MOOSE.KACHAT — @THEMOOSEISLOOS5

Satoshi’s Engine 2

Bounded self-regulation

Book One gives states. This expansion opens the gradients beneath them through coins, miners, machines, fees, volume, waiting and cycles.

CONTENTS

Foreword

What Is Taken as Given

Part One — The Base and the Heat

1. The Base That Grows

2. Nobody Exits at the Quoted Price

3. What They Are Buying Is Time

4. Urgency Plays Like Poker

Part Two — The Two Streams

5. Debasement Without Inflation

6. The Factory Answers Late

7. Where the Pressure Settles

Part Three — The Handover

8. What Pays the Bill

9. Nobody Schedules the Handover

10. Knowing Is Not Moving

Closing

What Remains

FOREWORD

What Is Taken as Given

This book continues Satoshi’s Engine: Carnot local, Brownian global. It does not stand on its own and assumes Book One has been read.

Book One described the engine in states: hot and cold, expansion and contraction, issuance and fees.

Book Two follows the gradients underneath those states.

Beneath the cycle Book One left two ratchets, one in the machines and the work behind them, the other in the coins that stop being sold.

This is the difference between a loop and a ratchet. Carnot names the local engine; Brownian names the global movement that carries its residue forward.

What happens when existing coins begin to fund a material share of the security work that issuance has carried?

The cycles Book One described can be observed. Nothing in this book has been: no network has run a sustained regime in which fees carry the work. What follows is not a report of what has happened. It is a derivation of what self-interested participants and fixed physical rules produce if it arrives.

PART ONE

The Base and the Heat

These chapters establish the conditions under which fees can appear: the surviving base of users and connections, the market's ability to move coin and purchasing power, the time decisions made by participants and the sorting of urgency.

CHAPTER ONE

The Base That Grows

Every cycle leaves a base of demand behind it: users, connections and hubs that can produce activity even when the visible flow is quiet. It is not a promise that every participant remains active. It is the accumulated ability to start again.

A cold market is not an empty network

When price falls, activity contracts. Holders stop selling, buyers wait, miners switch off and fees decline as fewer actions remain urgent. The coin also buys less: goods and services cost more of it than before. From a holder’s point of view, spending the same balance now purchases less.

The cold phase does not erase the network’s possible uses. A person who stops transacting still knows how to transact. A merchant who sees fewer customers has not forgotten the payment rail. A developer who pauses a project has not destroyed the route already built.

A user may still spend through a hub when it offers something that cannot be reached through another route. Falling purchasing power reduces optional activity first; it does not erase demand for a unique service or opportunity.

The visible flow can shrink while the base that can produce a later flow remains.

Connections multiply the surface

One user adds one possible participant. A second adds a possible connection between them. A third adds several more pairings, depending on what each person can offer or obtain.

The network becomes useful through connections, not headcount alone. The particular use can change without destroying the underlying ability to connect.

The informational ratchet is the third, and that larger surface is what it preserves: learned routes, tested limits, surviving connections and builders who no longer start from zero. Once a route has been learned and a use case proved, the network does not return to not knowing it when the route closes. Price can reset locally; the network does not return to having only the connections that existed at the previous bottom.

The later fee regime does not appear from nowhere. It requires the base preserved by the informational ratchet: routes and surviving connections from which recurring urgency can arise.

Hubs carry the connections

Connections need somewhere to meet. A hub can be a service, a market, a merchant network or a programmable structure built on the base layer.

A hub can serve existing demand, but it can also create the first reason for demand to appear. In an early network there may be nowhere to price the coin and nothing to do with it. A place to exchange it gives the coin a price. A place to spend it gives that price a reason.

The earliest hubs are often the most fragile and the most important. Their first function is to prove that a use case exists. That proof brings builders toward the next, better hub. A failed hub can leave users who now know what they wanted and may take that demand elsewhere. What the first hub proved can outlive the hub itself. Survival is not the measure of whether it mattered.

A failed hub can also reveal a limit. If users want the interaction but the route cannot bear its cost, delay or complexity, the failure shows what the next route must change. The network does not need to rediscover that limit by repeating the same attempt. Success proves that a path can be walked; failure can show where the ground gives way.

