begin to be answered empirically.

I will return to this in Part IV. For now I want you to hold the thought: a market-discovered interest rate, on a sound-money base, is the missing primitive that the modern monetary order has not produced for fifty years. Bitcoin-native credit infrastructure is what makes its discovery possible. That is not just a financial product opportunity. That is the operationalization of Hayek's argument in working code.

What Bitcoin-native credit operators are building is, in some sense, the most boring possible thing in financial markets: writing loans against collateral, on terms agreed by the parties. It is also, viewed against the background of the last fifty years of monetary history, one of the most consequential things being built anywhere. It is the quiet reconstruction of the price-discovery mechanism for capital itself.

I will make that case in detail in Part IV. The work of the next three chapters is to show, with data, what fifty years of running an economy without that price-discovery mechanism has produced. The Cantillon

effect, the productivity capture, and the international mirror are three views of the same phenomenon: an economy that has been allocating capital, for half a century, against a corrupted set of signals.

CHAPTER THREE

The Cantillon Engine

✦

Richard Cantillon was a French banker, born in Ireland, who made and lost a fortune in the speculative manias of the early eighteenth century and was murdered in his London townhouse in 1734, possibly by a disgruntled servant, possibly by an ex-creditor; the historical record is unclear. He left behind a single major manuscript, the *Essai sur la Nature du Commerce en Général*, which was published posthumously and which contains, among many other valuable observations, the original statement of what we now call the Cantillon effect.

The observation is simple. New money introduced into an economy does not enter everywhere at once. It enters at a particular point — through a particular channel, in the hands of particular people. Those who receive the new money first can spend it at the prevailing prices, before the new money has had time to bid up those prices. Those who receive the new money later, after it has worked its way through the economy and prices have risen, find that their money buys less than it used to. The first recipients capture purchasing power; the last recipients lose it. The transfer is invisible, because no one is robbed in any visible sense, and because the prices on the supermarket shelf may not change immediately or uniformly. But the transfer is real, and it is structural to the way new money enters any monetary system.

In Cantillon's day, the new money entered through whoever was nearest to the goldsmith's vault or the king's mint. In our day, the new

money enters through a more elaborate but structurally similar channel. The Federal Reserve creates reserves by buying securities — usually Treasury bonds and, in the post-2008 era, mortgage-backed securities — from a small group of primary dealers. The primary dealers receive newly created reserves in exchange for the securities. Those reserves enter the financial system at the apex of the credit hierarchy, in the hands of the largest banks, whose customers are by and large the largest holders of financial assets. From there, the new liquidity diffuses outward through the financial system, bidding up the prices of assets along the way before, eventually, working its way into goods and services prices.

This is not a conspiracy. It's a feature not a bug. If you were going to design a system in which new money entered at the apex of the wealth distribution and benefited those closest to the spigot first, you could not design one more efficient than the one we have. There is no smoke-filled room. There is just a transmission mechanism, operating exactly as designed.

If a discretionary monetary expansion is going to occur at all, the most defensible distribution would arguably be a direct airdrop to every participant's account in equal proportion, which would at least eliminate the Cantillon-effect transfer of purchasing power from the financially distant to the financially proximate. That this is not how monetary expansion is actually conducted tells you something about whose interests the existing system is structured to serve.

The data are the data. From the trough of the financial crisis in early 2009 through the end of 2024, the M2 money supply in the United States expanded by roughly 3.5x. The S&P 500, in nominal terms, expanded by roughly 9x from its March 2009 low. The Case-Shiller national home price index roughly doubled. Corporate earnings, measured by trailing operating earnings of the S&P 500, expanded by roughly 3.5x. Real

median household income, by the Census Bureau's measurement, grew by roughly 15 to 20 percent over the same period.

Hold those numbers in mind. M2 up 3.5x. S&P up 9x. Median wages up roughly 0.2x. The same monetary expansion that more-than-tripled the money supply only modestly improved the position of the median wage earner, while it nearly decupled the position of the asset holder. That gap is not noise. That gap is the Cantillon effect made visible. The new money entered through the asset markets, bid up asset prices, and only weakly transmitted to wages and goods. The dividend went somewhere. It went to whoever was closest to the spigot.

This is also the answer to the question that vexes mainstream macroeconomics: where did the inflation go? For the first decade after the financial crisis, the Federal Reserve expanded its balance sheet by trillions of dollars, and the consumer price index barely budged. The traditional monetarist prediction — that an expansion of the monetary base of that magnitude would produce a comparable expansion in goods prices — failed. A great deal of ink was spilled trying to explain why. The simple answer, hiding in the data, is that the inflation did not fail to occur. It occurred in assets. The S&P, the Case-Shiller, art prices, collectibles, fine wine, vintage cars, urban real estate — everywhere you looked at the asset class consumed by people closest to the credit spigot, you saw the inflation. The CPI basket — which weights heavily toward goods consumed by ordinary households, and which has had its measurement methodology adjusted multiple times over the past forty years — did not capture it.

The case I want to make is not that the CPI is fraudulent, although there are reasonable arguments to that effect that I do not need to litigate here. The case is that the CPI is the wrong instrument for measuring monetary debasement. Monetary debasement, in a Cantillon-effect economy, manifests not as uniform price increases across all goods and

services but as differential price increases concentrated where the new money first enters. The right measure of debasement is, in my view, the growth of the money supply itself, not the price-index proxy for that growth that the official statistics emphasize. To measure debasement by looking at the CPI is like trying to measure the rainfall in a city by sampling only the storm drains. You will find some water, but you will miss most of the volume.

The right way to measure monetary debasement in a Cantillon-effect economy is to look at the gap between what the financially proximate own and what the wage-earning median earns. That gap is the size of the debasement, measured directly in distributional terms. By that measure, the debasement of the past fifteen years has been catastrophic. The wealth share held by the top one percent has expanded; the wealth share held by the bottom fifty percent has contracted. Those who owned assets in 2009 saw those assets multiply in value. Those who did not own assets in 2009 — disproportionately younger, disproportionately working-class, disproportionately the people who would have been first-time home buyers in normal times — found themselves locked out of an asset market that was being levitated by monetary policy faster than their wages could catch up.

The political consequences of this are now visible everywhere, in both the developed and developing world. The stress on social cohesion, the rise of various forms of populism on both the left and the right, the collapse of trust in institutions, the visible disengagement of large parts of the population from the legitimacy of the existing order — all of this has many causes, but I am persuaded that the structural distributional consequence of the post-1971 monetary regime is one of the largest of those causes — in my view, the root cause and by far the most important.

People know they are being cheated. They can feel it. They cannot articulate exactly how, because the mechanism is invisible. But they know.

I do not write about politics in this book, and I am not going to start here. I mention the political consequences only to note that the Cantillon effect is not an arcane concept relevant only to monetary economists. It is the mechanism by which, over decades, the social contract is hollowed out. If you want to understand why politics in the developed world has become unmanageable, you have to understand the monetary arithmetic underneath the politics.

Bitcoin's relevance to this is twofold, and the second part is the one most often missed. The first part is the obvious part. A fixed-supply asset that is exogenous to the monetary regime offers a hedge. If the issuance rate of dollars is uncertain and the issuance rate of Bitcoin is fixed, then over a long enough horizon Bitcoin should appreciate in dollar terms by approximately the differential in issuance, adjusted for changes in demand. This is the price-target story most Bitcoiners are familiar with, and I do not have much to add to it. It is correct but not particularly interesting on its own.

The second part is more interesting. A fixed-supply asset, denominated in itself, makes the Cantillon effect legible in a way that fiat denomination cannot. As long as you measure your wealth in the same unit that is being debased, the debasement is invisible to you. Your wage is rising, your house is rising, your portfolio is rising, and you experience this as your getting wealthier — even when, measured against a non-debasing unit, you are running in place or losing ground. Bitcoin, used as a unit of measurement rather than as an investment, is the instrument that makes the debasement legible. When you measure your wealth in Bitcoin terms, you can see — for the first time in fifty

years — what monetary debasement is actually doing to your purchasing power relative to a constant.

I want to be cautious with one part of this argument. For the first fifteen years of Bitcoin's existence, much of the price action has been driven by adoption and network-effect dynamics from a very small base, rather than purely by the mechanical relationship between fiat issuance and Bitcoin's fixed supply. The two effects compound, but they are not the same effect, and serious analysis should not conflate them. I expect the network-effect contribution to diminish in importance as the base grows. The historical returns cannot be fully trusted as a guide to future returns, for a reason that is mechanical rather than speculative: they have been earned off a tiny base. A move from a market capitalization of $10 million to $100 million is a tenfold return that absorbs $90 million of incremental capital. The same tenfold return from a $1 trillion base requires absorbing $9 trillion of incremental capital. The percentage returns achievable from a small base are not the percentage returns achievable from a large one, regardless of how favorable the underlying monetary thesis remains. This is a Kelly-input adjustment any honest practitioner has to make: the implied edge from historical CAGR has to be discounted to reflect the much larger base future returns will be earned from. The discount is not a haircut for skepticism. It is a haircut for arithmetic. Strip out the small-base contribution and the pure monetary-anchor mechanic remains as the dominant driver of Bitcoin's price relative to fiat over longer horizons — but the magnitude of the future move, in percentage terms, will be smaller than the historical move was.

This is what Saifedean has called "the unit of account function," and it is the function that, over time, will matter most. Bitcoin as a store of value is a transitional argument; Bitcoin as a unit of account is the destination. When you can see your wages, your assets, your debts, and

your savings in a unit whose supply is not under deliberation, you begin to make different decisions. You stop confusing nominal appreciation with real wealth creation. You begin to allocate against a true measuring stick rather than a sliding one. That is what the rest of this book is about: how to think about portfolio allocation when you have, for the first time in your professional lifetime, access to a measuring stick that does not move.

But before we get to the measuring stick, I want to spend two more chapters on the diagnosis. The next chapter will look at the productivity story, which is the most damning evidence in the entire fiat-era ledger. And the chapter after that will look at the international mirror, which shows that this is not an American story.

CHAPTER FOUR

Productivity, Captured

✦

Between 1947 and 1971, output per hour worked in the United States grew at an average annual rate of approximately 2.8 percent, and the real median compensation of nonsupervisory workers grew at an average annual rate of approximately 2.5 percent. The two series tracked each other almost perfectly. When the country got more productive, workers got more prosperous. The relationship was so close that, plotted on a single chart, the two lines are barely distinguishable from each other.

In 1971, President Nixon closed the gold window, ending the last vestige of the Bretton Woods system and inaugurating what is sometimes called, with some euphemism, the “post-Bretton Woods era” but is more honestly called the era of pure fiat money. From 1971 through the present, U.S. labor productivity has continued to grow, more slowly than in the prior period but still substantially. Total productivity gains since 1971 are on the order of 140 percent. Real median compensation over the same period has grown by something on the order of 15 to 20 percent.

Two and a half lines on a chart now run wildly apart. Productivity continues to climb. Wages flatten. The wedge between them — the productivity dividend that, in the previous era, accrued to workers in the form of higher real wages — went somewhere else after 1971. It did not disappear. It was captured.

This chart is one of the most reproduced charts in modern American economic history, and it is also one of the most contested. A great deal of econometric labor has been done to explain away the divergence, to argue that the dividend really did flow to workers if you only adjust for benefits, or use the right deflator, or restrict the comparison to the right subset of the labor force. Some of those arguments have merit and modestly narrow the gap. None of them eliminate it. After every reasonable adjustment, a substantial wedge remains. Workers in the post-1971 economy did not capture the productivity gains of their own work to anything like the extent that workers in the prior era did.

Where did the dividend go? It went to the financial layer of the economy. Asset prices, financial sector profits, and the wealth share of the top decile all expanded substantially over the same period. Corporate profits as a share of GDP rose from a long-run average around 6 percent to peaks above 12 percent in the post-2008 period. The financial sector, which constituted about 4 percent of GDP in 1971, now constitutes between 7 and 8 percent. The total market capitalization of U.S. equities went from approximately 60 percent of GDP in 1971 to peaks above 200 percent in the post-2020 period. The household wealth held by the top 10 percent of households rose from about 65 percent of total household wealth in the early 1980s to about 75 percent today, with the top 1 percent's share rising from about 23 percent to about 31 percent over the same period.

Each of these numbers has some controversy attached to it, and I am not going to pretend the data are clean. But the directional story is not in dispute. The productivity dividend, in the post-1971 economy, has flowed disproportionately to the asset-holding layer of the economy, and within the asset-holding layer, disproportionately to the apex.

The conventional macroeconomic story for why this happened invokes globalization, technology, the decline of unions, and various

policy choices around taxation, trade, and labor. Each of those factors is real and contributed to the outcome. I do not dismiss any of them. But I want to put a different factor on the table, which I think the conventional story understates: the role of the monetary regime itself.

Consider what happened to credit availability and asset prices after 1971. With the gold constraint removed, the dollar issuance rate became a discretionary variable. Credit expanded much faster than the underlying economy could productively absorb. The total credit market debt to GDP ratio in the United States went from approximately 150 percent in 1971 to approximately 350 percent today. Federal debt alone went from about 33 percent of GDP in 1971 to about 120 percent today. Household debt, corporate debt, and the financial-sector liabilities all expanded by similar multiples. The monetary regime made credit cheap, abundant, and continuously available. Asset prices, which are determined in part by the cost and availability of credit, rose accordingly.

Wages do not rise the same way. Wages are set in labor markets, which are far less responsive to monetary policy in the short and medium run. A worker bidding for higher pay is bidding against other workers, against employers' productivity and pricing power, against the threat of substitution and offshoring. A buyer of a financial asset is bidding against other buyers, with leveraged credit available at policy rates that have, for most of the past forty years, been below the rate of asset-price appreciation. The two markets respond to monetary expansion on radically different timescales and at radically different elasticities. Monetary expansion lifts asset prices in real time; it lifts wages, if at all, only after long lags and only partially.

The result, compounded over fifty years, is the chart that does not lie. Productivity grew. Wages stagnated. Assets compounded. The dividend was captured.

This is what I think we should mean when we say the post-1971 monetary regime "produced inflation." It did not produce uniform inflation in the CPI sense, particularly not after the early 1980s when the Fed got serious about its mandate to control measured CPI. It produced a kind of differential inflation that we have not had a clean name for, in which the prices of the things owned by the asset-holding class rose much faster than the prices of the things consumed by the wage-earning class. A worker in 2024 buying the same basket of groceries he bought in 1971 pays roughly six times as much. A worker in 2024 buying the same house he could have bought in 1971 pays something like fifteen times as much. A worker in 2024 buying the same share of the same company he could have bought in 1971 pays something like fifty times as much. These are wildly different rates of inflation, and they correspond to wildly different rates of asset and wage appreciation across the income distribution. The monetary regime produced this differential. It is a policy outcome, not an unpredictable accident of nature.

Bitcoin's contribution to this story is simple and, I think, profound. By providing a fixed-supply unit of account, Bitcoin allows the differential inflation of the past fifty years to be measured against a single non-moving denominator. When you price an S&P 500 share in Bitcoin, or a house in Bitcoin, or a barrel of oil in Bitcoin, you see the asset's price approximately collapse over a long enough time horizon. The asset is not actually deflating in real terms; the dollar in which it is normally priced is debasing. Conversely, when you price wages in Bitcoin, the picture is much grimmer than the conventional wage-stagnation story already implies. Real wages, measured against Bitcoin, have been catastrophically debased. The wage-earning class has, in Bitcoin terms, become poorer at a rate that is concealed by the fiat denomination.

Investors understand that their stock portfolio rose three hundred percent over fifteen years. They have not previously had a clean way to ask: rose against what? When they price the same stock portfolio in Bitcoin, they see that, in the unit that is not being debased, their portfolio is roughly flat or modestly negative. They are not getting wealthier. They are running in place against a debasing currency. The nominal appreciation is a ratio, and they have been looking at only one half of the ratio.

The next chapter shows that this is not just an American story. It is a global one, and the pattern in some of the more dramatically debased currencies is so clean that it functions as a natural experiment. Then we will move on to the response.

CHAPTER FIVE

The Recent Mirror

✦

You do not have to fly to Istanbul. You do not have to read about Argentina or Lebanon. The same mechanism I have been describing across the previous four chapters just finished playing out in your own kitchen, in your own grocery store, in the listing price of the house down the street, and in the statement balance of every retirement account you own. The COVID-era monetary and fiscal expansion of 2020 through 2022 was the largest, fastest, and best-documented monetary intervention in American history. It compressed into roughly two years a Cantillon mechanism that had previously taken decades to manifest, and it did so in conditions where everyone who lived through it could see it happening in real time. If you want to know what the analysis in the previous four chapters actually feels like in lived experience, you do not have to look abroad. You only have to remember.

Consider the numbers, because they matter. In February of 2020, the M2 money supply in the United States stood at approximately $15.4 trillion. By April of 2022, M2 had reached approximately $21.7 trillion. The expansion was on the order of 40 percent in two years, an increase that, in the entire previous history of M2 measurement, had no comparable peacetime episode. The expansion was not merely the Fed's doing. The federal government, through the CARES Act of March 2020 and successive fiscal packages, deployed approximately $5 trillion in stimulus, much of it delivered as direct payments to households and businesses, and most of it financed by Treasury issuance that the Federal

Reserve effectively monetized through its asset purchase programs. The combined monetary and fiscal expansion was the textbook conditions for a Cantillon-mode debasement, executed in real time in front of a population that, for the first time in living memory, was actually paying attention to monetary policy.

What happened next was exactly what the framework predicts.

Asset prices reacted first, as they do under Cantillon mechanics, because the new money entered through the financial system and began bidding up financial assets before it could work its way through to goods and services. The S&P 500, which bottomed below 2,200 in March of 2020, was above 4,000 by April of 2021 and above 4,800 by the end of that year — nearly doubling from the trough in less than two years. The Case-Shiller national home price index, which had been growing at roughly 4 percent annually before the pandemic, accelerated to peak year-over-year growth of more than 20 percent in 2021, an unprecedented rate in the history of the series. Bitcoin, which had begun the cycle around $7,000, peaked above $69,000. Collectible markets, art markets, used cars, even baseball cards — every asset class that an asset-holding household could plausibly own experienced rapid and substantial price appreciation. The Cantillon effect was operating in plain sight, in real time, in the most data-rich monetary expansion ever recorded.

Then, with the lag the framework predicts, the new money began to spill into goods and services prices. CPI inflation, which had been running below 2 percent before the pandemic, accelerated through 2021, reached 7 percent by year-end, and peaked at 9.1 percent in June of 2022 — the highest reading since 1981. The price of eggs, gasoline, used cars, rent, and a hundred other items in the daily basket of ordinary households rose at rates that those households experienced as obvious and painful, regardless of how the official statistics characterized them. The wage-earning class, whose nominal wages were also rising but at a

substantially slower rate than the goods they had to buy, experienced a real-wage decline that the Fed's own measures eventually confirmed.

The political response was visible in real time and is still being processed. Households watched their grocery bills climb while their paycheck did not. They watched the equity in their homes, if they owned a home, increase by amounts that would have seemed implausible before. They watched their stock-market accounts, if they had them, double. They watched first-time buyers locked out of a housing market that was being levitated faster than wages could compete. And they watched the wealth distribution, already badly skewed, accelerate further toward the top. The COVID-era expansion was, in distributional terms, a transfer of purchasing power from wage earners and savers to asset holders and the financially proximate, executed at a scale and speed that has produced legible second-order political consequences across the developed world.