A route can still fade. A path worn across grass grows back when nobody walks it; an interface is forgotten and a habit lapses when the urgency moves elsewhere. The route decays faster than what it proved. The ratchet does not reverse, but its surface can become harder to reach. If no urgency returns anywhere, even the proof can eventually go quiet.

Competition between hubs can therefore replace a route without erasing the base. Hubs reveal what users want and what interactions are possible. Fees reveal how much users will pay to settle those interactions within a particular time.

Adoption arrives in waves

The first wave can arrive because price moved, but price does not always precede the first wave. The hub can come first, and the price with it.

Later users may arrive because a merchant accepts the coin, a market opens, a service becomes easier to reach or someone they already deal with is present. These are different reasons for entering the same network.

Speculative demand can disappear quickly in a cold market. A useful connection can remain available at lower heat, even if it is used less often. The later wave is more durable because it does not require every participant to be waiting for a higher price.

The base is not the flow

A larger base does not guarantee a higher price or a rising fee market. Participants can remain inactive, connections can fail and hubs can be outcompeted. The base is a condition for future flow, not the flow itself.

The base still produces a residual flow between larger waves: small transfers, occasional settlement and activity that is too weak to change a miner’s hardware decision. That is noise until its size and persistence make it a signal. A fee becomes a signal when it is meaningful enough to justify continued operation, capital expenditure or new machines.

The accumulated base matters because it lowers the cost of starting again. The next wave can use existing knowledge, connections and surviving routes instead of explaining the whole network from zero.

What survives the cold phase

The cold phase moves excess out of sight without destroying it. Weak miners leave, holders withdraw supply and hubs lose users or close. The visible flow becomes smaller.

What remains may still include more participants, possible connections and accumulated experience than the previous cycle had at the same price. The local loop returns to cold. The global starting point moves.

A network can therefore be quiet without being empty, and a later wave can move through it faster than the first one did.

CHAPTER TWO

Nobody Exits at the Quoted Price

Price and machines move through the cycle in plain view. The missing variable is how much of the system can change hands at the current price.

Market throughput is how much coin and purchasing power can change hands without forcing price sharply away from where it was. Settlement capacity is how much transaction demand the protocol can finalise in a given time. One determines how much value can move before repricing; the other determines how much movement can be completed without the charge rising. Either can become the constraint while the other still has room.

The market has to receive the heat

A large holder cannot exit simply because the quoted price is high. The position has to be absorbed. If willing buyers are scarce, selling lowers the price before the whole position changes hands.

The reverse is true at the bottom. A small amount of buying can move price a long way when little supply is available. That movement is not proof of strong demand. It may be the visible effect of a market too thin to distribute the order gradually.

What appears to be strength is low throughput.

One-way flow through changing volume

Trading is often described from the fiat side: someone buys the asset, waits for appreciation and sells it for a gain. From the asset’s side, nothing has been reversed. The buyer moved into the asset while the seller moved into purchasing power. Both positions changed, and the next trade begins from those new conditions.

The market is a crossing of continuous flows. Coin units move toward holders, users and miners while purchasing power moves toward whoever releases them. When the flows balance at the margin, price can remain

stable. When one side moves faster, price travels through the available inventory to find the next participant.

Persistent demand therefore does not produce the same price movement in every market. A thick market offers a wide path. A thin market offers a narrow one, so the same purchasing power must move price further to find supply. Available volume is the width of the path; price impact appears when it narrows.

The path is narrow by default. What is held, and could be released, is larger than what trades by orders of magnitude, so price is set at any moment by a small fraction of the holders. Thinness is not a brief condition on the way to a thick market. It is where a cold market rests.

Holding removes coins from active circulation. Price impact rises, and persistent demand produces a larger movement without necessarily becoming stronger. The response steepens as the path narrows.

The direction can reverse when older holders reach their profit tolerance and release inventory, and the release grows faster than the price that drew it out. Volume expands, price impact falls and the same demand meets more supply. Nothing stopped the move from outside. It ended itself.

In a cold market, holders have reasons to wait. Someone below cost does not want to realise the loss; someone still above it may not want to sell into a market unable to receive the position. Miners remain the exception because electricity arrives as a continuous bill. They sell enough to continue until the work no longer pays.

How buyers compound volume

When heat returns, the first movement may still be thin. The larger shift begins when volume returns. Higher prices give older holders a reason to sell. Their sales create inventory for new buyers, and new buyers make it possible for still larger holders to sell.