Then came the Fed's response. Faced with inflation it had insisted for most of 2021 was "transitory," the Fed eventually reversed course and began the most aggressive rate-hiking cycle in forty years. The federal funds rate went from effectively zero in March of 2022 to above 5 percent by mid-2023. The bond market reacted as the duration math required: long-dated Treasuries had one of the worst years in the history of fixed income in 2022, with the long bond losing more than 30 percent of its value. The 60/40 portfolio, supposedly the safe default, had its worst year on record because both halves declined simultaneously — a possibility I will return to in Chapter 10. The Fed's tightening achieved a partial resolution of the goods-price inflation, but it did not, and could not, undo the asset-price inflation that the prior expansion had produced. Houses did not return to 2019 prices. Equities did not return to 2019 prices. The Cantillon-mode purchasing-power transfer was, by and large,

locked in. Only the goods-price symptom was treated; the structural distributional consequence remained.

This is the United States in 2020 through 2024. This is not Turkey. This is not Argentina. This is the most observed monetary expansion in human history, in a developed economy with deep institutions and currency credibility, executed by people who genuinely believed they were doing the right thing, and the result was a textbook Cantillon-mode wealth transfer compressed into 24 months. Anyone who lived through it has the lived experience necessary to evaluate the framework I have laid out in the previous four chapters. The framework is not abstract. It is the explanation for what just happened to you.

The international evidence, which I want to discuss more briefly than I would have if the COVID-era episode were not so vivid, simply shows that the same mechanism operates in starker form in countries where the monetary regime is less restrained.

Turkey, between 2003 and 2024, lost approximately 99 percent of the lira's value against the U.S. dollar. The BIST 100 stock index, denominated in lira, rose by roughly 10,000 percent over the same period; in dollar terms it was roughly flat. Turkish equity holders captured nominal gains that did not represent any real wealth creation; they were merely tracking the debasement of the unit they were measured in. Turkish wage earners had no equivalent escape and watched their real purchasing power collapse.

Argentina has been doing some version of this for decades, with periodic redenominations of the peso wiping out successive nine-zero blocks of value. Argentine real estate is conventionally quoted in U.S. dollars rather than pesos, because no serious transaction will accept peso quotes for the relevant durations — the local currency has been demonetized for serious purposes by its own users.

Lebanon, in 2019 and 2020, demonstrated the discontinuous version of currency failure. The pound, pegged to the dollar at 1,500 since 1997, broke and went into freefall, eventually trading at a fraction of one percent of its previous value. Depositors who had trusted the Lebanese banking system lost most of their savings. Those who held real estate, gold, or by the end Bitcoin, preserved purchasing power.

Japan is the inverse case, and it deserves more attention than the simple "money printing equals inflation" thesis allows. Japan, since 1990, has run by any reasonable measure the most aggressive monetary policy of any major central bank for the longest period. The Bank of Japan now owns roughly 7 percent of the Japanese equity market through its ETF purchases. Yet Japan has experienced low or zero CPI inflation for thirty years. The framework I have laid out predicts this outcome correctly: the Cantillon mechanism in Japan has played out almost entirely in the asset and credit channels and almost not at all in the goods channel. Japanese asset prices, particularly the Nikkei since 2012, have responded exactly as the framework predicts. Japanese real wages have stagnated for decades. The wealth distribution outcomes are similar to what we see elsewhere; only the price index looks different. Japan is not a refutation of the framework. Japan is a case where the framework manifests in pure form, with the goods-price symptom suppressed by demographic and structural factors that do not apply elsewhere.

The international evidence collectively confirms what the COVID-era U.S. evidence already demonstrated. The Hayekian knowledge problem applied to money produces, in every monetary regime where central banks operate, some combination of asset-price inflation, real-wage stagnation, capital misallocation, and political fraying. The exact tempo and the exact mix vary by country and by episode. The mechanism does not. It is universal, because the underlying problem — discre-

tionary setting of monetary aggregates by committees that cannot in principle possess the knowledge required to set them well — is universal.

This is the diagnosis. It is not a story about emerging markets. It is the story of every country with a fiat currency, told at different speeds. The United States version of the story, between March of 2020 and the present, is the cleanest and most data-rich version that has ever been observed. It has happened in your lifetime. You do not need to take it on the authority of any economist. You only have to remember the past five years.

The response, then, has to be a unit of account that is not an output of policy deliberation. That is what Part II is about.

PART II

The Anchor

CHAPTER SIX

The Exogenous Unit

✦

The diagnosis in Part I was negative. The remainder of this book is positive. I want to begin the positive case by stating, as plainly as I can, what kind of monetary asset is required to escape the failure mode the diagnosis describes.

The asset must have a supply that is not subject to political deliberation. Not merely "limited" supply, not "scarce" supply, not "predictable" supply. Exogenous. The supply rule must be outside the system of variables that any political body can adjust through any normal procedure. If the supply rule can be changed by a central bank board, by a legislature, by a constitutional convention, by a treaty, or by any other deliberative body acting through a deliberative process, then the supply rule is endogenous to politics, and the asset is structurally subject to the same failure mode that fiat currencies are subject to. The failure may be slower, the deliberation may be more constrained, but the structural defect is present.

This requirement is more demanding than it first appears. Almost no asset humanity has used as money meets it. Gold's supply was constrained by mining technology rather than by deliberation, which is closer to the requirement, but the monetary use of gold has been repeatedly suspended, modified, and abandoned by political decision: the U.S. abandonment of gold convertibility in 1933 and again in 1971; the British abandonment in 1931; the various monetary reforms of the nineteenth century; the demonetization of silver. Even when gold's

supply was not subject to political control, gold's role as money was. The gold standard was a policy choice, and it could be — and was — unchosen.

A central bank's commitment to a fixed money supply rule could in theory satisfy the requirement, but in practice every such commitment has been broken by every central bank that has ever adopted one. Milton Friedman's k-percent rule, the Maastricht criteria, the various inflation targets, the gold standard itself — every monetary rule promulgated by a deliberative body has been suspended or modified when the deliberative body found the rule inconvenient. This is not a critique of the people who suspended the rules other than to say maybe they didn't really understand the importance. It is the structural problem with rules that depend on a deliberative body's continued willingness to abide by them. The rule is only as strong as the body's willingness, and the willingness is itself an output of deliberation.

Bitcoin is the first asset in human history whose monetary properties are not contingent on the continued willingness of any deliberative body to maintain them. The supply schedule is enforced by the protocol, the protocol is enforced by the network, and the network has no central operator who can change the rules. A change to the supply schedule would require a coordinated action by a global, decentralized population of node operators and miners, the vast majority of whom have powerful economic incentives not to ratify any such change. The supply schedule is, for all practical purposes, exogenous to politics in the strong sense the requirement demands.

I want to be careful here, because there is a long literature, much of it produced by Bitcoin's critics, suggesting that this exogeneity is an illusion — that "the protocol can be changed by hard fork," that "miners or developers could collude," that "regulatory action could compromise it." These claims need to be taken seriously rather than dismissed. The

honest answer is that Bitcoin's exogeneity is a function of the social and technical equilibrium around the protocol, not of any single inviolable feature. The 21 million cap could in principle be changed; the question is whether such a change could ever achieve the network consensus required to be ratified, given that ratification would require the assent of a population whose holdings would be diluted by the change. The answer, after sixteen years of operation, is that no such change has been ratified, multiple proposed changes far less consequential than a supply increase have failed to achieve consensus, and the social equilibrium around the 21 million cap appears to be one of the strongest commitments in the network's culture. Bitcoin's supply is exogenous to politics not because it cannot be changed in some narrow technical sense, but because the social and economic equilibrium around the supply rule is too strong for any normal political process to overturn.

This is, in its way, the strongest possible kind of exogeneity. A rule that is mathematically immutable in some narrow technical sense but socially fragile would offer false comfort. A rule that is technically mutable but socially adamant offers real assurance, because it has been tested by sixteen years of attempted modification and has not budged. Bitcoin is the latter. The 21 million cap has survived attempts to change the protocol's most basic parameters, has survived multiple regulatory crackdowns in major jurisdictions, has survived the collapse of essentially every centralized exchange and lending counterparty, has survived its own price collapses, has survived an ongoing campaign of academic and journalistic delegitimization, and remains intact. There is no equivalent monetary commitment in human history that has survived equivalent stress.

Once you have an exogenous unit of supply, several things follow.

First, the unit becomes a credible long-horizon store of purchasing power, not because its purchasing power is constant in the short run but

because there is no political mechanism to dilute it. Bitcoin's price will fluctuate. It has fluctuated dramatically in the past and will continue to. But the fluctuations occur against a fixed-supply backdrop, which means that on a long enough horizon, the integral of price action is integrated against a denominator that is not being expanded. This is structurally different from any fiat currency, where the long-horizon integral is being calculated against a denominator that is itself expanding at a rate that no holder can verify or control.

Second, the unit becomes the natural numeraire for measuring debasement in other currencies. I made this point in Chapter 4, and I want to develop it slightly further here. The most useful chart in modern monetary economics, in my view, is not the price of Bitcoin in dollars. It is the price of any asset, denominated in Bitcoin, plotted against time. When you do this exercise systematically — pricing the S&P 500 in Bitcoin, pricing gold in Bitcoin, pricing real estate in Bitcoin, pricing labor in Bitcoin — you discover that almost every asset class on Earth has been deflating against Bitcoin for the entirety of Bitcoin's existence — though, again, the small-base caveat from Chapter 3 applies, and this argument should be made with appropriate caution about the underlying drivers. This is not, primarily, because the assets are losing value in any real sense. It is because Bitcoin is the only major asset whose denominator is not expanding. The deflation of all other assets against Bitcoin is the visible measurement of monetary debasement in the unit each asset is conventionally priced in.

Third, the unit becomes the basis for credit and capital structures that are not subject to the failure modes of fiat-denominated credit. This is the argument that Part IV will develop in detail. For now, the point is that an exogenous unit of money is the necessary precondition for an exogenous unit of credit, and an exogenous unit of credit is the necessary precondition for a yield curve that is not under the control of any

monetary authority. The exogeneity of the unit propagates upward into the entire capital structure that can be built on top of it.

The remaining chapters of Part II address the obvious objections. Chapter 7 addresses the question of why gold, which has done a version of this job before, cannot do it again. Chapter 8 addresses the philosophical question of whether rules-based money is, on reflection, what we want at all, given that even Hayek toward the end of his life favored competitive private monies over a single rules-based money. Chapter 9 addresses the practical question of what changes for an investor when his unit of account is no longer subject to political adjustment.

If you accept the diagnosis from Part I, the requirement for an exogenous unit follows almost automatically. The remaining question is whether Bitcoin in particular meets the requirement, or whether something else might do the job better. I am persuaded that Bitcoin meets the requirement and that nothing else currently in existence does. The next chapter explains why.

CHAPTER SEVEN

The Evolution of Money

THE RIGHT WAY to think about money is evolutionary. Nick Szabo made this argument most clearly in his 2002 essay *Shelling Out: The Origins of Money*, in which he traced the origins of money back to what he called "collectibles" — proto-monetary objects like shells, beads, and worked stones that early humans accumulated, exchanged, and used as stores of value tens of thousands of years before the first coin was struck. Szabo's central insight is that money did not appear suddenly when civilization became complex enough to need it. Money emerged gradually, by selection, from objects whose properties — scarcity, durability, recognizability, divisibility, portability — happened to suit them for the function. Each subsequent monetary asset humanity adopted was a refinement on the prior one, solving the bottlenecks that had limited its predecessor.

Shells gave way to worked stones, which gave way to metals, which converged on gold and silver because of their superior monetary properties — denser, more durable, more divisible, more recognizable. Gold itself displaced silver in international monetary use over the nineteenth century because gold's higher value-to-weight ratio made it more practical for large-scale international settlement. The gold standard was, in turn, displaced by fiat money in the twentieth century, not because fiat money was monetarily superior — it manifestly is not — but because the political flexibility fiat offered was attractive to states that wanted to finance war and welfare without the discipline gold

imposed. Each transition in this sequence was, in Szabo's evolutionary frame, a response to a specific bottleneck in the prior monetary technology, and each transition solved the bottleneck while creating new ones.

Bitcoin, in this frame, is the next evolutionary step. Not because it is "magic internet money," not because it is a speculative asset that has performed well, but because it solves the bottlenecks that limited its predecessors while preserving the monetary properties that made gold the durable winner of millennia of selection. The case for Bitcoin over gold is not a case against gold's monetary properties, which remain excellent. It is a case that Bitcoin solves three specific deficiencies of gold that, taken together, prevent gold from being the monetary base of a digital, globally distributed economy.

The first deficiency is verification. Verification of physical gold by an end user is expensive and difficult. A retail holder cannot easily confirm that a gold bar is what it claims to be without specialized equipment and expertise. Counterfeiting — sometimes through tungsten cores in apparently genuine bars — is a recurring problem. The history of gold-backed monetary systems is in large part a history of issuers who claimed to hold the gold they were redeemable in but did not actually hold it. The paper-gold-to-physical-gold ratio in modern bullion markets is, by various estimates, somewhere between fifty and one hundred to one. Most "gold" exposure in the modern financial system is paper claims on gold rather than gold itself. Verification — by which I mean an end user's ability to confirm, without trust, that the asset he holds is the asset he believes he holds — is a property gold lacks at scale. Bitcoin is verifiable by any user, in real time, at zero marginal cost, by running a node and validating the chain. The verification is mathematical rather than physical. This is a property gold cannot have, because verification of physical gold inherently requires physical handling.

The second deficiency is portability and seizure resistance. Gold is heavy. A million dollars of gold is approximately fifteen kilograms; a billion dollars is fifteen tonnes. Moving large amounts of gold is logistically difficult, expensive, and visible. Storing large amounts requires physical infrastructure — vaults, security, insurance — that is itself centralized and politically vulnerable. The history of gold confiscations, most famously the 1933 U.S. confiscation under Executive Order 6102, is a history of governments seizing gold from custodians because the gold was physically locatable. An asset whose effective use requires custody by a small number of large counterparties is hostage to the political climate facing those counterparties. Bitcoin can be moved across any distance in minutes, at near-zero marginal cost, in any quantity. Bitcoin can be self-custodied with no physical infrastructure beyond a hardware device the size of a USB stick. The keys controlling any quantity of Bitcoin can be memorized, written down, distributed across multiple parties, or held in any number of cryptographic schemes that have no analogue in the physical world.

The third deficiency, and the one most underappreciated by the gold side of the debate, is programmability. Gold cannot enter into a smart contract. Gold cannot be the collateral in a cryptographically enforced loan. Gold cannot be moved conditionally on the outcome of an external event, paid out automatically against a price feed, or escrowed by mathematical proof. Every financial structure built on gold has historically required trusted intermediaries to enforce its conditions. The intermediaries are themselves the failure points of the system. Bitcoin can do all of these things. Discreet log contracts, which I will discuss at length in Part IV, allow two parties to enter into a contract whose payout depends on an externally observed outcome, with the contract enforced cryptographically rather than by any trusted intermediary. This is what

makes Bitcoin not just gold-on-the-internet but something structurally more capable than gold has ever been.

There is a thoughtful version of the gold argument, advanced most clearly by Lyn Alden, that frames Bitcoin and gold as complementary rather than competitive. Gold is a store of value with a five-thousand-year track record; Bitcoin is a store of value with a sixteen-year track record. Allocators with long horizons and risk tolerance for asymmetric outcomes might want both. The gold case is also the case that requires the least imaginative leap from an existing institutional allocator: gold is familiar, gold has been in central bank reserves for the entire modern era, gold's monetary thesis does not require believing anything about the future of digital systems. For any investor whose constraints rule out a meaningful Bitcoin allocation today, gold remains the next-best instrument for hedging monetary debasement. This is a position I have considerable sympathy with, and one I would defend against the maximalist objection that gold is somehow disqualified.

But the question this chapter is asking is not whether gold is good money. Gold is excellent money, and was for thousands of years. The question is whether gold is the right monetary base for the next century. I do not believe it is. The next monetary base will need to be digitally native, mathematically verifiable, programmable, and self-custodial at every scale. Bitcoin is the only candidate that is all of those things, and I am persuaded that no future candidate can be more than incrementally better, because Bitcoin is approaching the technical and social limit of what monetary properties an asset can possess.

Money evolved from shells to silver to gold to fiat. The next step in that sequence is digital, scarce, neutral, and verifiable. Bitcoin is the candidate that fits the evolutionary slot. The question is not whether the evolution will continue — the bottlenecks of fiat are too severe for any thoughtful observer to believe the current arrangement is stable — but

which candidate will occupy the next stage. I do not see another one with a credible claim.

The next chapter takes up the philosophical question that the previous two chapters have been circling: should the monetary base be a single thing, governed by inviolable rules, or should it be a competitive market of issuers? This is a question Hayek himself wrestled with throughout his life, and one I think Bitcoiners need to engage with seriously rather than assume the answer to.

CHAPTER EIGHT

Rules Over Discretion

✦

The most influential argument in modern monetary economics for binding rules in monetary policy was made by Milton Friedman, who argued for a fixed annual growth rate of the money supply — the famous "k-percent rule" — on the grounds that central bank discretion produced more harm than good, even when the central bankers were well-intentioned and competent. Friedman's argument was Hayekian in spirit even though Friedman himself was a monetarist rather than an Austrian. The argument was that the variability of discretionary policy contributed more noise to the economy than the policy contributed signal, and that a binding rule, even a suboptimal one, would produce better outcomes than the best-case discretionary policy because it would eliminate the noise.

A parallel argument has been made for centuries in the constitutional tradition, most powerfully in our era by James Buchanan, who founded a school of thought called constitutional political economy. Buchanan distinguished between the level of choice over rules and the level of choice within rules. At the constitutional level, we choose the rules under which we will subsequently make ordinary political and economic decisions. At the post-constitutional level, we make those ordinary decisions within the constraints of the rules we have already chosen.

Buchanan's claim was that many things which are properly decided at the constitutional level are improperly subjected to ordinary post-constitutional choice. Money is one of them. The rule by which money is

issued is a constitutional question, in Buchanan's sense — it sets the framework within which all other economic decisions are made. To leave it to ongoing post-constitutional discretion is to allow the framework itself to be continuously renegotiated by the actors who operate within it. This is, in Buchanan's view, a structural error in the design of modern monetary institutions, and it produces predictable failures regardless of which particular discretionary decisions are made within the bad framework.

Bitcoin, viewed through this Buchanan-style lens, is a constitutional monetary commitment that has been made unilaterally, by a population of voluntary participants, against the wishes of every constituted monetary authority. It is a kind of stateless constitutional money — money governed by binding rules that are enforced by the participants rather than by any external authority. The rules are not optimal in any obvious sense. They were chosen by Satoshi Nakamoto, a pseudonymous individual or group, in 2008 and 2009, and they have remained essentially unchanged since. Whether the supply schedule should asymptote to 21 million, whether the block time should be ten minutes, whether the difficulty adjustment should occur every 2,016 blocks — these are not choices that can be defended as optimal. They are choices that have become Schelling points: the values that everyone has settled on, that everyone expects everyone else to continue to settle on, and that have therefore become resistant to change because changing any of them would require coordinating an entire network around a new equilibrium.

This is, in a deep sense, what it means for a monetary rule to be constitutional. The rule does not have to be optimal. The rule has to be inviolable. The inviolability is the source of the rule's value, because it is the inviolability that prevents the rule from being subjected to ongoing post-constitutional negotiation by the actors who would benefit from changing it.

There is a counterargument from Hayek himself, which deserves an honest hearing. In *Denationalisation of Money*, Hayek did not argue for a single rules-based money. He argued for competitive private currencies — many different issuers, each producing their own money, with users choosing which money to use, and the market disciplining issuers who debased their currencies. Hayek's vision was not a single inviolable Bitcoin-like rule. It was a competitive market in monetary discipline.

There are good reasons to take Hayek's preference seriously. Single monies, even rules-based ones, have failure modes that competitive systems do not. A rules-based monetary system that turns out, in retrospect, to have a flaw in its rules — the wrong issuance schedule, the wrong governance assumptions, the wrong technical parameters — has no internal mechanism for correction. A competitive system, in which issuers can be displaced by better issuers, has built-in correction mechanisms that a single-issuer system lacks.

I think the Hayekian competitive-currency vision is theoretically attractive but has, in practice, lost the argument to the rules-based single-currency model for an empirical reason that Hayek could not have anticipated: network effects in monetary systems are extremely strong, and the costs of using multiple competing monies are significant enough that users tend to converge on a single dominant money even when alternatives are technically available. The history of money shows this convergence over and over: at any given time, in any given economic region, the overwhelming majority of monetary use is concentrated in a single asset, even when other monetary assets exist. Competitive currencies, at the limit, tend to produce a single winner.

If the network-effect dynamics push toward a single dominant money anyway, then the Hayekian vision of competitive monetary discipline is, in practice, a transitional argument rather than a stable equilibrium. The question becomes not which competing money will win but which type

of single money — discretionary fiat, gold-style commodity, or rules-based digital — is most likely to emerge as the dominant settler. I am persuaded, for reasons developed in the previous chapter, that the rules-based digital answer is the most stable of the three, and that Bitcoin specifically is the rules-based digital money most likely to occupy the position.

This means accepting a tradeoff. We are choosing inviolability over correctability. We are choosing the discipline of binding rules over the responsiveness of competitive markets. In Buchanan's framework, we are accepting that some decisions are properly removed from ordinary choice and lodged at the constitutional level, where they are stabilized against renegotiation.

I find this tradeoff acceptable, and I want to be clear about why. The history of discretionary monetary regimes is uniformly bad. Every fiat currency that has ever existed has either failed catastrophically or is currently in the process of failing more slowly. The history of rules-based monetary regimes — the gold standard, in its various incarnations — is much better, but the rules in those regimes were always vulnerable to political suspension, and they were always suspended when politics required it. The gold standard worked when it worked because political authorities chose to let it work; it failed when it failed because political authorities chose to suspend it. A rules-based monetary regime whose rules are not vulnerable to political suspension is something that has never previously existed in human history. Bitcoin is the first instantiation.

Whether the specific rules Bitcoin has chosen are optimal is an interesting question that I do not propose to answer, because I do not think it matters very much. What matters is that the rules are stable, that the stability has been demonstrated under stress, and that the stability removes a class of discretionary decisions from the political process

where those decisions have, historically, been made badly. We are giving up the option to make better discretionary monetary decisions in the future. We are gaining the certainty that we will not make worse ones.

For a portfolio allocator, this is a feature, not a bug. The job of an allocator is to manage capital under uncertainty over long time horizons. An asset whose monetary properties are guaranteed not to be modified by political deliberation, over the relevant time horizons, is an asset that simplifies the allocator's problem in a category-defining way. There is one less variable to forecast, one less risk to hedge, one less political contingency to handicap. The asset is what it says it is, indefinitely.

This is the value proposition that the financial planning case for Bitcoin ultimately rests on. We are not betting that Bitcoin will appreciate in any particular timeframe. We are recognizing that Bitcoin removes a category of risk — the discretionary monetary risk — from the portion of the portfolio allocated to it, and that this removal is itself worth a non-trivial allocation in any portfolio that is exposed to discretionary monetary risk anywhere else. Which is to say, in any portfolio that holds any other asset.

The self-sovereignty case for Bitcoin is also a financial-planning case in its own right, and I will return to it in Chapter 12. The ability to hold an asset whose properties cannot be modified by any institution is not just a libertarian-philosophical preference. It is a structural risk-management feature with substantial implications for how an allocator should think about the holding.

The next chapter develops this point in practical terms: what changes, for an investor, when his unit of account is no longer subject to political adjustment? The answer is more than most investors initially realize.

CHAPTER NINE

The Measuring Stick

✦

Imagine a world in which the standard meter — the literal physical object that defines the length of a meter — was modified, by international agreement, every several years. Imagine that the modifications were not announced in advance, that they were sometimes large and sometimes small, that they were made in response to political pressures rather than scientific ones, and that the direction of the modifications was always to make the meter slightly shorter. Imagine that you were a structural engineer in this world, designing buildings whose load-bearing requirements were specified in meters.

You would have a problem.

You could still do your job, in some sense. You could still design buildings. You could still hire contractors who used the same official meter you used, and the buildings would still stand. But you would lose, in a fundamental way, the ability to compare your designs across time. A building you designed twenty years ago, using the meter as it was then defined, would no longer be specified in the same units as a building you were designing today. Your comparisons across time would be confounded by the changing length of the unit of measurement. The longer the timescale of comparison, the less informative any measurement would be.

This is the situation we are in with money. Every fiat currency in the world has a definition that is, in effect, modified by political decision on a continuous basis. The dollar of 1971 and the dollar of today are not the

same unit, in any meaningful sense. They are nominally the same unit, in that they share a name and there is a continuous historical record connecting one to the other. But the purchasing power of one dollar, and therefore the meaning of any measurement made in dollars, has changed dramatically and in ways that are deliberately obscured by the issuer.

The result, in financial markets, is that comparisons across time made in fiat units are systematically misleading. A house that cost $50,000 in 1971 and costs $500,000 in 2024 has not "appreciated tenfold." It is, by most measures, the same house, with similar utility, similar location, similar physical characteristics. The unit it is being priced in has lost most of its value. The same house has not become ten times more valuable; the dollar has become roughly one-tenth as valuable. The conventional fiat-denominated measurement obscures this entirely. Investors look at long-term returns expressed in fiat and conclude that they have made remarkable wealth, when in fact they have, in many cases, merely tracked or modestly outpaced the debasement of the unit they were measuring against — and, in many cases, paid taxes on the nominal gains as a result.

Bitcoin offers something that no other asset in modern financial markets offers: a unit of measurement that does not move. It is the structural engineer's unmodified meter. When you measure your portfolio's value in Bitcoin terms, you are measuring against a denominator whose supply is not being expanded. The number you get is, for the first time in your investing life, an unconfounded measurement.

I want to give you a concrete exercise to make this real. Take any asset you own — a stock position, a real estate holding, a retirement account balance — and look at its history denominated in Bitcoin rather than in dollars. The chart will look very different from the chart you are used to. Most assets, charted in Bitcoin terms over the past decade, look like long sloping declines or sideways-to-down patterns. The S&P 500,

denominated in Bitcoin, has lost approximately 95 percent of its value over the past decade. Real estate, denominated in Bitcoin, has lost a similar amount. Gold, denominated in Bitcoin, has lost about 90 percent. Wages denominated in Bitcoin have lost approximately 99 percent.

These are not statements about how good Bitcoin is as an investment. They are statements about how badly the conventional units of account have debased over the period in question. The S&P has not actually lost 95 percent of its value; it has merely failed to keep pace with a non-debasing alternative. The chart of the S&P denominated in Bitcoin is the chart of the dollar's failure to maintain its purchasing power against the only major asset whose supply is not expanding.

There is a related argument worth noting briefly: the outperformance of companies that aggressively reduce share count through buybacks is, in part, a fixed-supply effect at the equity level — the same shares chasing growing earnings. The analogy to Bitcoin's monetary thesis is real but partial, and I will not press it here, because I want to avoid grounding too much of the case on Bitcoin's price performance off a small base. The structural argument about fixed supply matters; the historical performance is a noisier evidence channel.

When investors first see these charts, they often have a strong reaction. Some find them clarifying — a kind of "Oh" moment in which a great deal of confusing information about the past decade suddenly resolves. Others find them disorienting, because the picture is so unflattering to assets they have been told for their entire careers are good investments. Both reactions are honest. The picture is clarifying, and the picture is unflattering. Both things are true.

What these charts demonstrate, in a way no amount of macroeconomic argumentation can, is that the post-2008 monetary regime has produced enormous nominal wealth and very little real wealth. Investors who own assets are wealthier in dollar terms and approximately stable in

Bitcoin terms. Investors who hold cash are wealthier in dollar terms and dramatically poorer in Bitcoin terms. Workers whose only asset is their human capital, paid in fiat, are nominally stable and dramatically poorer in any non-debasing unit. The differential outcomes across these groups, charted in Bitcoin terms, are the cleanest possible visualization of the Cantillon effect at work.

There is a second use of the unmoving measuring stick that is, I think, even more important. It is the use of Bitcoin as the unit in which long-term financial planning is conducted. Most retirement planning, financial goal setting, and intergenerational wealth transfer is conducted in fiat units, with assumptions about returns, inflation, and purchasing power that are themselves expressed in fiat. The result is that the entire planning exercise is conducted in a unit whose meaning will have substantially changed by the time the plan is supposed to mature. A person who plans to retire on $5 million in thirty years is making a meaningless statement, because no one knows what $5 million will buy in thirty years, and the issuer of the dollar has every incentive to make sure $5 million buys substantially less than the planner expects.

A person who plans to retire on a target Bitcoin balance is making a statement that does not depend on anyone's monetary policy. A target of, say, ten Bitcoin in retirement specifies an actual quantity of monetary assets, with known and immutable supply characteristics, that will exist regardless of what any central bank does in the intervening decades. The plan is denominated in a unit that cannot be debased away. This is a property no fiat-denominated retirement plan has ever had, and it is one of the most consequential properties Bitcoin offers to long-horizon investors.

I will return to the planning implications of this in Part III, when I discuss portfolio construction in detail. For now, the point I want to leave you with is that the value of an unmoving measuring stick is not just a

matter of intellectual clarity. It is a matter of practical financial planning. An investor whose horizon is multi-decadal needs a unit of measurement that will still mean the same thing at the end of the horizon as it does at the beginning. There is no fiat currency that satisfies this requirement, and there has not been one for any extended period in modern history. Bitcoin satisfies it.

This is, in the end, the simplest way to state the case I have been building across Part II. Bitcoin is the response to the diagnosis from Part I because Bitcoin is the only asset whose unit of account is structurally exogenous to the political process that produces the failure mode the diagnosis describes. The exogeneity is what makes Bitcoin a measuring stick. The measuring stick is what allows long-horizon allocation to be conducted against a stable denominator. The stable denominator is what restores the missing precondition for serious long-term financial planning.

We are now ready to talk about how to actually build a portfolio against this measuring stick. That is the work of Part III.

PART III

The Allocation

CHAPTER TEN

The Death of 60/40

The 60/40 portfolio — sixty percent equities, forty percent bonds — has been the default model for serious institutional and retail allocation for more than half a century. It is the model embedded in target-date funds, in endowment baselines, in the textbooks every CFA candidate is required to internalize. It is the framework against which every alternative allocation has had to argue itself. And it is, as a default model for a long-horizon portfolio in the post-2020 monetary environment, dead.

Saying the 60/40 portfolio is dead is not a fashionable thing to say in mainstream financial planning circles. It cuts against an enormous amount of professional infrastructure, and it threatens a great deal of revenue tied to managing portfolios in the conventional way. I am not saying it for shock value. I am saying it because the assumptions on which the 60/40 model rests have, one by one, broken in ways that are not recoverable within the existing monetary regime.

The 60/40 model rests on three foundational assumptions. The first is that equities and bonds are negatively or weakly correlated, so that a balanced portfolio achieves diversification benefits and reduced drawdowns relative to either asset alone. The second is that bonds, particularly sovereign bonds, offer a positive real yield over a reasonable holding period, so that the bond portion of the portfolio is doing real work — generating return — in addition to its diversifying function. The third is that the central bank stands behind the system as a backstop, so

that in periods of severe market stress, the bond portion will appreciate while the equity portion declines, smoothing the overall portfolio path.

Each of these assumptions has broken.

The first assumption — equity-bond negative correlation — broke in 2022. The standard 60/40 portfolio had its worst year in modern history, with both stocks and bonds falling substantially in the same calendar year. This was not a small data point or a one-off anomaly. It was the consequence of the Federal Reserve being forced to choose between two horns of a dilemma it had created for itself: continue accommodative policy and risk runaway inflation, or tighten aggressively and crash both stock and bond markets simultaneously. The Fed chose the latter. The result was the simultaneous decline of both halves of the 60/40 portfolio, in exactly the period when the diversification benefit was supposed to be most valuable.

This is not an artifact of one year's policy. It is the structural consequence of what happens when a central bank's primary policy lever — the discount rate — has to be moved aggressively to address an inflation problem of its own making. In such periods, equities and bonds are positively correlated, both responding to the discount rate change in the same direction. The diversification benefit of the 60/40 portfolio depends on the absence of such periods. Such periods are, going forward, going to be more common rather than less common, because the underlying conditions that produced the 2022 episode have not been resolved.

The second assumption — that bonds offer a positive real yield — has been intermittently false for most of the post-2008 period. Real yields on developed-world sovereign debt have been negative or barely positive for the better part of fifteen years. Investors holding long-dated sovereign bonds during the period from approximately 2012 through 2022 were locking in negative real returns over the holding period, often

without fully recognizing it because the nominal yield was positive. The bond allocation of a 60/40 portfolio, in this period, was not generating real return; it was generating real loss, partially offset by the diversification benefit, with the diversification benefit also intermittently failing.

The third assumption — that the central bank provides a credible backstop — was the foundation of the entire post-2008 risk-asset rally. As long as investors believed the Fed would step in to support asset prices in any serious downturn, the implicit put under risk assets justified higher valuations and lower required returns across all asset classes. This belief is now meaningfully impaired. When the Fed has to choose between defending asset prices and defending the currency, the Fed will, eventually, defend the currency, because a complete loss of currency credibility produces consequences that the Fed cannot remediate. The 2021-2023 inflation episode was the first time in the post-Volcker era that this tradeoff became visible in real time, and the Fed's response — aggressive tightening despite asset price damage — confirmed that the put is not unconditional. Investors now have to handicap, for the first time in a generation, scenarios in which the Fed lets asset prices decline materially in the service of currency credibility.

If all three foundational assumptions of the 60/40 model are weakened or broken, what replaces it?

The honest answer is that there is no single replacement model that has yet achieved the consensus the 60/40 model has enjoyed. We are in a transitional period, and the new default has not yet emerged. But the directional logic of what the new default has to look like is clear.

It has to include an asset whose return is structurally uncorrelated with the central bank's policy choices, because the central bank's policy choices are now the dominant source of correlated risk across the rest of the portfolio. It has to include a fixed-income alternative whose return

does not depend on the issuer of the sovereign bond honoring his real obligations, because the issuer of the sovereign bond has explicitly indicated that he will not honor those real obligations when push comes to shove. It has to include exposure that protects against the specific failure mode the existing model is most exposed to: the simultaneous decline of equities and bonds in response to central bank tightening forced by inflation the central bank itself produced. The same dynamic, in a Bitcoin-anchored future, would manifest differently — credit conditions would tighten endogenously through the operation of a fixed-supply monetary base, rather than through the discretionary action of a central bank responding with its own priors to its own priors.

Bitcoin, in self-custody, is the only asset class I am aware of that satisfies all three conditions. Bitcoin's return is structurally uncorrelated with central bank policy choices in the long run, because Bitcoin's supply is not adjustable by any central bank. Bitcoin is not exposed to sovereign credit risk because no sovereign issues it. Bitcoin's correlation with traditional risk assets has been variable but has shown, in the most stressed monetary periods, the kind of negative-to-uncorrelated behavior that made bonds valuable in the 60/40 model when the model worked.

Bitcoin-collateralized credit, as I will develop in Chapter 15 and again in Part IV, is the natural replacement for the bond allocation of the 60/40 model. It offers yield that is determined by market discovery rather than by central bank manipulation. It offers credit risk that is structurally collateralized by an asset outside the fiat system. And it offers a fixed-income exposure whose performance does not require the sovereign issuer to honor real obligations he has every incentive to dishonor.

The new model that I think is most likely to emerge as the post-2020 default looks something like: a substantially reduced equity allocation (perhaps thirty to forty percent rather than sixty), a meaningful self-

custody Bitcoin allocation (perhaps ten to twenty-five percent depending on the investor's risk tolerance and horizon), a Bitcoin-collateralized credit allocation replacing most or all of the traditional sovereign bond allocation (perhaps twenty to thirty percent), and the remainder in some combination of cash, gold, real assets, and selective alternative exposures. The exact percentages will vary by investor. The structural shift — out of sovereign bonds, into Bitcoin and Bitcoin-collateralized credit — is what defines the model.

I am not the first to make this argument. Versions of it have been advanced by Lyn Alden, by Greg Foss, by Michael Saylor in his more measured moments, and by various others working in the institutional Bitcoin space. What I want to add is the connection between the portfolio-level argument and the deeper Hayekian argument about why the underlying assets in the 60/40 model are no longer reliable. The 60/40 model is dead because the assets it depends on are products of a monetary regime that is itself failing. The new model has to be built on assets whose performance does not depend on the failing regime.

The remaining chapters of Part III work through the practical questions of how to build that new model. Chapter 11 takes up the position-sizing question — how much Bitcoin, how much credit, how much else. Chapter 12 takes up the form-of-holding question, which I think is the most underappreciated risk in the entire space. Chapter 13 names the elephant in the room — the structural conflict in the financial planning profession that, if left unaddressed, would prevent the recommendations of this book from ever reaching most clients. Chapter 14 takes up the question of why sovereign bonds need to be largely or entirely replaced. And Chapter 15 takes up the question of what they should be replaced with.

CHAPTER ELEVEN

Position Sizing Under Uncertainty

The question every Bitcoiner has to answer, and answer continuously, is "How much should I own?" It is the question we ask ourselves and each other in any honest conversation about the asset, and it is the question conventional financial planning is the worst at answering. Nothing in this chapter is financial advice — I do not run an advisory practice — but the framework I find useful is the one I would want presented to me, and the one I think most readers will find more interesting than the conventional treatment of asset allocation typically allows.

The conventional approach is mean-variance optimization, as originally formulated by Markowitz. Mean-variance assumes that asset returns are normally distributed, that the expected return and variance of each asset can be reasonably estimated from historical data, and that correlations between assets are stable over time. None of these assumptions hold for Bitcoin. Bitcoin's return distribution is highly non-normal, with very fat tails on both sides. Bitcoin's expected return is essentially impossible to estimate from historical data, because the asset is sixteen years old and has been in a very different monetary environment for most of that history than it is in today. Bitcoin's correlations with other asset classes have varied substantially over time and across market regimes. Mean-variance is the wrong tool for the job.

The right tool, with appropriate humility about its inputs, is some version of the Kelly criterion. To take it seriously, you have to understand where it came from.