Buyers therefore do more than bid price upward. They create the conditions under which more inventory can enter without stopping the rise immediately. Selling during a hot phase is not necessarily evidence

that heat has ended; it may show that the market has become thick enough for older positions to move.

Volume is what lets opposing intentions coexist.

The machine market has a slower channel

The machine market is another inventory market, but Book One already traced its full cycle: new hardware arrives late, used machines undercut it, and the used pool eventually runs dry. What matters here is the speed. A coin can change hands in seconds while new capacity takes months to design, build, ship and install.

Coin-market throughput and machine-market throughput therefore do not synchronise. Price can turn down while hardware ordered during the previous hot phase is still arriving. The two channels run on different clocks.

Settlement is a different constraint

Settlement capacity is limited in a different way. A coin can have deep market liquidity while settlement is expensive, or thin market liquidity while settlement is cheap. Market throughput is how much the market can receive. Settlement capacity is how much can be finalised, and the contest for it is sorted by urgency rather than by price alone.

The coin market prices the unit in which settlement is paid. A fee quoted in that unit carries the coin’s market value with it. When the coin becomes more valuable, the same nominal fee costs more to the user and delivers more purchasing power to the miner. When settlement becomes urgent, the fee market can add a further premium.

The coin market prices the settlement instrument. The settlement market prices immediate access to it.

What heat leaves behind

A hot market creates more participants, a wider distribution of positions and more machines that can later be redistributed. Coins sold into

strength become someone else’s held inventory. Machines sold at a discount become someone else’s productive capacity.

Cooling removes market throughput. Holding takes coins out of active circulation. Switching machines off takes capacity out of active production. The next movement may be easier to start because less inventory is in the way, but harder to absorb because fewer participants remain.

Price reports the current crossing. Volume determines how much of the system can pass through it.

CHAPTER THREE

What They Are Buying Is Time

Every participant in the cycle trades time for something, or something for time. A miner buys a shorter wait with capital. A holder buys the option to wait by refusing to sell. A user pays a fee when moving now is worth more than waiting.

The useful question is not whether time has value. It is who is paying to avoid it, and what they receive.

The miner pays first

Time is not a mood in this system. It is a cost with several owners.

An expensive machine commits more capital before it has produced a coin. The premium buys a shorter expected wait before the return arrives. That return is never guaranteed: new hardware arrives, difficulty changes, issuance thins and other miners make the same calculation.

A machine can be earning money while losing time, because each month of operation is also a month closer to obsolescence. Doing nothing is free for a holder and expensive for a miner, and the expense is not only the bill.

The buyer pays differently

A coin buyer does not need to build a machine. The work has already happened. The buyer pays for access to an existing result: a network, a history and an amount of work that would take time to reproduce.

The buyer may also be paying to avoid a future wait. If the network becomes more useful, harder to reproduce or more difficult to enter, owning the coin now keeps an option open. Its value depends on what the holder can do with it later, not only on what the miner spent earlier.

Holding is stored time

A coin not sold is a decision not yet made. What the holder keeps is not only the coin but the choice of a later moment.

Demand need not become enormous when few holders are ready to sell into it. A small wave can move through the remaining inventory. Every holder carries a private estimate of the price and time at which selling will become worthwhile. When price compensates them for ending the wait, their coins become inventory for newer buyers.

The seller receives present purchasing power. The buyer receives exposure to a future result. Neither side is buying the same thing.

The user pays to resolve time

A transaction fee buys settlement time. The user pays to reduce the delay before a desired result: a trade, a purchase, access to a service or an opportunity that disappears if it waits.

If the benefit is smaller than the fee and the delay, the transaction waits. If the benefit is larger, the user pays. Two transactions on one network can therefore pay very different fees. A coffee can wait. An arbitrage opportunity cannot.

The fee becomes miner revenue. The user buys time, and the miner receives payment for producing the settlement that closes it.

Time moves through the cycle

The same market prices two different things depending on its temperature. When it is quiet, price settles toward the work behind the coin, because competition drags it there and nothing is urgent enough to pull it away. When it is hot, price moves toward what time is worth to whoever cannot wait. Neither reading is the truer one. It is one mechanism, pricing whichever scarcity is binding.

At the bottom, time belongs to holders. Selling offers little benefit, while miners cannot postpone their bills.