In 1956, John Kelly Jr., a researcher at Bell Labs, published a paper titled “A New Interpretation of Information Rate” in the *Bell System Technical Journal*. The paper began as an analysis of how much capital a gambler with private information about a rigged horse race should bet on each race to maximize the long-run growth rate of his wealth. Kelly derived a formula — now called the Kelly criterion — specifying the optimal bet size given the gambler’s edge and the odds offered. The formula was elegant, surprising, and instantly recognized as the correct answer to the position-sizing problem under uncertainty whenever the bettor expected to play the same game many times.

The Kelly criterion left the world of horse racing almost immediately and entered the world of card playing. Edward Thorp, a mathematician at MIT, used Kelly sizing in his now-famous 1962 work *Beat the Dealer*, which demonstrated that blackjack — properly played, with card counting, and with bets sized by Kelly — could produce a positive expected return for the player. Thorp went on to apply the same framework to financial markets, where it became one of the foundational ideas in quantitative trading. The genealogy of position sizing in modern finance runs through Bell Labs and Las Vegas before it ever reached Wall Street, and it is worth remembering that the idea is fundamentally a card player’s idea, dressed up in slightly more respectable mathematical language.

The formula itself, in its simplest form, is $f = (bp - q) / b$, where f is the fraction of bankroll to bet, b is the net odds received on the bet (payoff per unit risked), p is the probability of winning, and q is the probability of losing ($q = 1 - p$). The intuition is that the optimal bet sizes up the size of your edge (numerator) relative to the variance of the

bet (denominator). Apply it to Bitcoin and the inputs are subjective rather than known: assume, for the sake of a worked example, that you assign a 60 percent probability to Bitcoin substantially outperforming the alternative use of capital over your holding period, with the substantial outperformance defined as a 3-to-1 payoff on the position size relative to your alternative. Then $b = 3$, $p = 0.6$, $q = 0.4$, and $f = (3 \times 0.6 - 0.4) / 3 = 1.4/3 \approx 0.47$. Full Kelly is a 47 percent allocation. Half-Kelly, the conventional safety margin, is a 23 percent allocation. Quarter-Kelly is a 12 percent allocation. The range that emerges from any plausible set of inputs for a high-conviction Bitcoin investor is roughly 10 to 50 percent of net worth, with where you land in the range a function of how confident you are in your inputs and how much margin of safety you want against being wrong about them.

Card players think about risk differently than financial planners do. A financial planner thinks of risk as variance — the standard deviation of returns, the probability of a drawdown. A card player thinks of risk as ruin: betting too much and going broke before your edge has time to play out. Variance is uncomfortable; ruin is fatal. The Kelly criterion is built around the distinction. Bet too little and you under-deploy your capital, accepting a slower compound growth rate than your edge entitles you to. Bet too much and you risk ruin in the bad outcomes that, given enough trials, will eventually occur. The Kelly fraction is the position size that maximizes long-term geometric growth, balancing these two errors. Bet smaller than Kelly and you leave compound growth on the table. Bet larger than Kelly and you increase the probability of catastrophic drawdowns to a point where they become more than just uncomfortable.

Here is the part of Kelly’s framework that the conventional financial-planning treatment of asset allocation systematically ignores. The risk of betting too small is real. It is not just an opportunity cost in some neutral sense. It is a positive error, comparable in importance to the error of

betting too large. Warren Buffett has said versions of this throughout his career: “the most extraordinary thing in investing is not the things you’ve bought but the things you didn’t buy.” Charlie Munger went further, calling errors of omission — the great investments he passed on — by far the largest errors of his investing life, much larger than the errors of commission that get the headlines. Buffett’s pantheon of regrets is not the bad investments he made. It is the great ones he saw and did not act on. He has said this repeatedly, in print and in person, for fifty years. It has not stopped most professional investors from continuing to behave as though only commission errors counted.

For a high-conviction, asymmetric-payoff asset, the asymmetry of regret is itself asymmetric. If you are right that the asset is going to multiply many times over the long run and you held a 5 percent position when you could have held 25, the magnitude of your error in foregone upside is many times larger than the magnitude of the error you would have made by holding 25 percent in an asset that, in the worst case, declined by half. Ruin and missed upside are both errors, but on a high-conviction asymmetric asset, missed upside is the bigger one. The Bitcoin community has a meme for the cumulative recognition of this asymmetry — call it the maximalist insight — that conventional financial commentary tends to dismiss as zealotry. It is not zealotry. It is a serious recognition of the Kelly mathematics applied to asymmetric assets where the conventional safety-margin assumptions are themselves a source of error.

This connects to a second thread, which Annie Duke makes most clearly in *Thinking in Bets*. Duke, a professional poker player turned cognitive scientist, distinguishes between decision quality and outcome quality. A good decision can produce a bad outcome (you bet on the better hand and lost on a coin flip), and a bad decision can produce a good outcome (you bet on the worse hand and got lucky). The two have

to be evaluated separately. Most investors evaluate the quality of their portfolio decisions by the outcomes that resulted, which is exactly backwards: the only reliable way to evaluate a decision is by its quality at the time it was made, given the information then available. An investor who held 5 percent Bitcoin in 2015 and watched it become 50 percent of his net worth by 2024 made a bad decision that turned out well. An investor who held 25 percent Bitcoin in 2015 and watched it become 80 percent of his net worth made a good decision that also turned out well. Both look the same in outcome terms. Only the second was a good decision.

This is the part of the Bitcoin community's full-Kelly or super-Kelly tendency that I think the conventional commentary does not engage with seriously. There is a coherent argument, taking Kelly seriously, that for a sufficiently asymmetric asset with sufficiently high subjective probability of substantial outperformance, the conventional 5-to-15-percent allocation range is itself a Kelly underbet. The half-Kelly safety margin that most allocators apply for "conservatism" is, on a high-conviction asymmetric asset, simply a different kind of error: the error of capping the upside in a position where the upside is the entire reason for the position. I do not endorse 100-percent positions for most investors, and I will explain in a moment why. But I want to be honest that the case for them, when the inputs are right, is mathematically more defensible than the conventional treatment admits.

Why I do not endorse them more broadly comes down to the difficulty of knowing your own inputs. The Kelly criterion is unforgiving when you are wrong about the inputs — your expected return, your probability of various outcomes, your tail risk. Overestimate your edge by a factor of two and full Kelly becomes a path to ruin. The half-Kelly or quarter-Kelly margin of safety most professionals apply is not just timidity. It is recognition that the inputs are themselves uncertain, and

that the cost of being wrong about the inputs is asymmetric in the direction of ruin. For investors who have very high subjective probability of Bitcoin's continued substantial appreciation, who have long horizons, who have no near-term obligations that would force a sale into a drawdown, and who can psychologically tolerate 70-percent peak-to-trough declines without making errors, full or near-full Kelly may indeed be defensible. For most investors, including most Bitcoiners, those conditions do not all hold, and the conventional fractional-Kelly result of a 5-to-25-percent allocation, sometimes higher for specific circumstances, is a more honest answer.

There is one Kelly-input adjustment that the conventional treatment of Bitcoin sizing systematically ignores, and that I want to make explicit because it changes the math meaningfully. Bitcoin's historical compound annual growth rate is not a defensible estimate of forward edge. The historical CAGR was earned off a vastly smaller base than future returns will be earned from. A 100-percent CAGR off a $10 billion market capitalization is a different financial event than a 100-percent CAGR off a $2 trillion market capitalization — the first absorbs $10 billion of incremental capital, the second absorbs $2 trillion. The historical record is dominated by small-base mathematics that will not repeat at scale. Any honest Kelly application has to discount the implied edge to reflect this. The discount is not a haircut for skepticism; it is a haircut for arithmetic. A practitioner who plugs historical CAGR into the Kelly formula is solving a different problem than the one his future returns will actually be drawn from.

There is a serious caveat for younger investors, which the conventional life-cycle finance literature has long acknowledged: leverage early in the life cycle, when human capital is the dominant component of net worth and financial capital is small, can produce substantially higher long-run wealth than the unlevered alternative. Ian Ayres and Barry

Nalebuff made this argument compellingly in *Lifecycle Investing*. The argument applies in modified form to Bitcoin. A young investor with high human capital and modest financial assets has a structurally different position-sizing problem than an older investor whose financial capital is already large. The young investor can plausibly justify Bitcoin allocations well above the conventional 5-to-25-percent range, sometimes including modest leverage, on Kelly-with-human-capital grounds. The framework is the same; the inputs and constraints are different.

For investors with high subjective conviction, long horizons, and the right life-cycle position, considerably higher is defensible on Kelly mathematics, with the appropriate caveats about input uncertainty and psychological capacity for drawdowns.

I have deliberately avoided giving you specific allocation ranges in this chapter. This is not advice. What I will say is that the conventional 2-to-5-percent allocation that the most "responsible" Bitcoin commentary recommends is not a number that survives any honest application of the framework I have laid out. It is a number that emerges from professional discomfort with the asset, dressed up in the language of prudence. I would invite the reader, the next time he encounters such a recommendation, to ask the recommender to defend it in the language of portfolio theory. Specifically: what subjective probability is being assigned to substantial Bitcoin outperformance, and against what alternative use of capital, and at what assumed expected return for that alternative, does a 2-percent allocation actually maximize expected long-run wealth? The answer will, in almost every case, reveal that the 2-percent allocation is not the output of a position-sizing calculation at all. It is the output of a career-risk calculation by the recommender, in which the cost of being wrong about Bitcoin in either direction is asymmetric-

ally borne. That is a real consideration for the recommender. It is not a financial-planning consideration for the client.

Two clarifications. First, these percentages are after taking into account the asymmetric upside that makes Bitcoin a Kelly-relevant asset in the first place. The sizing math implicitly assumes that the asset has a meaningful probability of substantial outperformance over the long run; if you do not believe this, the Kelly fraction is much smaller, and you would not be reading this book. Second, these percentages refer to total net worth, not investable assets. An investor with a paid-off home worth $1 million and $500,000 in investable assets has $1.5 million in net worth, and the Bitcoin allocation should be calculated against the larger number, not the smaller. Treating the home as outside the portfolio understates the effective position size and overstates how much risk you have actually taken.

There is a complication that conventional position-sizing analysis does not handle well, which is the form-of-holding question. A position in Bitcoin held in self-custody and a position in Bitcoin held through an ETF are not the same instrument, even if they offer the same notional price exposure. They have different risk profiles, different tail behavior, different counterparty exposures, different jurisdictional exposures, and different optionality with respect to the failure modes the Bitcoin allocation is partially intended to hedge. I will address this in detail in the next chapter. For position-sizing purposes, the relevant point is that an ETF-held Bitcoin position is in important respects a fiat-denominated derivative on Bitcoin's price rather than a true Bitcoin holding, and should be sized differently in a portfolio than a self-custody position. A useful heuristic is to count ETF Bitcoin as approximately half of a self-custody Bitcoin allocation for purposes of the hedging function, while counting it the same for purposes of the directional-exposure function.

There is a second complication, which is the Bitcoin-collateralized yield position. Once an investor has a meaningful Bitcoin holding, the question arises of whether and how to use that holding to generate yield, either by lending it out or by borrowing against it. This is a separate position-sizing question that compounds with the underlying Bitcoin sizing question. It also forces the question that most position-sizing discussions in finance silently assume away: every investor has a cost of capital, whether or not he names it. The decision to hold an asset rather than lend it is implicitly a decision that the asset's expected return exceeds the rate at which it could otherwise be put to work. The decision to borrow against an asset is implicitly a decision that the cost of borrowing is below the marginal use of the borrowed funds. Both decisions require an explicit cost-of-capital framework. I will address this in Part IV, and I will return to the larger point — that everyone, including governments, has a cost of capital, and that a Bitcoin-anchored monetary system equalizes that cost across participants in a way the existing system does not — in Chapter 19.

For now, the framework I want you to take from this chapter is straightforward. Position sizing for Bitcoin is a Kelly problem, applied with humility about inputs, with the recognition that errors of omission are real and often larger than errors of commission. The answer is not zero, even for skeptics who think it is a scam but understand math. The answer is rarely 100 percent for most investors, but it is also rarely 5 percent for those who have done the work. A defensible range for most readers of this book is somewhere between 10 and 30 percent of total net worth, with specific circumstances pushing higher or lower. Sizing is a function of the investor's specific situation, his subjective probability assessments, and his life-cycle position, not of the asset's general merit. An asset can be excellent and still warrant a moderate position in any

given portfolio; it can also warrant a much larger position in another portfolio, with neither answer being wrong.

The next chapter takes up the form-of-holding question, which I think is the single most underappreciated risk in modern Bitcoin allocation.

CHAPTER TWELVE

Paper Bitcoin and the Custody Question

✦

The introduction of spot Bitcoin ETFs in the United States in early 2024, and the subsequent proliferation of similar products in other jurisdictions, has been celebrated by much of the Bitcoin investment community as a milestone in Bitcoin's institutional adoption. The argument is that ETFs allow institutional capital to allocate to Bitcoin in a regulated, familiar wrapper, that they bring price discovery into traditional financial markets, and that they validate Bitcoin as an investable asset class on terms that institutional allocators can engage with.

All of this is true, and the ETFs are an important development. They are also, when used as substitutes for self-custodial Bitcoin holdings, a potentially serious misunderstanding of what Bitcoin is and what an allocation to Bitcoin is meant to accomplish.

I want to develop this carefully, because the argument cuts against the grain of much current commentary and because the implications are uncomfortable for a great deal of the institutional Bitcoin product complex.

Begin with the question of what an allocation to Bitcoin is intended to accomplish in a portfolio. The conventional answer is some version of: directional exposure to Bitcoin's price, with the expectation that Bitcoin will appreciate over the holding period, and with diversification

benefits relative to the rest of the portfolio's holdings. This is the answer that frames Bitcoin as essentially analogous to other risk assets — equities, real estate, commodities — but with different return and correlation characteristics.

This answer is incomplete. The full answer, for any investor who has worked through the diagnosis in Part I, is that an allocation to Bitcoin is intended to accomplish two distinct things: directional price exposure, and a hedge against the specific failure modes of the existing financial and monetary system. The two functions are related but they are not identical, and they have different implications for how Bitcoin should be held.

Directional price exposure to Bitcoin can be achieved through any wrapper that tracks Bitcoin's price reasonably closely. ETFs do this. Futures-based products do this. Custodial accounts at exchanges do this. For purposes of capturing the directional exposure, the wrapper does not matter very much for the directional exposure itself, although counterparty risk in the wrapper is its own concern that I will address shortly. What matters for tracking the directional exposure is tracking error, fees, and liquidity.

Hedging against the failure modes of the existing system, however, requires that the Bitcoin position itself be outside the existing system. This is a categorical requirement, not a marginal preference. An ETF share is a security. A security is held within the brokerage system. The brokerage system is regulated by the same authorities that regulate the rest of the financial system. The brokerage system depends on the same custodial, settlement, and clearing infrastructure as the rest of the financial system. The brokerage system is exposed to the same counterparty, settlement, and operational risks as the rest of the financial system. In the failure modes that an investor is trying to hedge against — sovereign credit stress, financial system seizure or freeze, regulatory

action against specific asset classes, capital controls — the ETF position is exposed to the same risks as the rest of the portfolio it was supposed to be hedging.

Self-custodial Bitcoin is, by contrast, structurally outside the existing system. The keys controlling the position are held by the investor directly, not by any intermediary. The position is not held through a brokerage, is not subject to settlement risk, is not exposed to custodian failure, is not subject to redemption procedures or trading suspension or any of the other operational interventions that affect securities held through traditional channels. The position can be moved across borders without crossing any banking system, used as collateral in cryptographically enforced contracts without any trusted intermediary, and accessed at any time of day or night without depending on any market being open. These are not marginal advantages. These are categorical differences in the risk profile of the holding.

The case I want to make is that ETF Bitcoin and self-custody Bitcoin should be thought of as two different instruments that happen to share a price feed. They are both Bitcoin price exposures, but they are not both Bitcoin.

There is a useful historical analogy here, which I touched on briefly in Chapter 7. The bullion gold market has, for decades, operated with a paper-claim layer that vastly exceeds the underlying physical gold. The London bullion market operates on what amounts to a fractional reserve basis, with paper gold trading at fifty to one hundred times the volume of physical gold. (I believe something analogous will eventually develop for Bitcoin-collateralized credit, but with the structural difference that on-chain settlement and cryptographic enforcement allow for a much higher degree of transparency than the gold paper market has ever had.) As long as no one tries to take physical delivery in unusual quantities, the system functions. When the system is stressed, it has historically

been the case that paper gold and physical gold can trade at meaningfully different prices, and that paper gold holders may not be able to convert their claims into physical metal at the prices the paper market implies. The paper market and the physical market are connected but they are not the same market.

The Bitcoin market is at the beginning of building an analogous paper-claim layer. ETFs are paper Bitcoin in the technical sense: they are securities representing claims on Bitcoin held by a custodian, with the custodian responsible for honoring the claim. As long as the system functions normally, the claim is good. In the failure modes that matter — system-wide stress, regulatory intervention, custodian failure — the claim may not be honorable on the terms the holder expected. The history of similar paper-claim systems suggests that this risk is real and that it materializes precisely when the holder needs the underlying asset most.

There is a second, more structural point that follows from this. The 21 million cap on Bitcoin is meaningful only if a critical mass of holders insist on owning actual Bitcoin rather than claims on Bitcoin. If the institutional Bitcoin holding base shifts heavily into ETF and other paper claims, the effective supply of "Bitcoin exposure" can expand without limit through fractional-reserve dynamics in the paper layer, even while the underlying chain still has only 21 million coins. The monetary thesis of Bitcoin depends on the actual scarcity of the asset, which depends in turn on the actual scarcity being respected by the holding base. A holding base that accepts paper claims as substitutes for actual coins compromises the monetary thesis at a structural level.

For these reasons, I argue that any serious Bitcoin allocation in a portfolio should be predominantly in self-custody, with ETF and other paper exposure used only for specific purposes — typically tax-advantaged accounts where self-custody is not legally available, or

institutional mandates where regulatory or custody policies prevent self-custody. For the bulk of an investor's Bitcoin allocation, self-custody is not a preference. It is what makes the allocation actually accomplish what the allocation is supposed to accomplish.

This has implications for product design and for institutional practice that the industry has not yet fully absorbed. The financial planning profession is structured around custody-based products: assets held through brokerages, advisor platforms, trust companies. The infrastructure for self-custodied assets in a financial planning context is much less developed. This is, I think, the next major frontier in financial planning. The credit and yield products that I will discuss in Part IV are designed specifically to allow Bitcoin holders to use their holdings productively without surrendering self-custody. This is not an incidental feature. It is the entire point. A credit product that requires the holder to give up self-custody is structurally incompatible with the reasons the holder has for owning Bitcoin in the first place.

The next chapter takes up a related question that, if left unaddressed, would prevent the framework of this book from ever reaching most clients of the financial planning profession. It is the structural conflict in advisor compensation that systematically discourages the very recommendation Chapters 11 and 12 imply. After that, Chapters 14 and 15 take up the question of what should happen to the conventional fixed-income allocation in a portfolio designed around the considerations of these chapters. The short answer is that sovereign debt, which has been the default fixed-income holding for generations, has to be substantially or entirely replaced. The reasons follow.

CHAPTER THIRTEEN

The Elephant in the Room

✦

There is a question I have so far avoided in this book, and it is the question that any honest reader of Chapters 11 and 12 will already be asking. If a serious portfolio in the current monetary environment ought to include a meaningful self-custody Bitcoin allocation, replacing a substantial portion of the conventional fixed-income sleeve with Bitcoin-collateralized credit, then why is essentially nobody in the financial planning profession recommending this to their clients? The framework, if correct, has been available in some form for years. The conclusions follow from premises that are not particularly controversial once stated. And yet the average advisor recommendation for Bitcoin, where one is made at all, is a two-to-five percent ETF allocation that, on the analysis I have laid out, is not a position size at all but the smallest defensible gesture toward the asset.

The conventional explanation is that the profession does not yet understand Bitcoin, that it is appropriately conservative about new asset classes, and that recommendations will catch up as research and product infrastructure mature. There is something to this, and I do not want to dismiss it. But I think the real answer is structural, and I think it is worth naming plainly, because naming it is the first step toward solving it.

The financial planning profession in its dominant modern form is compensated through a percentage of assets under management. The standard fee for retail and mass-affluent practices is on the order of one percent annually. The fee is charged against the assets the advisor

custodies on the client's behalf, typically through a brokerage platform that holds the client's securities in accounts the advisor has read-and-trade access to. Assets that are not on the platform are not, generally, billed against. They are not part of the AUM number that determines the advisor's compensation, and they are not part of the AUM number that determines the practice's enterprise value at sale.

Self-custody Bitcoin sits entirely outside this structure. By design. A position whose private keys are held by the client, on a hardware device in the client's possession, is not on any platform, is not visible to any custodian, is not included in any quarterly statement the advisor's compliance department reviews. From the advisor's perspective, it is an asset that has left the room. An advisor who recommends that a client move twenty-five percent of net worth into self-custody Bitcoin is recommending a twenty-five percent reduction in his own AUM, and therefore in his own compensation, and therefore in the value of the practice he has spent his career building.

This is the elephant in the room. It is the largest unaddressed conflict of interest in the modern financial planning profession, and it dwarfs the conflicts that fee-only fiduciary advisors have correctly resolved relative to commission-based brokers. The fee-only model resolves the conflict between recommending products that pay the advisor a kickback and products that do not. It does not resolve the conflict between recommending assets that stay on the advisor's platform and assets that leave it. The first conflict is about which platform-held product to recommend. The second conflict is about whether to recommend that wealth leave the platform at all. They are different problems, and the second has not been seriously engaged with by the profession because, until very recently, there was no asset class that genuinely required leaving the platform to function as intended.

I want to be careful about the framing here, because I think the way the Bitcoin community has typically discussed this conflict has been counterproductive. The profession is not full of bad actors hiding the truth from clients to protect their fees. The profession is full of conscientious people operating inside a compensation structure that systematically penalizes a category of recommendation, in a regulatory environment that makes self-custody operationally awkward to advise on, with a peer group whose practice norms have not yet shifted, and with clients who in many cases would not adopt the recommendation even if it were made. Each advisor is making locally rational decisions inside this set of constraints. The aggregate result is a systematic underrepresentation of self-custody Bitcoin in client portfolios that, on the framework I have laid out, ought to include it. But the result is structural, not moral, and the path to changing it has to be structural too.

The closest analogy in recent professional history is the slow correction of the home-country bias in equity allocation. For decades, U.S. advisors recommended portfolios dramatically overweighted to U.S. equities relative to what global market capitalizations would have prescribed. Every individual recommendation was defensible. The aggregate pattern was a systematic distortion, driven in part by the structure of available products and the career risk of recommending allocations that diverged from peer practice. The bias was not corrected by client-by-client argument. It was corrected by the proliferation of low-cost global index products that made the conventional U.S.-overweight defensible only as an active bet, requiring the advisor to articulate why he was making it. The product infrastructure changed the conversation. The profession then followed.

I think the self-custody question is going to follow a similar arc, but the path will look different because self-custody is, by its nature, not a product that can sit on a platform. The resolution will not come from a

better Bitcoin ETF or a more sophisticated managed account. It will come from the development of an economic relationship between the advisor and the client around the self-custodied position itself — a relationship in which the advisor is genuinely valuable to the client whose Bitcoin is off-platform, and is compensated for that value, without the compensation depending on the asset returning to the platform. The productive layer that Part IV describes is, among other things, the infrastructure for that relationship.

Consider what an advisor can usefully do for a client with a meaningful self-custody Bitcoin position. He can advise on overall portfolio construction in a world where one of the largest asset positions is not on his platform. He can model the client's full balance sheet, including the off-platform Bitcoin, in financial planning software that treats self-custody positions as first-class citizens rather than as awkward write-ins. He can advise on the use of Bitcoin-collateralized credit to generate yield against the position without forcing a sale, or to access fiat liquidity without triggering a tax event, or to hedge the position with put structures of the kind described in Chapter 17. He can advise on the tax treatment of various productive uses of the holding, on the estate planning architecture that survives the client across generations without losing the keys, and on the operational discipline — multi-signature setups, seed phrase management, inheritance protocols — that distinguishes a serious self-custody practice from an amateur one. Each of these is real work, requiring real expertise, and each of them is genuinely valuable to the client. None of them require the Bitcoin to be on the advisor's platform.

The question is how the advisor is compensated for this work. The AUM model, applied to assets the advisor does not custody, does not work cleanly. Some practices have begun charging a fee against total net worth rather than against custodied AUM, which is a step in the right

direction but raises its own questions about how to verify net worth that is not on a platform and how to calibrate the fee level to the actual work being done. Others have moved to flat-fee, hourly, or retainer-based pricing for planning services, with a separate (typically smaller) AUM fee for the assets the advisor does still custody. Others are experimenting with project-based engagements for specific complex tasks — the initial setup of a self-custody architecture, the structuring of a Bitcoin-collateralized lending program, the integration of an estate plan that includes off-platform digital assets. The profession is, slowly, working out what the new fee architecture looks like.

I want to encourage this work, and I want to encourage advisors who are reading this chapter to engage with the framework as collaborators rather than as adversaries. The framework laid out in Parts I through III implies a substantial restructuring of client portfolios. That restructuring needs advisors. The clients who will follow this framework are clients who have a great deal of work to do, and they will do it better with competent professional help than without it. The profession that figures out how to advise on self-custody Bitcoin holdings, on Bitcoin-collateralized credit, on the integration of these positions into comprehensive financial plans, will be the profession that captures the next twenty years of growth in serious portfolios. The profession that holds onto the AUM-against-platform-assets model and quietly steers clients toward ETF allocations because that is what the platform supports will, I think, find that growth happening elsewhere.

The productive layer that Part IV describes is part of why this transition is more attractive than it might initially appear. A client with a self-custody Bitcoin position that is being used productively — generating yield through Bitcoin-collateralized lending, hedged with options, integrated into an income plan — is a client with an ongoing financial relationship that an advisor can participate in and add value to.

The advisor is not asked to refer the client to an off-platform asset and then have nothing more to say. The advisor is asked to help the client construct and manage a productive holding that has its own continuous economic activity, with its own decisions about tenor, loan-to-value, hedging, reinvestment, and integration with the rest of the portfolio. There is more advisory work in a productively deployed Bitcoin position than there is in a passive ETF holding. The compensation structure has not yet caught up with this fact, but it will, because the work is real and clients will pay for it.

What follows from this for the reader of this book?

If you are an investor, you should understand the conflict frankly. The advice you receive about Bitcoin from a conventional AUM-based advisor passes through a filter that systematically discounts the property of the asset that requires it to leave the platform. This does not mean the advice is dishonest or that the advisor is acting in bad faith. It means the advice is shaped by structural pressures that the advisor may not himself be fully aware of, and that you, as the client, should weight accordingly. The two-percent ETF allocation that your advisor is comfortable with has more to do with the compensation structure he operates inside than with the position-sizing analysis I laid out in Chapter 11. Adjust your interpretation of the recommendation accordingly, and have the conversation with your advisor explicitly. A good advisor will engage with the structural question honestly. An advisor who waves it away is telling you something useful about how to interpret the rest of his advice.

If you are an advisor, you should understand that the framework laid out in this book implies a restructuring of client portfolios that you are well-positioned to participate in if you are willing to engage with the structural question about your own compensation. The conversation with clients about self-custody is awkward in the current model. It does not have to be. There are practitioners working out new fee architectures,

new service offerings, new ways of being valuable to clients whose largest holdings are not on a platform. The infrastructure that makes those holdings productively useful — Bitcoin-collateralized credit chief among it — creates the basis for an ongoing economic relationship that justifies meaningful fees without requiring the underlying asset to move. The profession that gets there first will, I think, look back on the transition the way the global-equity-allocation profession looked back on the home-country-bias correction: as obvious in retrospect, valuable to the clients who experienced it, and a competitive advantage to the practitioners who led rather than followed.

The transition will not be easy and will not happen all at once. But it is starting. There are advisors today who have done the work, who advise on self-custody competently, who have moved their fee structures to accommodate off-platform holdings, and whose clients are better off as a result. They are still a small minority of the profession. They will not be in ten years. The AUM-against-platform-assets model, like the home-country-equity-overweight model before it, will erode under pressure from a better way of doing the same work, and the practitioners who lead the erosion will be better-positioned than the practitioners who resist it.

The next chapter takes up the related question that this one has set up: what should actually replace the conventional fixed-income allocation in a portfolio designed around the considerations of Parts I through III. The argument I make there about sovereign debt is, in its way, structurally parallel to the argument I have made here about the AUM model. Both are inherited conventions that worked reasonably well in a different monetary regime and that have to be reconsidered in the one we now inhabit. The difference is that sovereign debt has no obvious constituency to defend it beyond its own institutional inertia, while the AUM model has a profession built around it whose participa-

tion in the transition I am hoping to encourage rather than to circumvent. The case for the transition, in both, rests on the same underlying analysis. The path through it, in the case of the profession, depends on the willingness of practitioners to do the work. I hope this chapter has been a useful prompt for that work to begin, where it has not already.

CHAPTER FOURTEEN

Sovereign Debt Is Not Risk-Free

✦

In the conventional financial planning framework, sovereign debt — particularly U.S. Treasury debt — is treated as the risk-free asset against which all other returns are measured. The capital asset pricing model takes the risk-free rate as an input. The Sharpe ratio is calculated relative to the risk-free rate. Bond ladder strategies, retirement income planning, and most institutional liability matching all use sovereign debt as the foundation. The phrase "risk-free rate" is so ubiquitous in financial language that it is rarely questioned.

It needs to be questioned. The risk-free rate is not actually risk-free. It is risk-free in a specific and limited sense — namely, that the issuer can always print the unit in which the obligation is denominated, so that nominal default is unlikely — and it is risky in a far more important sense, namely, that the real value of the unit in which the obligation is denominated is being continuously eroded by the issuer's actions. To call this rate "risk-free" is to use the term in a way that misleads more than it clarifies.

Consider what risk-free actually has to mean for an investor whose objective is to preserve and grow purchasing power over time. It has to mean that the investor's purchasing power is preserved with high confidence over the holding period. A nominal-dollar obligation that is fully honored does not satisfy this requirement if the dollar's purchasing

power has substantially declined over the holding period. The investor receives back what was promised in nominal terms and is poorer in real terms. The obligation was nominally riskless and economically destructive. (There are inflation-protected sovereign instruments — TIPS in the United States, Real Return Bonds in Canada, similar issues elsewhere — that partially mitigate this exposure, but only to the extent that the holder's actual inflation rate matches the official CPI used for the indexation. For most allocators in most cycles, the actual inflation rate they experience differs substantially from the CPI used to adjust the bond, and the protection is correspondingly partial.)

This is not a hypothetical concern. Investors who held long-dated U.S. Treasury bonds from 2020 through 2022 experienced one of the worst real returns in the history of U.S. fixed income. Long-duration Treasuries lost approximately thirty percent of their nominal value in the 2022 rate-hiking cycle. Real returns over the same period, accounting for inflation, were substantially worse. The "risk-free" asset, held through what was supposed to be a routine period of monetary tightening, produced losses that any private credit investor would consider unacceptable from a single counterparty.

The 2022 episode is not anomalous. It is the predictable consequence of holding long-duration sovereign debt during a period when the sovereign issuer is forced to raise rates aggressively to contain an inflation problem the issuer itself produced. As long as fiscal positions remain unsustainable, as long as central banks remain the marginal buyer of new issuance, and as long as inflation episodes continue to require defensive tightening, sovereign debt will continue to produce these outcomes. The duration risk in long-dated sovereign debt is not a temporary feature of one policy cycle. It is the structural consequence of the issuer's position. There is, in fact, a gambler's-ruin element to the sovereign's position that the conventional treatment understates: a

sovereign that monetizes its own debt is, in effect, doubling its bets each round to stay in the game, with the implicit assumption that it can keep doubling indefinitely. That assumption holds until it doesn't. When it doesn't, the unwind is non-linear.

There is a second category of risk in sovereign debt that the conventional framework even more thoroughly mis-prices, which is what we might call sovereign credibility risk. The U.S. Treasury, like every major sovereign issuer, has obligations measured in trillions of dollars, denominated in a unit the issuer controls, against a tax base that is large but not infinite. The current trajectory of U.S. fiscal policy implies federal debt growing faster than GDP for the indefinite future, with debt-service costs already exceeding the defense budget and projected to exceed Social Security in coming years. There is no plausible path to fiscal sustainability that does not involve some combination of substantial spending cuts, substantial tax increases, substantial inflation, or substantial financial repression — any of which would impose significant real costs on holders of sovereign debt.

The investor holding a Treasury bond is, in effect, betting that the issuer will continue to honor the real value of the obligation in the face of fiscal pressures that make doing so increasingly difficult. This bet has worked, mostly, for the past forty years, with intermittent periods of inflation eroding the real return without producing a nominal default. There is no reason to believe it will continue to work as well over the next forty years. The fiscal trajectory is worse, the debt levels are higher, the tax base is more constrained, and the political will to address the imbalance is, by any reasonable measure, lower.

The honest reframe is that sovereign debt is a credit instrument, not a risk-free instrument. Like any credit instrument, it has counterparty risk. The counterparty in this case is the issuer's willingness and ability to honor the real value of the obligation. The willingness is a function of

political conditions; the ability is a function of fiscal conditions. Both are uncertain. Both are deteriorating. The conventional treatment of sovereign debt as risk-free obscures the fact that the holder is taking on a non-trivial credit exposure to a counterparty whose fundamentals are weakening.

If sovereign debt is a credit instrument with deteriorating fundamentals, then the conventional fixed-income allocation in a portfolio is not what it claims to be. It is not a low-risk anchor providing diversification and modest real returns. It is a substantial credit position, in a single counterparty, with poor real return prospects and high duration risk. Treated honestly, this position would not survive any normal portfolio review. It survives in conventional portfolios only because of inherited convention and because there has not, until recently, been a credible alternative.

The credible alternative now exists. Bitcoin-collateralized credit, structured properly, offers most of the functions a conventional fixed-income allocation is supposed to perform — current income, diversification from equity exposure, lower volatility than risk assets — without the structural deficiencies of sovereign debt. The yield is determined by market discovery rather than by central bank manipulation. The credit risk is collateralized by an asset that is structurally outside the failing fiat system. The duration risk can be managed by keeping loan tenors short. The counterparty risk can be eliminated through cryptographic enforcement of the collateral agreement. The result is a fixed-income exposure that is, in most respects, structurally superior to the sovereign debt it replaces.

I want to be careful here, because I am aware that I am proposing a substantial change to a model that has been the foundation of institutional and retail financial planning for generations. I am not proposing the change lightly, and I do not propose it as a marketing exercise. I propose

it because I am persuaded, on the underlying analysis, that the conventional model is structurally broken — that it is going to produce poor outcomes for investors over the next several decades, and that the alternative is now sufficiently developed to be a real option for serious portfolios.

The next chapter discusses what the new fixed-income allocation actually looks like in practice — what kinds of Bitcoin-collateralized credit products exist, how they are structured, what their risk profiles look like, and how they should be incorporated into a portfolio. This is the bridge from the diagnosis-and-allocation discussion of Parts I through III into the productive-layer discussion of Part IV.

CHAPTER FIFTEEN

The New Fixed Income

✦

In the previous chapter I argued that sovereign debt should be substantially or entirely replaced as the fixed-income anchor of serious portfolios. In this chapter I want to describe what should replace it.

The replacement is Bitcoin-collateralized credit, structured to preserve the properties that make Bitcoin valuable as a holding while generating market-discovered yield. This is a category of product that is still in its early development but is, I believe, the most important single innovation in fixed income in fifty years. I want to describe what the category looks like, what makes it structurally different from the failed crypto credit products of the previous cycle, and how it should be evaluated by allocators.

The basic structure of Bitcoin-collateralized credit is straightforward. A holder of Bitcoin posts the Bitcoin as collateral for a loan, typically denominated in fiat currency or a stablecoin. The lender provides the loan principal. The loan accrues interest at a market-determined rate. If the value of the Bitcoin collateral falls below a threshold relative to the loan principal, the loan can be liquidated to repay the lender. If the borrower repays the loan plus interest at maturity, the collateral is returned. This is a familiar credit structure, analogous to a margin loan or a securities-collateralized loan.

The structural innovation is in how the credit agreement is enforced. The previous generation of crypto credit products — BlockFi, Celsius, Genesis, and others — enforced credit agreements through trusted

custody. The borrower transferred Bitcoin to the platform's custody. The platform held the Bitcoin and made loans against it (and, problematically, often relent or rehypothecated the Bitcoin to other parties for additional yield). When borrowers wanted their Bitcoin back, they had to trust the platform to return it. When the platforms failed — and they all failed, in the 2022 contagion — the borrowers' Bitcoin was caught in bankruptcy proceedings, and many borrowers lost most or all of their collateral.

The new generation of Bitcoin-collateralized credit enforces credit agreements through cryptographic primitives rather than through trusted custody. The borrower's Bitcoin never leaves a wallet the borrower controls. The collateral arrangement is enforced by a discreet log contract — a cryptographic structure in which the conditions of the loan are encoded into a signed commitment that pays out automatically based on externally observed price feeds. If the Bitcoin price stays above a threshold and the borrower repays the loan, the contract returns the collateral to the borrower. If the price falls below the threshold, the contract pays out the collateral to the lender, with the surplus returned to the borrower. The contract is enforced by mathematics, not by trust in third parties.

This is the categorical difference between the new generation of products and the old. The old products required the borrower to trust the platform with custody. The new products require the borrower to trust only the underlying cryptographic primitives, which have been heavily studied and are widely understood. The platform, in this design, does not custody the borrower's Bitcoin and cannot abscond with it, rehypothecate it, or freeze it. The platform's role is to facilitate the matching of borrowers and lenders, to provide the price feeds that the contracts depend on, and to provide the user interface that makes the cryptographic primitives accessible to ordinary users. The platform is not a

custodian, and so the platform's failure cannot result in the loss of user collateral.

The reasons previous operators did not take this approach are partly that the cryptographic primitives required were not yet mature, partly that the custodial model was easier to build and operate, and partly that the custodial model offered larger profit margins to the operators because it allowed for rehypothecation and relending. The combination of immature cryptography, easier engineering, and higher margins produced a generation of products that were structurally vulnerable to the failure modes that materialized in 2022. The new generation of products eliminates the structural vulnerability by eliminating the custody model that made the vulnerability possible.