On the climb, time becomes more valuable to buyers. Price rises before new hardware can arrive, and coin ownership becomes the only available route to exposure for demand that would otherwise have gone into machines.

At the top, miners commit to machines to shorten the return while holders sell because price is finally high enough to end their wait. In the fall, buyers demand a larger discount for accepting the wait and miners sell more of their production because their bills arrive on schedule.

The cycle redistributes who can wait and who must pay to avoid waiting. What sustained redistribution then does to the floor is a separate question; it begins with this decision.

The cost that cannot be seen on a receipt

Every decision has a visible price and an invisible schedule. The machine price says what was paid upfront, but not how long the machine will remain competitive. The coin price says what ownership costs now, but not how long the holder will wait before selling. The fee says what settlement costs now, but not what opportunity would have been lost by delay.

A miner may be buying months, a holder selling years of optionality and a user buying minutes. The engine compares all three through the same question: what is the cost of waiting, and who can bear it?

CHAPTER FOUR

Urgency Plays Like Poker

Fees do not wait for a completely full network. They rise as reasons to move compete for limited settlement capacity, and the response steepens when capacity becomes the pinned variable.

At low density, most transactions can be carried without competing strongly. A fee still deters marginal activity, but users often pay little because few are bidding against one another. As activity becomes denser, the fee measures urgency.

The first pressure is selective

Not every transaction values confirmation equally. A coffee can wait longer than an opportunity that disappears in seconds. The fee market lets each user compare the benefit of moving with the cost of delay.

The low-value transaction waits. The high-value transaction pays. Between them is a continuous range of urgency, not a switch between use and no use.

The charge is paid per settlement and not per unit of value, so a small transfer pays as much as a large one. That is what batching avoids: fewer settlements, each carrying more.

This is selective deterrence: it keeps every possible action from demanding immediate settlement at the same time.

Urgency plays like poker

Every participant makes a wager against the clock. An urgent participant raises because the opportunity may disappear. Another holds because waiting is acceptable. A low-value transaction folds when its benefit no longer covers inclusion. Someone can leave the table entirely by deferring the action or abandoning the opportunity.

The pressure can point in either direction. One user pays to enter before access closes; another pays to exit before a position deteriorates. Both turn time into a fee.

Density changes the response

Early demand can be absorbed with little change. The next wave may find fewer convenient places or shorter waiting times and begin bidding against activity already in motion.

The network is still carrying transactions. What changes is the price of insisting on a particular time. A fee market can matter before every available unit of capacity is occupied.

Capacity is not measured only by transaction count. Transactions can consume different amounts of settlement capacity, and users can attach different value to being included now. A few heavy or urgent transactions can therefore create meaningful fee pressure while ordinary transfers still appear plentiful.

Pressure raises the boiling point

Put water and air in a sealed vessel at ordinary pressure and heat it. The water warms, then boils. Added heat now goes into changing state rather than raising temperature, which is the kink: the same input stops producing the same effect.

But the vessel is sealed. The steam has nowhere to go, so pressure rises, and water under pressure boils at a higher temperature. To keep it boiling the temperature has to keep rising. The vessel has answered the heat by making heat less effective.

The base of latent demand is the water. Urgency is the heat. Transactions entering the queue and bidding against one another are the boil, and the clearing fee is the pressure they produce.

Raise urgency and more transactions compete. The fee rises. A higher fee means a transaction now needs more urgency than before to be worth sending at all. The bar has moved, exactly as the boiling point moves.

Sustaining the same volume of bidding takes more urgency than it did an hour ago.

Take the heat away and the steam condenses. Deterred transactions are not destroyed. They wait, they batch, they are sent later. They return to the pool they came from, and the pressure falls back toward where it started.

Nothing switched to a new mechanism at any point. The pressure is the same thing throughout. What changed is how much urgency it takes to add to it. The vessel fixes the relation between its pressure and its boiling point. Nothing fixes the fee market's, so the shape carries over and the quantities do not.

Urgency can keep the market hot

Saturation does not mean every transaction stops. A profitable trade can justify a fee that makes an ordinary transfer irrational. The fee can become highly deterrent without becoming absolute: it excludes marginal movement while permitting movement whose benefit exceeds the charge.