For an allocator considering Bitcoin-collateralized credit as the new fixed-income sleeve, the central question is: what is the actual risk profile of the position?

The answer has several components. The first is collateral risk: the risk that the Bitcoin collateral falls in value faster than the liquidation mechanism can capture, leaving the lender with insufficient collateral to be made whole. This is a real risk, but it is bounded by the loan-to-value ratio of the loan and by the speed and reliability of the liquidation mechanism. A loan made at 50 percent loan-to-value, is well-protected against all but the most catastrophic Bitcoin price declines over a one year time horizon. The protection can be extended further, as I will discuss in Part IV, by purchasing a put option on Bitcoin with a strike at or below the liquidation threshold, which converts the residual gap risk into a known cost.

The second is counterparty risk: the risk that the borrower defaults on the loan in some way the contract does not handle. In well-designed DLC-based products, this risk is essentially eliminated. The contract enforces the collateral agreement automatically; the borrower cannot

default in the conventional sense, because the borrower's only obligation is to repay the loan to recover the collateral, and failure to repay simply means the collateral is liquidated. There is no equivalent of a credit-card default, because there is nothing to default on beyond the collateral itself.

The third is operational risk: the risk that the cryptographic primitives, the price feeds, the user interfaces, or the protocols underlying the product fail in some way that affects the position. This risk is real but is shrinking over time as the primitives mature and as the operators gain experience. It is also bounded by the design of the system: the worst-case operational failure should leave both the borrower's collateral and the lender's principal in a recoverable state, even if the recovery requires manual intervention.

The fourth is regulatory risk: the risk that some jurisdiction takes action against the product or the operators in a way that affects the position. This is genuinely difficult to handicap, and it varies substantially by jurisdiction. The general direction in serious jurisdictions is toward regulatory frameworks that accommodate Bitcoin-native credit, but the path is uneven, and I do not pretend to be able to forecast it. This is genuinely not my lane. The risk that operators in this category have to manage most carefully, and the one that is most within their control, is operational risk.

Aggregating these risks, a well-structured Bitcoin-collateralized credit position is, in my view, structurally lower-risk than a comparable position in long-dated sovereign debt. The collateral risk is bounded and hedgeable. The counterparty risk is essentially eliminated. The operational risk is real but small and shrinking. The regulatory risk is uncertain but no worse than the regulatory risk attached to large positions in any cryptographic asset. By contrast, the sovereign debt position is exposed to substantial duration risk, substantial real-return

risk from monetary debasement, and a sovereign credibility risk that the conventional framework does not even recognize.

What does this look like in a portfolio? A reasonable starting point for a sophisticated allocator with a meaningful Bitcoin holding is to redirect a substantial portion of what would otherwise be the sovereign debt allocation into a Bitcoin-collateralized lending position. This generates current yield on Bitcoin collateral that the investor already holds, without requiring the investor to liquidate the Bitcoin position to access fiat liquidity. It replaces an asset class with deteriorating fundamentals (sovereign debt) with an asset class with structurally superior risk-return characteristics (collateralized lending). And it does all of this without compromising the self-custody of the underlying Bitcoin position, which, as I argued in Chapter 12, is essential to the function the Bitcoin allocation is meant to perform.

The mechanics of how this is structured, and the more subtle question of how the lending position's yield can be hedged to manufacture a near-risk-free rate, are the subject of Part IV. This chapter has aimed only to establish the strategic case: that Bitcoin-collateralized credit, properly structured, is the natural successor to sovereign debt as the fixed-income anchor of serious portfolios, and that the transition is not a marginal product substitution but a categorical upgrade in the structural soundness of the fixed-income sleeve.

The portfolio implications, taken together with the equity reduction implied by Chapter 10 and the self-custody Bitcoin allocation implied by Chapters 11 and 12, are substantial. The investor who follows the framework described in Part III emerges with a portfolio that looks very different from the conventional 60/40 model and that is, in my view, structurally much better suited to the monetary environment of the next several decades.

We turn now to the productive layer that makes this all possible.

PART IV

Completing the System

CHAPTER SIXTEEN

Time Value of Money on Bitcoin

✦

Sound money without credit is hoarding. This is the fact about Bitcoin that the maximalist version of the Bitcoin argument has, in my view, never fully grappled with. A money that cannot be productively lent and borrowed, that cannot earn interest in the unit it is denominated in, that cannot serve as the basis for time-structured contracts, is an incomplete monetary asset. It is a store of value, perhaps. It is not a complete monetary system. The same arguments can be made for payments — spend and replace.

Money has always been three things: a medium of exchange, a unit of account, and a store of value. Bitcoin's progress on each of these dimensions has been uneven. The store-of-value function was the first to develop, and it is the function that most early Bitcoin holders engaged with. The unit-of-account function is developing more slowly, as I described in Chapter 9, but is the function that ultimately matters most. The medium-of-exchange function is developing through Lightning and other settlement-layer innovations that are outside the scope of this book.

There is a fourth function, less often enumerated but no less important, which I want to call the intertemporal coordination function. This is the function by which money serves as the basis for credit — for the lending of present purchasing power against future repayment, with interest paid as compensation for the time preference of the lender and

the credit risk of the borrower. The intertemporal coordination function is the function that connects monetary assets to capital structure and to the productive economy. Without it, money cannot finance long-duration production, cannot fund investment, cannot serve as the basis for the time-structured contracts that constitute most of what we call finance.

For the entirety of human monetary history before Bitcoin, the intertemporal coordination function was performed through trusted intermediaries. Banks took deposits and made loans. The deposit-loan structure constituted credit creation. The bank, as intermediary, bore the credit risk of the borrowers, performed the maturity transformation between short-term deposits and longer-term loans, and stood between depositor and borrower as the trusted counterparty for both. The entire architecture of fractional reserve banking, central banking, deposit insurance, lender of last resort, and prudential regulation grew up around the management of the risks introduced by this intermediary structure.

Bitcoin presents a problem and an opportunity. The problem is that the trust-based credit architecture is incompatible with Bitcoin's foundational value proposition. Bitcoin is a money that does not require trust. Building a credit system on top of Bitcoin that requires trust — by handing custody of Bitcoin to an intermediary who then makes loans against it — reintroduces the trust dependency that Bitcoin was specifically designed to eliminate. The previous cycle of crypto credit products did exactly this, and exactly the predictable consequence followed: the intermediaries failed, and the depositors lost their Bitcoin.

The opportunity is that Bitcoin, as a programmable monetary asset, offers the possibility of credit structures that do not require trusted intermediaries at all. Credit can be enforced through cryptographic primitives rather than through trust. Loans can be collateralized by Bitcoin held in the borrower's own wallet, with the conditions of the loan encoded into a cryptographic contract that pays out automatically

based on observable conditions. The intermediary disappears as a credit-bearing entity and is replaced by mathematics.

This is what discreet log contracts, or DLCs, accomplish. A DLC is a cryptographic structure invented by Tadge Dryja in 2017 and developed substantially in the years since. The structure allows two parties to enter into a contract whose payout depends on the resolution of an externally observed event — for example, the price of an asset at a specified time — with the contract enforced through pre-signed Bitcoin transactions that are revealed and broadcast based on the event's outcome. Neither party needs to trust the other. Neither party needs to trust an intermediary. The contract is enforced by the same Bitcoin protocol that secures the underlying asset, with the same trust-minimization properties.

The relevance to credit is direct. A loan secured by Bitcoin collateral can be structured as a DLC in which the contract pays out the collateral to the lender if the borrower fails to repay, and pays out the collateral back to the borrower if the borrower repays. The conditions are encoded in the contract; the enforcement is automatic; no intermediary is required to custody the Bitcoin or to enforce the agreement. The borrower retains effective control of the Bitcoin throughout the life of the loan, in the specific sense that the Bitcoin is committed only to the agreed-upon contract is freed when the contract settles by repaying the loan. The lender has high-confidence collateral protection, in the specific sense that the contract pays out the collateral automatically in the event of default. Both parties have what they need; neither party needs to trust the other or any intermediary.

This is the structure that the next generation of Bitcoin-native credit operators are converging on. The structure is not, fundamentally, an innovation in credit; lending against collateral is one of the oldest financial structures in existence. The innovation is in the enforcement mechanism. We have moved the enforcement of the credit agreement

from trust-based custody to cryptographic primitives, and in doing so we have eliminated the failure mode that doomed the previous generation of crypto credit products.

What does this enable? It enables, for the first time, the development of a credit market on a sound monetary base. Bitcoin holders can put their holdings to work without surrendering self-custody. Lenders can deploy capital against high-quality collateral without depending on the operational integrity of any custodian. Borrowers can access fiat liquidity without selling their Bitcoin and without trusting an intermediary not to lose, hypothecate, or freeze their position. The basic operations of credit — the lending of present purchasing power against future repayment, the matching of savers and borrowers, the discovery of market interest rates — can all happen on a Bitcoin base, with the cryptographic enforcement that the Bitcoin base makes possible.

I want to emphasize how new this is. The infrastructure to do this at scale did not exist three years ago. It exists today in early but real form, and it will become substantially more developed over the next several years. The implication for Bitcoin as a monetary system is that Bitcoin is in the process of acquiring the fourth function — intertemporal coordination — that has historically been provided only through trust-based intermediation. When this function is fully developed, Bitcoin will be a complete monetary system: medium of exchange (via Lightning and other settlement layers), unit of account (via the slow process of denomination shift), store of value (already established), and intertemporal coordination (via DLC-based credit and its successors).

The completeness matters. A complete monetary system can support the entire architecture of modern finance — capital markets, credit markets, derivatives, structured products, time-structured contracts of every kind — on a sound-money base. An incomplete monetary system supports only a subset, and forces participants who need the missing

functions to use parallel systems built on unsound money for those functions. Most Bitcoin holders today still rely on fiat-denominated credit for any of their borrowing needs, because Bitcoin-denominated credit at scale does not yet exist. This forces the holder to operate in two monetary systems simultaneously, with the corruption of the fiat system continuing to affect the holder's affairs even on the margins where his Bitcoin allocation is supposed to insulate him.

The full vision of Bitcoin as a complete monetary system is one in which a holder can conduct essentially all his financial affairs — saving, borrowing, lending, contracting, transacting — within the Bitcoin system, without ever needing to trust a fiat-denominated counterparty for any of his economic activity. We are not at this point yet, and certainly most people will not be at this point for some years. But the trajectory is clear, and the productive-layer infrastructure that will get us there is being built now.

The next chapter takes up a specific and important sub-question of the productive layer: how can lenders in this new credit system manufacture a near-risk-free yield, comparable to the function sovereign debt has historically performed in portfolios? The construction is elegant and important.

CHAPTER SEVENTEEN

The Synthetic Risk-Free Rate

✦

In Chapter 14, I argued that sovereign debt is not actually risk-free, and that the conventional fixed-income allocation that depends on it should be substantially replaced. In this chapter I want to describe how the replacement can be structured to actually deliver the function sovereign debt was supposed to perform — namely, a low-risk, predictable yield that anchors the conservative end of a portfolio.

The construction I am going to describe is straightforward but not, to my knowledge, widely understood outside the small community of practitioners building Bitcoin-native credit products. It is the construction by which a Bitcoin-collateralized loan, combined with an appropriate hedging position, produces a synthetic near-risk-free yield in fiat terms. The construction has direct analogues in the prime brokerage and securities lending world, where similar structures have been used for decades to manufacture synthetic risk-free rates from securities collateral. The novelty is that the construction can now be performed on Bitcoin collateral, without the prime broker, without rehypothecation, and without sovereign credit underneath.

Here is how it works.

To make this concrete using a representative 50 percent LTV product as the example, the typical structure works as follows. The borrower contributes $120,000 worth of Bitcoin as collateral and the lender contributes $100,000 in Bitcoin collateral, with the contract enforcing the agreed-upon rate and term. The asymmetry in collateral contribution

reflects the over-collateralization required to give the lender appropriate protection against price declines, in addition to the interest. In this example, the lender is expecting to receive $110,000 at maturity from a collateral pool that was worth $220,000 at inception.

From the lender's perspective, the position has the following risk profile. The lender receives 10 percent annualized yield on the $100,000 deployed. The lender's principal is protected as long as the collateral remains worth more than $110,000. The lender's main residual risk is gap risk: the risk that Bitcoin's price falls and the collateral at a level sufficient to fully repay the principal at maturity. This is a real risk for very large or very fast price moves, and it is the only material risk in the structure once collateral risk and counterparty risk are properly handled.

Now consider what happens when the lender purchases a put option on Bitcoin with a strike price at or near the liquidation threshold. The put gives the lender the right to sell Bitcoin at the strike price, regardless of the actual market price. If Bitcoin's price collapses below the liquidation threshold faster than the contract can liquidate, the put pays out the difference between the strike and the actual market price, recovering the lender's gap exposure. The put converts the residual risk into a known cost — the premium paid for the put — which can be deducted from the loan yield to produce the lender's net return.

For Bitcoin at typical long-tenor implied volatilities, a one-year put with a strike at 50 percent of spot is meaningfully out of the money but not absurdly so, and its cost is highly sensitive to the assumed volatility. Running Black-Scholes at a 5 percent risk-free rate, with spot of 100 and strike of 50, the one-year put is worth roughly 0.3 percent of notional at 40 percent implied volatility, roughly 1.0 percent at 50 percent vol, roughly 2.0 percent at 60 percent vol, and roughly 3.4 percent at 70 percent vol. Bitcoin's listed long-tenor implied vols have generally sat in the 50-to-70 percent range, putting the realistic cost of the hedge in the

1-to-3 percent range, with 2 percent a defensible mid-point estimate at current market conditions. The exact cost varies with prevailing volatility and with the specific strike chosen.

The lender's net return is then 10 percent (loan yield) minus 2 percent (put cost), or 8 percent annualized, with the residual risk substantially eliminated. The lender has manufactured a near-risk-free 8 percent yield from Bitcoin collateral, with no exposure to sovereign credit, no duration risk beyond the one-year loan term, and no dependence on any counterparty's continued operational solvency. Varying the loan-to-value ratio across loans of the same term effectively triangulates the implied risk-free rate at that term: a portfolio of loans at different LTVs, with the corresponding put hedges, produces a converging estimate of the rate at which the market is willing to lend against Bitcoin collateral, free of gap risk.

I want to emphasize what this produces, because it is structurally categorical. The lender now has a fixed-income position that is, in nearly every dimension, structurally superior to a comparable position in U.S. Treasury debt. The yield is comparable to or better than current Treasury yields. The credit risk is collateralized by an asset structurally outside the failing fiat system, hedged for residual gap exposure, and enforced cryptographically rather than through trust. The duration risk is bounded by the loan term, which can be chosen to match the lender's preferred horizon. The counterparty risk is essentially eliminated by the cryptographic enforcement structure. The position is, in short, what sovereign debt has been claimed to be but is not: a low-risk, predictable yield, anchoring the conservative end of a portfolio, without the structural deficiencies that make conventional fixed income increasingly inadequate.

There are several refinements to the basic construction that are worth noting briefly.

First, the put hedge can be partial rather than complete. A lender willing to bear some residual gap risk can choose a put strike further out of the money, paying a lower premium and accepting a higher residual risk. This produces a higher net yield with a slightly worse risk profile, and the choice of strike becomes a portfolio decision based on the lender's overall risk tolerance.

Second, the loan-to-value ratio can be varied. A lower LTV (say, 30 percent rather than 50 percent) provides more collateral cushion and reduces the put cost required to fully hedge gap risk, but typically also produces a lower loan yield because borrowers will pay less interest for less leverage. A higher LTV (say, 70 percent) produces higher loan yield but requires a more aggressive hedge to maintain comparable safety. The optimal choice depends on the prevailing relationship between loan yields and option premiums at any given time, and this relationship will evolve as the market matures. (This is the triangulation point on the synthetic risk-free rate that I referred to earlier: by varying LTV and observing the corresponding loan rates, the market collectively discovers the term-structure of risk-free yields against Bitcoin collateral.)

Third, the position can be laddered. Rather than holding a single loan, the lender can structure a portfolio of loans of staggered tenors — one-month, three-month, six-month, twelve-month, twenty-four-month — providing continuous reinvestment opportunities and exposure to multiple points on the emerging Bitcoin yield curve. This is the structure that traditional Treasury portfolio managers use, and it translates naturally to the Bitcoin-collateralized credit context.

Fourth, the position can be denominated in Bitcoin rather than in fiat. A lender willing to lend Bitcoin against Bitcoin collateral can earn a yield in Bitcoin terms, which is a fundamentally different proposition from earning a fiat yield. A Bitcoin-denominated yield compounds the holder's Bitcoin position; a fiat-denominated yield converts to additional

fiat purchasing power. Both have their uses, and a sophisticated lender will likely want exposure to both.

The point of this chapter is not to provide a complete operating manual for Bitcoin-collateralized credit construction. It is to demonstrate that the basic construction exists, that it produces a structurally sound near-risk-free yield, and that it is the natural replacement for the sovereign debt allocation in a serious portfolio. The construction is not exotic; it is a straightforward application of well-understood financial primitives — collateralized lending, options hedging, laddered tenors — to a new collateral asset whose properties make the application particularly attractive.

What this chapter has described is a single point on a yield curve. The next chapter takes up the larger structure — the emergence of an entire market-discovered yield curve in Bitcoin-collateralized credit, and what this means for capital allocation in the broader economy.

CHAPTER EIGHTEEN

The Bitcoin Yield Curve

✦

The yield curve is not a financial product. The yield curve is the most important price signal in the entire economy. It is the price signal that coordinates capital across time. When the yield curve is corrupt, the entire intertemporal structure of capital is corrupt. When the yield curve is honest, the intertemporal structure can be coordinated honestly. Restoring an honest yield curve is the most important thing that needs to happen for sound capital allocation to be possible at all.

I want to make this case carefully, because it is the chapter on which the entire book turns.

What is a yield curve? A yield curve is the schedule of interest rates at which borrowers can borrow money for different lengths of time. A typical yield curve plots the rate as a function of maturity — overnight, one month, three months, one year, five years, ten years, thirty years. The shape of the curve contains an enormous amount of information about the economy. The level tells us about the general cost of capital. The slope tells us about expectations of future rates and about the liquidity premium between short and long maturities. The shape — normal, flat, inverted, humped — has historically been one of the most reliable indicators of economic conditions and of investors' collective view of the future.

What does a yield curve coordinate? It coordinates the time structure of capital. A project that requires capital deployed for ten years, at returns that justify a ten-year cost of capital, can be financed at the ten-

year rate. A project that requires only short-term capital can be financed at the short-term rate. The relative pricing of capital across different maturities determines which projects get funded and which do not. Long-duration capital-intensive projects — infrastructure, factories, research and development, real estate development — are particularly sensitive to the long end of the curve, because their viability depends on the cost of long-duration capital. Short-duration projects — inventory financing, working capital, short-term trading — are sensitive to the short end.

When the yield curve is determined by market discovery — by the actual willingness of savers to defer consumption against the actual demand of borrowers for capital across time — the curve provides accurate price signals about the economy's true intertemporal preferences. Capital flows to the projects that the curve indicates are economic. Capital is withheld from projects that the curve indicates are not. The economy's capital structure self-organizes around the curve, in a process that is the time-extended analog of what Hayek described for spot markets in *The Use of Knowledge in Society*.

When the yield curve is determined by central bank policy — by deliberate setting of the short rate, by quantitative easing or tightening of the long rate, by forward guidance about future policy — the curve no longer carries the same information. The curve carries information about the central bank's policy choices rather than about the economy's true intertemporal preferences. (One could argue that central banks generally just reflect and react to the expectations of the market they operate within, rather than truly setting rates against the market's wishes; even if so, the rate they set is still an output of deliberation rather than discovery, and the distortions that follow are the same.) Capital flows to projects that the policy-determined curve indicates are economic, but those projects may not be economic at the underlying market-determined

rate. The result is intertemporal misallocation: long-duration projects funded at policy-suppressed rates that would not have been funded at true rates; short-duration speculation incentivized by policy-suppressed short rates; capital structures that look profitable under current policy but that will be revealed as malinvestment when the policy changes.

This is the Austrian critique of central banking, in its purest form, and it is the critique that I think has been most undervalued in mainstream macroeconomic discussion. The 2008 financial crisis was, on this view, the consequence of fifteen years of policy-suppressed rates after the dot-com collapse. The 2020s inflation episode is the consequence of fifteen years of policy-suppressed rates after the 2008 collapse. Each cycle of policy-suppressed rates produces a generation of capital structure that depends on the rates remaining suppressed; when conditions force rates higher, the capital structure has to be liquidated, often catastrophically. This is not bad luck. This is the predictable mechanism by which a corrupted yield curve produces intertemporal misallocation that has to be eventually unwound.

Since 1971, no major economy in the world has had a market-discovered yield curve. Every major sovereign debt curve traces back to a policy rate that is deliberately set, with central bank purchases or guidance shaping the longer maturities. The most important price signal in the global economy has been, for half a century, an administrative output rather than a market discovery. The cumulative cost of this, in terms of intertemporal misallocation, is not knowable in any precise sense, but it is, on any reasonable accounting, enormous.

What I am going to argue is that the next several years will witness the emergence of the first market-discovered yield curve in any major economic system since the demonetization of gold. The curve will not emerge in fiat-denominated markets, because fiat-denominated markets remain downstream of central bank policy choices. The curve will

emerge in Bitcoin-collateralized credit markets, where neither the borrower nor the lender depends on any central bank for their participation, and where the rates are determined by the actual willingness of capital to be deployed against actual demand for capital across time.

The mechanism is straightforward. As Bitcoin-collateralized credit products mature and scale, loans of varying tenors will be transacted continuously between many borrowers and many lenders. Each loan represents a market-clearing rate at the specific tenor for the specific loan-to-value and specific risk profile. Aggregating across many loans, at many tenors, with many participants, produces an empirical yield curve in Bitcoin-collateralized credit. The curve will have all the features of a normal sovereign yield curve — a level, a slope, a shape — but it will be derived from market discovery rather than from policy choice.

This curve will be, when it emerges, the first true yield curve in any major economy in fifty years. I do not say this with any rhetorical exaggeration. I say it because it is, on the analysis I have laid out, literally the case. The fiat yield curves that exist today are policy outputs. The gold-era yield curves of the past were market outputs but in a much narrower, less liquid, less globally integrated market than the Bitcoin-collateralized market will be at maturity. A globally accessible, deeply liquid, market-discovered yield curve in a sound-money base does not exist anywhere in the world today. It will exist, in functional form, within the next several years.

What does the existence of such a curve enable? Several categorical things.

First, it enables honest pricing of long-duration capital, for the first time in fifty years. Projects that depend on the cost of long-duration capital — infrastructure, capital-intensive industry, multi-decade investment — can be evaluated against a curve whose long end is not being suppressed by central bank action. The viability of these projects

can be assessed against true market rates rather than against subsidized rates. The Austrian intertemporal coordination problem is solved, in the specific sense that capital can flow to the projects that are actually economic at the actually market-determined cost of capital across time.

Second, it enables the entire derivative apparatus of modern finance to be reconstructed on a sound base. Forward rates, swap curves, options on rates, interest rate derivatives of every kind, depend on the existence of an underlying yield curve to be priced and hedged. As the Bitcoin yield curve emerges, all of these instruments become possible in Bitcoin-denominated form. The financial superstructure that has been built on policy-distorted fiat rates can be replicated on market-discovered Bitcoin rates, with vastly better information content and vastly better risk properties.

Third, it enables monetary economics, as a discipline, to make empirical claims that have not been possible to make in fifty years. The natural rate of interest — the rate at which the supply of savings clears the demand for borrowable funds, in the absence of central bank intervention — is currently a theoretical construct, estimable only through indirect inference. Once a Bitcoin-collateralized credit market exists at scale, the natural rate becomes empirically observable, in real time, in continuously updated form. Hayek's argument that the price of money is a market discovery ceases to be a normative claim and becomes an operational fact.

Fourth, and most importantly for the purposes of this book, it completes the Bitcoin monetary system. Bitcoin, with a yield curve, is a complete monetary asset. It supports the full range of intertemporal coordination functions that modern finance requires, on a sound-money base, with cryptographic enforcement, without trusted intermediaries. There is no remaining function of a complete monetary system that Bitcoin cannot perform. The construction of the system, which began in

2009 with the publication of the Bitcoin protocol, is finished — not in the sense that there is no further work to do, but in the sense that all the necessary primitives are in place and all the necessary functions are operational.

This is, I think, the most important thing happening in monetary economics today, and almost nobody is talking about it in these terms. The press coverage of Bitcoin focuses on price action, ETF flows, regulatory developments, and the personal dramas of various major holders. The press coverage misses the structural fact that, beneath the surface, the conditions for the first market-discovered yield curve in fifty years are being assembled. When the curve emerges as an observable, named, traded market structure — it will be the most consequential single development in financial markets since the closure of the gold window in 1971.

That is a strong claim, and I want to defend it directly, because the obvious candidates for the most consequential post-1971 development do not, on inspection, match what a market-discovered yield curve would represent. The rise of derivatives over the past several decades has been transformative in scale and complexity, but every derivative trades on top of an administered yield curve; the underlying base rate is a policy output, and the derivative inherits that policy output as an input. The development of ETFs and the broader securitization of risk have been important for distribution and access, but ETFs are wrappers around assets whose underlying prices are themselves denominated in policy-managed money. The 2008 financial crisis was a stress event of historic magnitude that reshaped institutional structure and regulatory practice, but it did not restore market discovery to any administered price; it confirmed and entrenched the central bank's role as the marginal pricer of risk in the system. Each of these is consequential in its own way. None of them is the structural restoration of price discovery to the

most important price in the economy. A market-discovered yield curve in a sound-money base is categorically different. It is not an innovation built on top of the administered system. It is the reconstruction, outside the administered system, of the price-discovery mechanism for capital itself. That is what makes the comparison to the 1971 break apt rather than rhetorical. 1971 was the moment a market-discovered monetary anchor was abandoned. The emergence of a Bitcoin-collateralized yield curve is the moment that abandonment is structurally reversed.

The closing chapter of this book draws together the threads of the entire argument. We have moved from the diagnosis of the failed monetary regime, through the identification of Bitcoin as the response, through the portfolio implications for serious allocators, to the productive layer that completes the system. The final chapter places this entire argument in the longer context of the Hayekian project and explains what I mean by the title of the book.

CHAPTER NINETEEN

Everyone Has a Cost of Capital

✦

There is an implicit assumption embedded in almost every conversation about money, credit, and finance, and it is so universal that most participants do not realize they are making it. The assumption is that the cost of capital is something that applies to corporations, to investment projects, to leveraged buyouts and infrastructure financings, but not to the people doing the talking. Households do not think they have a cost of capital. Sovereigns do not think they have a cost of capital — they think they have a "policy rate" or a "borrowing cost" or an "interest expense line." The notion that every economic agent on Earth has, at every moment, a cost of capital that determines what he can and cannot afford to do is somehow not visible in most ordinary economic discussion.

It should be. Everyone has a cost of capital. The only question is whether it is honest.

A cost of capital is the rate of return at which an actor is indifferent between deploying capital today and waiting. For a household, it is the rate at which paying down a mortgage becomes a worse use of marginal cash than investing in a higher-returning asset. For a small business, it is the rate at which expanding inventory or hiring an additional employee becomes uneconomic. For a public corporation, it is the weighted average cost of capital — the blended return required by debt and equity

holders — that justifies any given investment project. For a sovereign, it is the rate at which issuing more debt to fund another spending program becomes prohibitively expensive given the existing debt-service burden. In every case, the cost of capital is the price signal that determines what gets built, what gets funded, what gets foregone.

The current monetary regime distorts this price signal across the entire structure of capital users, in a hierarchy that has been so naturalized in modern finance that almost nobody questions it. At the top of the hierarchy sits the sovereign, who borrows at the policy rate set by his own central bank, with the central bank as the marginal buyer of his issuance during periods of fiscal stress. The sovereign's cost of capital is, in effect, whatever the central bank decides it should be. If the central bank wants the sovereign's cost of capital to be 4 percent, it is 4 percent; if the central bank wants it to be 0 percent, it is 0 percent; if the political pressure to keep it low becomes overwhelming, it can be made arbitrarily low through balance sheet operations the central bank is empowered to conduct.

Below the sovereign sit the largest financial institutions, which borrow at the policy rate plus a small spread, and which serve as the primary distribution mechanism for the central bank's monetary policy throughout the rest of the economy. Below them sit investment-grade corporations, who borrow at the policy rate plus a more substantial spread, with the spread reflecting the conventional credit risk premium of the issuer. Below them sit smaller corporations and high-yield issuers, with larger spreads still. Below them sit ordinary households, who borrow at policy plus a substantial spread for a mortgage, plus a much larger spread for an auto loan, plus a still larger spread for credit card debt. Below them sit small businesses, paying SBA-loan rates or worse. At the very bottom sit borrowers without access to formal credit at all,

paying payday-loan rates that would be illegal under usury laws of any prior century.

This hierarchy is not a function of true credit risk. It is a function of distance from the central bank's spigot. The sovereign's claim to a low cost of capital does not rest on his superior creditworthiness, which by any honest accounting is deteriorating as fiscal positions worsen. It rests on his ability to issue debt that the central bank stands ready to monetize. Every other participant in the economy is paying a higher cost of capital not because their credit is worse but because they are further from the source. The hierarchy of cost-of-capital is, in significant part, a hierarchy of access to monetary subsidy.

This has consequences that ripple through every layer of the economy. The sovereign, paying a politically subsidized cost of capital, can fund things that would not be funded if he had to pay a market rate. He can fund wars whose returns no honest investor would underwrite at the true cost of capital required by the risk. He can fund welfare programs whose actuarial liabilities would not be sustainable at any market rate. He can fund the maintenance of an empire whose ROI is, on any honest calculation, deeply negative. He can do all of this because his cost of capital is not market-discovered; it is administratively set by an institution whose policy choices are influenced by his own fiscal needs.

The corporate sector, paying a cost of capital several hundred basis points above the policy rate but still substantially subsidized by the central bank's overall stance, can fund stock buybacks at scale, can engage in leveraged dividend recapitalizations, can pursue M&A whose value-creation rationale would not survive a true cost of capital. The proliferation of zombie firms in the post-2008 era — companies whose interest expense exceeds their operating earnings, sustained only by the willingness of credit markets to refinance them at policy-suppressed rates — is the most visible evidence of corporate cost-of-capital

distortion. These firms would not exist at a true cost of capital. They exist because the cost of capital is administratively suppressed, and they consume capital that would otherwise flow to firms whose returns justify the deployment.

The household sector, paying a still higher cost of capital, makes worse decisions because of it. A household that faces a 6 percent mortgage rate and a 25 percent credit card rate is making borrowing decisions on a different basis than a household that faces those rates plus an inflation rate that erodes the real value of fixed nominal obligations. The misalignment between cost of capital and inflation expectation produces systematic over- or under-borrowing, depending on which way the misalignment runs in any given decade. Over a long enough time horizon, almost no household correctly handicaps its own true cost of capital, because the inflation-adjusted real cost varies more than ordinary financial planning can accommodate.

A Bitcoin-anchored monetary system, with a market-discovered yield curve in Bitcoin-collateralized credit, does something to all of this that is difficult to overstate. It equalizes the cost of capital across participants, in the specific sense that the same market rate is available to anyone with the same collateral, regardless of his political proximity to any monetary authority. The sovereign, in such a system, is not a privileged borrower with access to subsidized rates. He is just another borrower competing for capital with every other participant in the market. If his credit is good, his cost of capital is competitive with the market rate. If his credit is poor, his cost of capital reflects that, and the projects he can fund are correspondingly constrained.

This is the structural change that I think the conventional Bitcoin commentary has not fully grappled with. Bitcoin is sometimes discussed as if its consequence for the state would be to make state borrowing impossible. That is not quite right. The consequence is more subtle and, I

think, more important. The state would still be able to borrow. It would just have to borrow at a market rate, against a market-determined assessment of its creditworthiness, in a market it does not control. The cost of capital it faces would no longer be a function of its own central bank's willingness to monetize its issuance. It would be a function of what savers and capital allocators, given their own opportunity sets, are willing to accept in exchange for lending to a sovereign issuer with the underlying credit characteristics of the modern fiscal state.

For most modern sovereigns, with debt-to-GDP ratios well above 100 percent and unfunded entitlement obligations many times larger than reported debt, that market rate would be substantially higher than the policy rates they currently pay. Many of the spending programs they currently fund would become uneconomic. Many of the wars they currently finance would become unaffordable. The fiscal capacity of the modern state would be substantially constrained — not by any moral or political judgment, but by the simple fact that the state would, like every other capital user, face an honest cost of capital.

There is a deeper consequence of this dynamic that I want to make explicit, because it is the structural pattern compounded over decades that the surface argument about specific programs does not capture. A sovereign that faces no real budget constraint — because his cost of capital is administratively set rather than market-discovered — will predictably expand into every economic niche where a political constituency exists for him to do so. This is not a moral failing of any particular sovereign. It is the natural behavior of any actor with access to an apparently free input. He will keep deploying the input until something forces him to stop. Without market discipline on his cost of capital, nothing forces him to stop except the eventual loss of currency credibility, which arrives slowly and is denied loudly until it is undeniable.

The result, observed over decades in every developed economy, is the steady displacement of private economic activity by sovereign-funded activity. Government programs grow into spaces where private enterprise once operated. Subsidized credit flows to politically favored projects, leaving market-tested projects starved of capital. Regulatory bandwidth expands to cover behaviors that had previously been left to private judgment. The cumulative effect is not just that the state grows, but that the state takes the oxygen out of the room — every dollar of subsidized sovereign borrowing is a dollar that could have funded private investment, every regulatory expansion is a constraint on private experimentation, every program that crowds in becomes a program that crowds out. Over a long enough horizon, the economy fills with sovereign-funded and sovereign-mediated activity, and private innovation, risk-taking, and productive use of capital atrophy. The visible consequence is that nothing meaningful happens at scale outside the sovereign's perimeter, because nothing meaningful at scale can happen outside the sovereign's perimeter when the sovereign's perimeter is funded at a subsidized rate that no private actor can match.

A market-discovered cost of capital reverses this dynamic mechanically. The sovereign that has to borrow at a market rate has to choose what to fund, the way every other borrower has to choose. Programs that would have been funded at the subsidized rate become uneconomic at the true rate. The crowding-out reverses. Private capital flows back into the spaces it had been displaced from. The oxygen returns to the room. This is, in the end, what the cost-of-capital reframing is actually about. Not the abstract correctness of price signals, but the concrete consequence for whether anything productive can happen outside the sovereign's perimeter.

This is the principal mechanism by which Bitcoin matters at the macro-political level, and it is the mechanism most often misdescribed.

The argument is not that Bitcoin overthrows the state. The argument is that Bitcoin restores the cost-of-capital discipline that the state has, since 1971, been systematically exempt from. The state remains. The state's privileged access to subsidized capital does not.

The corollary, which matters as much as the direct effect on sovereigns, is that every other participant in the economy gains access to the same market rate that everyone else has. The household with good credit does not pay the policy-rate-plus-banking-spread that he currently pays. He pays a rate that reflects his actual creditworthiness against a market-discovered base. The small business does not pay an SBA rate that reflects political programs. It pays a rate that reflects its own credit profile against the same base. The marginal effect for households and small businesses is, in many cases, a substantial reduction in cost of capital, because the policy-rate-plus-spread structure the existing system imposes on them is itself a distortion driven by the banking system's privileged distribution role.

The same logic — that everyone has a cost of capital and that Bitcoin disciplines the price of capital across participants — has a second-order consequence that I think is the most underappreciated game-theoretic feature of the entire transition. Sovereigns are not merely passive receivers of the cost-of-capital discipline that a Bitcoin-anchored system would impose. Sovereigns can also be active participants on the other side of the trade. The same machinery that allows a corporation with a low cost of capital to outperform competitors over time — issuing debt cheaply and deploying the proceeds into higher-returning assets — is available, at a much larger scale, to any sovereign who chooses to use it.

Consider the sovereign whose short-end borrowing rate sits at, say, 7 percent in his domestic currency. Such a sovereign is not unusual; many small and mid-size sovereigns have nominal fiat coupon costs in this range, reflecting some combination of fiscal credibility, currency risk,

and the absence of reserve-currency advantages enjoyed by the largest issuers. That sovereign has a choice he has not historically been forced to think clearly about. He can issue debt in his own currency at his coupon rate, fund whatever programs he funds, and accept the standard fiscal consequences. Or, on the margin, he can issue debt at the same coupon rate and convert a portion of the proceeds into a reserve asset whose purchasing power is not being systematically eroded — Bitcoin — and earn, over time, the spread between his nominal coupon cost and the rate at which the reserve asset compounds in real terms against the fiat unit he is issuing the debt in.

This is structurally identical to the corporate balance sheet operation that the most aggressive Bitcoin-treasury corporations have been executing for several years. Issue cheap debt, deploy the proceeds into a non-debasing reserve asset, capture the spread. The mechanics work at the sovereign scale just as well as they work at the corporate scale, and probably better, because sovereigns generally have access to far larger debt markets at far longer durations than any corporation does. A sovereign that holds Bitcoin as a reserve asset and continues to issue fiat-denominated debt is not destroying his currency. He is, on the contrary, strengthening his balance sheet against the debasement of his own currency, and capturing the asymmetry that is currently being captured by private holders, by corporations, and increasingly by sovereign-wealth funds operating in the same space.

The game theory follows directly. A sovereign that does this first gets to acquire the reserve asset at relatively favorable prices, against a competitive set of other allocators that does not yet include other sovereigns at scale. As more sovereigns recognize the trade and begin to execute it, the price of the reserve asset rises, the spread available to later movers shrinks, and the structural advantage of having moved early compounds. This is the standard first-mover dynamic of any reserve-

asset transition, played out in real time over the next several decades. The sovereigns who move first will, over a long enough horizon, have substantially stronger fiscal positions than the sovereigns who move last, regardless of the initial size or credibility of either. The advantage is not a function of currency status; it is a function of who acquired the appreciating reserve asset earlier in its monetization curve.

There is a second adjacent move that the game theory suggests will become important, which is sovereign participation in Bitcoin-native credit markets directly. A sovereign with cheap fiat funding and an accumulating Bitcoin position is naturally positioned to be a lender into Bitcoin-collateralized credit, deploying its capital base as a Bitcoin-secured lender at market-determined yields. This produces additional balance-sheet income, expands the sovereign's effective monetary toolkit, and contributes to the depth and stability of the emerging Bitcoin yield curve. The sovereign that does this is not merely accumulating a reserve. It is participating, on profitable terms, in the construction of the Bitcoin financial system itself. The first sovereigns to occupy this position will accumulate substantial structural advantages in the resulting equilibrium.