The deterrence also limits the fee. Nobody pays more to settle than the settlement is worth, and nobody pays indefinitely when waiting, batching, abandoning the action or using another route costs less. As the charge approaches the value of moving, marginal users leave the queue. The fee market therefore weakens the demand that created the fee.

The network is not closed. It is selective.

Density is not headcount

More users do not automatically mean more fee pressure. Density depends on how many want to move, how often they move, how many connections they can make and how urgently they need the result.

A large population holding quietly can produce little fee revenue. A smaller population trading intensely can produce much more. A hub can increase density by giving users a reason to interact, while a failure can reduce it even if the users remain.

The base creates the possible surface of activity. Density appears when enough of that surface becomes active at once, and that activity competes for settlement.

When pressure falls, some activity was merely waiting and returns. Other activity disappears because its opportunity passed. The fee stream reveals which reasons to move survive their own price.

The informational ratchet shapes which reasons survive. A network that has already proven certain routes and connections does not generate the same fee pressure as one that is still discovering them. Learned uses produce recurring reasons to move. Speculative or experimental uses produce more sporadic ones. The base that survives a cold phase therefore does more than preserve headcount; it preserves the kinds of urgency that can later fund security.

Density creates pressure. Urgency decides which pressure becomes payment.

PART TWO

The Two Streams

These chapters follow the central mechanism: redistribution changes the monetary floor, sustained fees can reach the machine frontier, and cooling pressure settles into machines, held coins or new capacity.

CHAPTER FIVE

Debasement Without Inflation

Fees can lower the monetary floor without expanding the stock of coins. The distinction rests on two flows underneath that movement.

Book One could speak of the floor as one thing because issuance carried nearly all of the miner’s reward. Once fees carry part of it, the physical work and the monetary burden placed on the coin price have to be counted separately. The machines still require the same joules, but the revenue can arrive through issuance or fees.

The joule limit belongs to the machine, and the market does not set it. The monetary floor is the coin price required to carry the current work after fee revenue is counted. When fees diminish, the fee offset shrinks and the monetary floor moves back toward where issuance alone would put it. When fees grow, the monetary floor can fall without the joule limit moving.

Two rates, two burdens

There are two rates to keep separate.

The distribution rate is the rate at which new coins enter through issuance. The redistribution rate is the rate at which existing coins move from users to miners through fees.

One changes the stock. The other changes who controls the stock and how much of it is active.

Each rate leaves a burden. Issuance alone sets what every coin must carry; fees decide how much of it is paid by coins that already exist. The burden is what a coin has to carry. The floor is the coin price at which that burden can be carried.

The stock does not tell the whole story

Stock-to-flow describes how quickly new issuance can add to existing supply. It does not show how quickly existing coins move through the

security budget. Issuance can decline while activity grows: a fee adds no coin to the stock, but moves an existing coin from a user to a miner. To the miner, both units pay the same electricity bill.

Debasement without inflation

Two different things are called inflation and they are not the same thing. One is growth in the supply of money: more units in existence than there were. The other is a rise in the prices of goods and services: each unit buying less than it did. A currency with a finite supply can suffer the second without the first.

This book uses the word in the first sense only. Where these pages say inflation, they mean coins created. Where a coin buys less than it used to, that is a change in purchasing power, which has its own causes and is not what the title of this chapter denies.

Suppose fees become equal to issuance while the block work stays the same.

The miner now receives twice as much revenue for finding the block: half from new coins and half from recycled coins. The physical work has not become cheaper. The joules per unit of work have not changed. But the miner no longer needs the same coin price from issuance alone to pay for that work.

That is debasement without inflation.

The effect is not automatically a lower market price. Demand can grow, holders can absorb the redistribution and miners can reinvest it. The claim is narrower: the price required to sustain a given level of mining has fallen, all else equal.

The machine cannot distinguish the two streams. Issuance changes the stock; fees change its path. Both pay for the same joules.

How long the stock takes to pass through miners

The redistribution rate can be expressed against the existing stock.

If fees recycle one percent of the total stock each year, the stock passes through miners once every hundred years at that simple rate. At two percent, once every fifty years. This is a comparative turnover measure, not a literal claim that the dormant portion of the entire stock will pass through miners on that schedule. Fee-paying coins come disproportionately from active circulation.

Those are turnover times, not inflation rates. No new coins are created by the recycling itself. The comparison matters because it shows how much purchasing power the active economy is transferring into the security budget.

The market decides what is enough

There is no predetermined optimal redistribution rate. If an opportunity is worth resolving, users can pay more to move. If moving is worth less than its charge and delay, the transaction waits or disappears. The fee stream observes unresolved demand; it is not a policy target.

When redistribution becomes the signal

At the beginning, issuance is the clearest signal to miners. The protocol offers new coins, and the market determines what those coins are worth.

Later, fees can become large enough to matter more than issuance. The transition is not defined by a date. It begins when fee revenue exceeds issuance revenue for long enough to influence machine decisions.

Issuance becomes the noise that fees once were.

If the fee persists, the extra margin attracts machines. Competition divides the work, the advantage narrows and the temporary debasement is carried into a larger productive base. If the fee disappears quickly, the fast user response is all that remains.

The redistribution rate can therefore lower the monetary floor now and raise the work behind future coins later.

What the concept does not claim

Debasement without inflation does not mean every fee market harms holders. The fee may be paid because the network has become more useful. Nor does it mean structural scarcity has vanished: redistribution changes the path through the stock, not the number of coins that exist.

The title denies growth in the supply of one scarce thing, and the protocol makes two: coins, and settlement capacity. Fees add to neither. A fee moves purchasing power from a holder to a miner; the value stays inside the system and the work it pays for continues.

The precise claim is simpler:

Distribution changes the stock.

Redistribution changes the flow through the stock.

Debasement changes the monetary floor without requiring inflation.

Fees are therefore not merely an extra reward. They are the market’s way of changing the burden carried by the coin price while keeping the security work in motion.

CHAPTER SIX

The Factory Answers Late

The mining-side ratchet runs underneath a local loop. Price can reverse, but machines ordered during a profitable phase arrive later and compete in conditions that may already have changed. Fees add a second stream to that delay. That stream acts on three clocks. The monetary clock moves first, the capital clock next, the factory clock last.

The machine frontier is the leading edge of deployed efficiency: the newest machines operating at the lowest joule limit available to the network. It advances only when a revenue signal survives long enough to become capacity.

What passes one way

The profitability gap changes sign. When price sits above the cost of production, machines are added. When it falls below that cost, marginal machines leave. Price can reverse through a full cycle; accumulated work and productive capacity do not reverse with it.

Two features of the system produce that result. Difficulty adjusts, so work that leaves is redistributed to the work that stayed: an operator who shuts down hands its share to every operator that did not. The machine is already bought, so the hurdle for continuing is the electricity bill and not the price of the hardware.

Book One named the pawl that keeps accumulated work from unwinding. These two are what keep it running while price falls.

The floor is a level the system returns toward, not a barrier the price cannot pass. It is set by what the marginal machine must cover, so it moves when that machine changes. A price below the cost of production is not a contradiction. It is a statement about which machines are still running.

When weak operators leave, the survivors are more efficient and the floor can move with them. How far it moves depends on how closely their costs sit together.

Where operators of similar efficiency stand close, one departure hands the rest a share so small that the next marginal operator is postponed rather than relieved. The floor holds while the fleet is packed, and it can fall a long way at once when the packing runs out. Where costs are spread more widely, the floor simply follows the price down, and nothing gives way at any particular level.

The lag makes the ratchet visible. Machines are ordered against one profitability gap and arrive against another. Delayed production lets each profitable phase load the next one.

The monetary clock

When fees enter a block, the miner receives more coins for the same search. The machine has not become more efficient and the joule limit has not moved. Another source of purchasing power is carrying part of the bill, so the monetary floor moves first.

A miner compares the expected value of issuance and fees with the electricity, hardware and operating cost of the work. A persistent fee stream lowers the coin price required for that comparison to clear.

The miner can sell, hold, pay costs or commit capital. A fee event can be large without being durable. The current miner responds immediately; the machine frontier waits.

The capital clock

A persistent fee stream changes the comparison between old and new machines. It can keep older machines online because their operators no longer need the coin price from issuance alone to cover the bill. That protects existing capacity and delays replacement.

If the stream remains visible, it can also shorten the expected payback of new hardware. Orders arrive, manufacturers build and new capacity competes for the same reward. The advantage of the new generation is converted into more work competing for each coin.

Fees can therefore purchase persistence or progress. Which one dominates depends on duration, machine cost, used supply and confidence that the activity will remain when hardware arrives.