The largest sovereigns, those whose currencies enjoy reserve status, face a more complicated incentive structure. Their fiat issuance is held by foreigners as savings; the deal that has sustained the U.S. dollar's reserve role for the past half-century requires that the dollar be a reasonable store of purchasing power, even if not a great one. A reserve-currency sovereign that begins to accumulate Bitcoin signals that it is hedging its own currency, which threatens the very reserve status that has subsidized its borrowing costs for decades. This is not a trivial constraint, and it explains, I think, why the largest sovereigns have been the slowest to adopt. But the incentive to defect is strong, and grows stronger as smaller sovereigns demonstrate that the Bitcoin-treasury

operation is profitable and sustainable. The reserve-currency sovereign that holds out the longest faces, eventually, the worst version of the trade — accumulating the reserve asset at prices substantially higher than its earlier-moving competitors, after its currency's reserve premium has already begun to erode through the same dynamics it failed to participate in.

This is what I mean when I say the cost-of-capital reframing makes the game theory real. Sovereign adoption of Bitcoin is not a moral question, not a political question, not even primarily a monetary question. It is a balance-sheet question, of the kind every CFO understands and every corporate treasurer is trained to think clearly about. The framework of issue-cheap, deploy-into-appreciating-collateral is centuries old. The novel feature is that it now applies to sovereigns as well as corporations, and that the game-theoretic equilibrium it implies has very few stable resting points other than near-universal adoption. The path is uncertain. The endpoint is, in this specific structural sense, more inevitable than the conventional commentary acknowledges.

There is an important caveat that I want to be explicit about, because it cuts against the maximalist reading of this argument. The cost of capital that ultimately matters for any individual decision is not the market rate. It is the personal cost of capital — the rate at which the specific individual, with his specific opportunities, his specific time preferences, his specific tolerance for risk, is indifferent between present and future use of capital. The market rate is an input to that calculation; it is not the calculation itself. An individual whose own subjective cost of capital is higher than the market rate should borrow until his marginal use of borrowed funds equals the market rate. An individual whose subjective cost of capital is lower should lend. Either way, the personal calculation is what governs the decision.

What a Bitcoin-anchored cost-of-capital structure does is make the personal calculation honest. It removes the systematic distortions that the existing system imposes on different classes of actors. It puts everyone on the same footing with respect to the market price of capital, leaving the personal optimization to the individual. The market rate becomes a neutral reference. The individual decision becomes a function of individual circumstance. Both are improvements over the current arrangement, in which the market rate is itself a policy output and individual decisions are systematically distorted by the resulting incentive structure.

This is, I think, the single most consequential implication of the Bitcoin-anchored monetary order, and it is the implication that most directly connects the macro-political discussion of Parts I and II to the practical portfolio discussion of Parts III and IV. Bitcoin is sometimes described as “freedom money” in a vague libertarian sense. The specific freedom it provides, at the level of capital allocation, is the freedom from a systematically distorted cost of capital. That is not a small freedom. That is the freedom that determines what gets built, what gets funded, what gets foregone, across the entire economy. Restoring it changes what is possible at the level of every actor in the system, from the household to the sovereign.

The next and final chapter draws all of these threads together into the larger Hayekian argument that the entire book has been building toward.

CHAPTER TWENTY

The Hayekian Endgame

✦

The title of this book is *Endgame*. I want, in this final chapter, to be clear about what I mean by it.

I do not mean the endgame of fiat money, in the sense that fiat money will end on any particular timeline. Some uses of fiat will likely persist indefinitely — government employee salaries, tax collection, certain local payment functions — long after the macro-monetary architecture of the system has shifted. I have been careful throughout this book to avoid predictions of monetary collapse, currency reset, or any of the other dramatic scenarios that populate the more apocalyptic edges of the Bitcoin commentariat. Fiat regimes can persist for a very long time, even when their underlying conditions are catastrophic. Markets — and the institutions that sustain them — can remain irrational longer than any individual participant can remain solvent betting against them. They can persist longer than the working life of any allocator who is currently making decisions. The endgame I am describing is not a sudden discontinuity. It is a long, structural transition, of which we are perhaps a decade into a process that may take three or four decades more.

What I mean by endgame is the endgame of the project that began with Hayek's *The Use of Knowledge in Society* in 1945, was extended into monetary theory in *Denationalisation of Money* in 1976, was given a constitutional foundation by Buchanan and his colleagues in the 1980s, and is being operationalized, in working code and working markets, by the generation of Bitcoiners and Bitcoin-native financial infrastructure

builders working today. The endgame is the completion of the Hayekian project. Sound money plus market-discovered credit equals restored intertemporal capital coordination. The Hayekian critique of central planning, applied at the deepest layer of the economy, is being instantiated in technology and institutions that did not exist when Hayek was writing and that he could not have imagined in detail.

This is, I think, the right scale at which to understand what is happening. We are not living through a price cycle in a new asset class. We are living through the emergence of the institutional infrastructure for a different way of organizing capital allocation. (The network-effect adoption rates and the compound-annual growth in users and balances are real, even allowing for the small-base caveat that I have flagged in earlier chapters.) The infrastructure has been visible only in fragments. I have tried, in this book, to describe the fragments and to show how they fit together.

The fragments, restated:

A monetary asset whose supply is structurally exogenous to political deliberation, satisfying the requirement that no Hayek-compliant monetary system can do without.

A self-custody architecture that allows holders to participate in this monetary asset without depending on any custodian, intermediary, or institutional gatekeeper, satisfying the requirement that the asset's properties not be compromised by the holding structure.

A cryptographic credit infrastructure — DLCs, adaptor signatures, multi-party computation, and the rest of the rapidly developing Bitcoin-native financial primitive set — that allows credit to be enforced by mathematics rather than by trust, satisfying the requirement that the credit system not depend on the same trusted intermediaries whose failure modes have characterized every previous credit cycle.

A market-discovery process in collateralized lending that, as it scales, will produce the first market-discovered yield curve in any major economy since 1971, satisfying the requirement that the price of money over time be an output of market process rather than of policy deliberation.

A portfolio framework that incorporates these instruments in serious allocations, displacing the failing 60/40 model with a more structurally sound model that is appropriate to the monetary environment of the next several decades, satisfying the requirement that practical allocation move in step with the underlying institutional development.

Each of these fragments exists today. None of them existed twenty years ago. Together, they constitute the infrastructure for a different monetary order — not by overthrowing the existing one, but by building a parallel one that, over time, becomes increasingly the natural choice for rational market participants who can choose. The transition is voluntary, gradual, and reversible at each step for any individual participant — and yet each step from getting off zero to optimizing one's allocation tends to leave the participant better positioned than the step before. It is not, however, reversible at the systemic level. Once the parallel infrastructure exists at sufficient scale, the dependency of individual participants on the existing system can be progressively reduced. Once enough individual dependencies are reduced, the system as a whole loses the captive base that has historically allowed it to perpetuate itself in the face of its own contradictions. The revolution, in this sense, will not be televised.

This is what Hayek meant, I think, by the spontaneous order. He did not mean that order arises from chaos by magic — though, as Arthur C. Clarke observed, any sufficiently advanced technology is indistinguishable from magic. He meant that order arises from the cumulative result of decentralized choices — and decentralization is, for me, the ultimate

point, the issue underneath all the other issues — made by many participants, each pursuing their own ends (in a formulation that perhaps sounds more like Adam Smith than Hayek, though both were making variants of the same point), within a framework of rules that is itself either spontaneously ordered or constitutionally fixed. The Bitcoin monetary order is spontaneous in this Hayekian sense. It was not designed by any committee. It is being built by a vast number of individual choices — to hold Bitcoin in self-custody, to use Bitcoin-collateralized credit, to denominate plans in Bitcoin terms, to construct portfolios with meaningful Bitcoin allocations — none of which is coordinated centrally. The cumulative result is an order that has been emergent, that nobody planned, and that is, on the analysis I have laid out, increasingly the natural attractor for participants who understand the alternatives. It's a live process connecting more and more people in real time.

The work of building this infrastructure is the most interesting work I have ever been part of. It is not interesting because it is lucrative, although it may turn out to be. It is interesting because it is the work of completing a project that great thinkers have been articulating for nearly a century, and that has, until recently, been unable to be operationalized for lack of the technical primitives. The primitives now exist. The institutions are now being built. The portfolios are now being constructed. The yield curve is now beginning to emerge.

If this analysis is right, the next twenty years in monetary economics will be more interesting than the previous fifty. I do not know whether we will see the Hayekian endgame fully realized in our working lives. Every day I work on this, I am more persuaded that we will see it substantially advanced, and that the financial planning profession, the institutional allocation industry, and the broader economic system will all be meaningfully transformed by what is being built. Bitcoiners who

position themselves and their portfolios accordingly will participate in that transformation. Those who do not will be participants in the existing system whether they intend to be or not.

The choice, for any allocator who has read this far, is what to do with the analysis. The framework is now in your hands. The position-sizing question, the form-of-holding question, the fixed-income replacement question, the productive-use question — each of these has answers that follow from the framework I have laid out, and each requires the allocator to make decisions for his own portfolio. I have not tried to make those decisions for you. I have tried to give you the framework that, applied honestly, produces decisions that you can defend — because you might have to.

That is the most a writer can do. The rest is up to you.

Afterword

A Note on Inevitability

I have deliberately avoided, throughout this book, the language of inevitability. I want to close by saying explicitly why.

The Bitcoin commentariat is heavily populated with claims that Bitcoin's success is inevitable, that fiat collapse is inevitable, that some particular price target is inevitable on some particular timeline.

Claims of inevitability mis-state the actual case for Bitcoin. The case for Bitcoin in a serious portfolio does not require inevitability. It requires only asymmetry. A non-trivial probability of substantial outperformance, against bounded downside, is sufficient to justify a meaningful position size in any portfolio that is exposed to the failure modes Bitcoin hedges against. The asymmetric-payoff framing is intellectually honest, falsifiable in principle, and sufficient to drive the portfolio decisions that actually matter. The inevitability framing is none of these things. If your investment thesis required inevitability to motivate the position, your investment thesis was probably weaker (or had a lower expected return) than you realized.

The framework I have laid out in this book does not depend on Bitcoin's success being inevitable. It depends on Bitcoin's success being plausible enough, against the failure modes of the existing system being severe enough, to justify the allocation. If the analysis in Part I is even approximately correct, the case for the allocation does not require any particular outcome. It requires only that the underlying conditions described — the Hayekian knowledge problem applied to money, the Cantillon mechanism, the productivity capture, the international

evidence — continue to operate as they have been operating. They show no signs of stopping.

I have also avoided, throughout this book, predictions about timing. I do not know when the transition I describe will be substantially advanced. I do not know when a Bitcoin-collateralized yield curve will be a recognized market structure that allocators routinely consult, available in custodied institutional portfolios, or legally recognize smart contracts as collateral. I do not know when the 60/40 portfolio model will be retired from the textbooks. I have my own working estimates, but they are estimates, and I would not want any reader to make decisions based on them. The framework is robust to a range of timing outcomes; the decisions an allocator makes today should not depend on any particular timing forecast being correct.

What I do believe is that the framework is durable, that the trajectory is clear in direction even if uncertain in pace, and that the asymmetric-payoff allocation case is sufficient on its own terms. An investor who acts on the framework is, in my view, well-positioned across a wide range of possible futures. An investor who ignores the framework is exposed to risks he may not fully appreciate.

That is the case I have tried to make. I have tried to make it without overclaiming, without leaning on inevitability, and without predicting any particular outcome. The framework either persuades you or it does not. If it does, the decisions follow. If it does not, you have at least encountered a serious version of an argument you may have only previously seen in weaker forms.

I look forward to the next several decades of this project, and I look forward to whatever role each reader chooses to play in it.

Kevin Bell Founder, Cadena Bitcoin

A final note, and then I am done.

Nothing in this book is financial advice. I do not give financial advice. I do not run an advisory practice, and I have specifically avoided framing this book as guidance to be acted on without your own analysis. The framework is a framework. The decisions are yours.

I want to close on something that has been implicit throughout the book and that I want to make explicit now. Markets are not efficient. Information is not perfectly distributed. Markets are aggregations of human judgment under uncertainty, and human judgment under uncertainty is sometimes wrong, sometimes systematically so, and sometimes for very long periods.

But here is the part that Hayek understood that the efficient-markets people did not. The price system is not perfect. The price system is the best information-aggregation mechanism humans have ever devised, and it is far better than any alternative, but it is not perfect. It contains systematic errors. It can be wrong for years. What the price system gives you is not a guarantee of correctness. It gives you the best available signal about what every other participant in the market currently believes, weighted by how much capital they have at risk on the belief.

That is enormous information. It is more information than is contained in any analyst report, any expert opinion, any government statistic. It is more information than is contained in this book. The price of an asset is a statement of value made by people putting their own money behind their statement, which is a far more reliable indicator of what they actually believe than anything they say in writing.

The price of Bitcoin, in particular, is a statement that should be listened to carefully. It has been telling a coherent story for sixteen years, through multiple cycles of skepticism, through repeated obituaries written by serious people, through every plausible scenario for its failure. The price has continued to ratchet higher over time, in long cycles of advance and retracement, as more and more capital has been

put behind the proposition that this asset matters. You can disagree with that proposition. You should, if you find yourself disagreeing with it, articulate why. But you should also recognize that the price contains information, and that the information is not random.

Listen to the price. Not in the sense of tracking the daily quote. In the sense of asking, periodically, what the long-term trajectory is telling you about the cumulative judgment of the millions of participants who have, with their own capital, been arriving at conclusions about what this asset is worth. The trajectory is the message. The trajectory has been remarkably consistent. The trajectory is the best evidence you have for whether the framework I have laid out in this book is, on balance, correct.

I am persuaded that it is. I am persuaded that the price has been telling the story for sixteen years. I am persuaded that listening to it carefully, and acting on it within the framework of one's own situation, will turn out, in retrospect, to have been the most consequential single financial decision most readers of this book will make in their working lives.

I could be wrong. I do not believe I am, but I might be. The price will continue to tell us.

Seriously, listen to the price.

Kevin Bell Founder, Cadena Bitcoin

Suggested Reading

The thinkers and writers I have built on, or argued against, in this book. Read them; argue with them; come back with sharper questions than the ones I have posed.

The Bitcoin Standard (2018) by Saifedean Ammous remains the foundational text for the Austrian-Bitcoin synthesis. Ammous is occasionally polemical, but the underlying analysis is rigorous, and the historical sweep of the argument has not been surpassed.

Broken Money (2023) by Lyn Alden is, in my view, the best single book on monetary history, monetary mechanics, and the case for a Bitcoin-anchored future. Alden writes with more nuance than almost anyone else in the space, and her treatment of debt dynamics and Eurodollar mechanics is essential reading.

The Price of Tomorrow (2020) by Jeff Booth makes the deflationary case for technology and the catastrophe of trying to fight that deflation with monetary expansion. The argument is short and sharp and has aged well.

Bitcoin Is Venice (2022) by Allen Farrington and Sacha Meyers is the most intellectually ambitious recent treatment of Bitcoin as a civilizational technology rather than as an asset class. Farrington's later essays extend the argument into territory adjacent to the one this book occupies.

The Use of Knowledge in Society (1945) by Friedrich Hayek is the foundational paper on the price system as an information-aggregation mechanism. Fourteen pages. Read it twice.

Denationalisation of Money (1976) by Friedrich Hayek extends the argument into monetary theory. Hayek imagined competing private

monies; Bitcoin instantiates a single rules-based money. Either way, the underlying insight is Hayek's.

Prices and Production (1931) by Friedrich Hayek lays out the Austrian theory of the business cycle and the central role of intertemporal coordination. Demanding but worth the work.

The Theory of Money and Credit (1912) by Ludwig von Mises is the foundational Austrian work on monetary theory. The English translation by H.E. Batson is the standard version.

Good Money (2008) by George Selgin reconstructs the history of free banking in Britain and the United States. The empirical case that money does not require central banking is here.

Free Banking in Britain (1984) by Lawrence White covers similar ground for the British case in particular. White and Selgin together constitute the modern free-banking school.

The Calculus of Consent (1962) by James Buchanan and Gordon Tullock founds the constitutional political economy tradition. The distinction between constitutional and post-constitutional choice is essential for understanding why rules-based money is not just preferable to discretionary money but categorically different.

Essai sur la Nature du Commerce en Général (1755) by Richard Cantillon is the original observation of what we now call the Cantillon effect. The Henry Higgs translation is widely available.

What Has Government Done to Our Money? (1963) by Murray Rothbard is the polemical short version of the Austrian monetary critique. Read it for the rhetoric; read Mises and Hayek for the analysis.

The Sovereign Individual (1997) by James Dale Davidson and William Rees-Mogg anticipated, in remarkable detail, much of the digital monetary order that is now emerging. Some of its specific predictions are dated; the structural analysis has aged well.

Layered Money (2021) by Nik Bhatia is the cleanest exposition of the layered monetary system — base money, bank money, shadow money — that any reader of contemporary financial commentary should understand.

Lifecycle Investing (2010) by Ian Ayres and Barry Nalebuff is the rigorous case for early-life-cycle leverage in long-horizon investing, and the framework I draw on in Chapter 11.

Thinking in Bets (2018) by Annie Duke is the cleanest treatment of decision quality versus outcome quality I know of, and is referenced in Chapters 11 and 14.

Shelling Out: The Origins of Money (2002) by Nick Szabo is the essay that grounds the evolutionary frame for monetary technology in Chapter 7. Available freely online.

There is also a growing body of technical literature on Bitcoin-native financial primitives — discreet log contracts, adaptor signatures, taproot, and the like — most of which lives in white papers and blog posts rather than in books. The interested reader is referred to the writings of Tadge Dryja, Lloyd Fournier, and the Bitcoin Optech newsletter as starting points. These sources will become dated quickly; the underlying ideas will not.

Acknowledgments

I owe substantial intellectual debts to the writers in the preceding section, particularly to Ammous, Alden, Booth, Farrington, Foss, and Szabo, whose work shaped my own thinking and whose books should be read alongside this one rather than as alternatives to it.

I owe debts of a different kind to my colleagues and collaborators at Cadena Bitcoin, who have built the infrastructure that makes the productive-layer arguments in Part IV concrete rather than theoretical. Any errors in the technical descriptions of DLCs, adaptor signatures, or related primitives are mine; any insight into how these primitives can be assembled into actual credit markets is theirs.

I owe debts to the clients and prospective clients of Cadena who have, over many conversations, sharpened my understanding of the questions that this book attempts to answer. Their questions were better than my initial answers, and the book is better for having had to address them.

I owe debts to the Bitcoin community broadly, which has been the most passionate, stimulating, and intellectually serious community I have ever participated in, and which has taught me most of what I now take for granted about money, credit, and capital.

Finally, I owe debts to my kids, my parents, and my brother, who have helped me through challenges while keeping me grounded in the reality that exists outside the work. The book exists in part because they made it possible for me to keep working on it.

Most importantly, thanks to God.

Kevin Bell Founder, Cadena Bitcoin