The factory clock

A fee spike can lower the effective cost of production in the current block while doing almost nothing to the machine frontier. The signal may disappear before a miner orders, a manufacturer builds or a machine is delivered and installed.

A brief event remains mainly a monetary adjustment. A sustained event becomes a capital signal. The machine frontier advances only when the signal survives the delay between revenue and deployment.

The delay is what creates the residue.

The ratchet is not a target

The market does not choose a correct amount of physical work. If fees are low, marginal machines leave. If fees become material, existing coins can fund the same search and, if the flow persists, the next generation of machines.

The source of purchasing power changes. Competition between machines does not. Fees can lower the monetary floor now while raising the physical work behind future coins later.

The first movement belongs to users and miners. The second belongs to factories, delivery schedules and machines that remain after the heat has gone.

CHAPTER SEVEN

Where the Pressure Settles

When a cycle cools, pressure does not disappear. It settles into destinations that leave different residues.

Call them cold sinks. The cold reservoir is where the production gradient ends; a cold sink is where excess pressure goes while the gradient is still operating.

Three sinks matter here: the used-machine market, held coins, and the capacity that fees themselves can fund. One preserves capacity, one removes willing supply, and the third creates new capacity.

The used-machine market

Machines arrive on a schedule that ignores the market's mood. They were ordered when a margin looked large enough to justify a factory decision; by delivery, it may have disappeared.

The used machine carries the previous owner's mistake. Once its discount undercuts new production, it becomes the cheaper route to the same capacity. Capital is destroyed, but the machine changes owner and continues doing work.

The used market preserves capacity. It does not create any. Book One traces what happens when the pool of idle machines is finally absorbed and the next generation becomes worth ordering again.

Held coins

The coins are not destroyed. They become delayed supply. The base can grow while the active flow shrinks, leaving more possible users and connections for the next reason to move.

The sink fills while price sits below what holders will accept, and it can empty faster than it filled. Held coins become willing supply again at a price high enough to end the holding, and what accumulated across a cold phase can return in a single movement.

Over a full cycle, held coins still carry the coin-side ratchet described in Book One.

The capacity fees can fund

Fees are a destination that can transfer urgency directly into new machine production.

A miner spends the same energy to find a block whether the reward contains only issuance or also contains user payments. The reward is larger and the joules are not, so the monetary burden assigned to each coin falls even though the machine's energy cost per unit of work has not changed.

A fee can therefore let a miner survive a lower coin price than issuance alone would support. If the revenue persists, payback shortens, machines are ordered and sustained demand pays for new productive capacity. What that competition leaves behind is this third sink.

Fees can let the cold phase go deeper before marginal machines have to leave, because part of the bill is being carried by a second stream.

Fees do not need a full block before they matter. What decides whether they matter is their size against issuance. Users pay for settlement when the value of moving exceeds the charge and delay. The resulting coins return to the base layer and its miners.

The demand that funds this sink is the demand least likely to persist. Urgency is the shortest-lived reason to move, and it is the one that buys the most durable residue.

What survives

Of the three, only fees can leave behind capacity that did not exist before, and that capacity joins the mining-side ratchet when it survives the cycle.

The same coin can pass through more than one destination over different turns: a holder can fund a machine, a miner can earn fees and buy one, and that machine can later enter the used market.

Nothing here absorbs forever. The used market and held coins run out of what they take in; the fee market stops sorting when the marginal user no longer receives enough benefit to justify moving. What they hold when they stop is what the next cycle starts with.

PART THREE

The Handover

These chapters ask what actually pays for security, how issuance, fees, miners and machines hand pressure from one clock to another, and where the third ratchet reaches a limit of its own.

CHAPTER EIGHT

What Pays the Bill

A continuing flow of purchasing power reaches miners, pays electricity bills and keeps competing machines at work.

Market value is a stock; security expenditure is a flow. They are related, but not interchangeable. The value of all existing coins does not itself pay an electricity bill.

———

Security is a flow

A miner receives a reward for producing a block. It can contain issuance, fees or both. The machine is indifferent to the accounting origin. It responds to revenue, cost and expected margin.

That revenue pays for electricity, hosting, labour, debt and the machine itself. If the margin disappears, the least efficient machines leave. If it persists, more machines arrive until the work is divided and the margin narrows.

Security is the real-world work that this competition can sustain. It has to be paid for again whenever the network is defended.

Present demand reaches the miner through two routes. Demand can bid the coin higher, increasing the purchasing power of each issued coin. Or it can express urgency through fees, transferring existing coins without expanding the stock.

The price route draws purchasing power from outside the existing stock. The fee route reallocates it within the stock. The accounting origin differs; the physical work does not.

———

How a fee stream grows

The fee route has two factors: how many transactions settle and what each one pays. They do not rise for the same reason. A charge rises only when movement is rationed, because excluding marginal movement is what a higher charge does.

Where the settled count cannot grow, a growing fee stream requires a rising charge, and a rising charge is deterrence reaching further into the queue.

Where the count can grow, the stream can grow at an unchanged charge. The same deterrence sorts the same way while more of the demand behind it settles. When demand falls instead, the sorting loosens and the charge falls with it. It does not fall to nothing. Below some level the network will not carry the movement at all, and that level is not set by the contest.

A network whose settled count cannot grow has only the price route left for growth in what it can pay its miners.

The selling ceiling

A miner can sell one hundred percent of newly mined production and no more of it. The ceiling belongs to that source alone. Once every coin received is needed to cover operating costs, a further decline cannot be met by selling a larger share. The next adjustment comes from machines switching off, cheaper power, lower machine prices or another revenue stream.

Fees relax that ceiling by adding a second revenue stream. If fees equal issuance while the joules per block remain unchanged, the miner temporarily receives twice the revenue for the same work. The price required from issuance alone falls, all else equal.

Switching off does not reduce production. Issuance is fixed, so the same coins appear; difficulty falls and survivors receive a larger share of the unchanged total. That larger share is the pawl at work, and it changes the schedule of the next shutdown far more than the amount being sold. The gain is in coins, and the price decline behind it has lowered the purchasing power of each coin.

Output stays constant. The forced flow does not fall with the price. The bill is owed in a unit the miner cannot produce. A lower coin price therefore means more coins must be sold to settle the same amount. Measured in coins, the forced flow rises as price falls, and it keeps rising until enough machines leave at once to break it.

The ceiling also declines as issuance declines. The largest forced sale the network can produce from newly mined coins becomes smaller on the protocol's schedule. Fees can offset that decline only by bringing purchasing power through the second route.

The gap changes hands

A ceiling that falls to zero leaves a question. If the largest forced sale the network can produce from issuance becomes nothing, what is left to answer a market running too hot?

Book One answered it with a different actor. When the profitability gap opened, holders sold coins to buy machines, and those coins went to market as liquid supply. That is abundance, and abundance has one effect on price. The heat produced the thing that ended it.

That mechanism never needed issuance. It needed someone holding a widening gap.

When fees carry the work, that someone is the miner. Fee revenue is largest when settlement is most contested. Where broad market heat is what creates that contest, the gap and the heat arrive together. Under issuance the forced flow was largest when price was lowest. Under fees it is largest when price is highest. The sign has flipped.

A miner holding that revenue has three uses for it and only two are choices: the bill has to be paid, and what is left can be held or spent on more machines. Only holding keeps the coins out of willing supply. Capex reaches the market as surely as opex does, because a manufacturer is not paid in hashrate, though it arrives in a lump and on the factory's schedule rather than the month's.

Nor is the choice free. A fee stream that lifts profitability slightly can be held, but it is also too small to need answering. A fee stream wide enough to justify hardware is visible to every miner at once, and nobody wants to be the one still deliberating. The discretion to hold survives only where there is nothing worth holding back.

So the heat funds the response to itself, and the response arrives as willing supply.

Selling is induced cooling. Nothing decides that the market should cool; the miner paying a bill is the cooling, and it arrives as supply when the heat is highest. The fee coin therefore works against itself, because its revenue is largest exactly when converting it does the most to remove the contest that produced it.

This is a brake and not a ceiling. It does not stop the market at a level. It acts against motion, it strengthens as the market moves faster, and it can be overwhelmed. Miners may hold longer than the mechanism expects. The fee has its own limit, since nobody pays more to settle than settling is worth.

Bounded below by what the work costs. Bounded above by what the heat pays for. Neither is a wall.

The coin and the service

Mining produces two things in one act. The coin is the durable output. Security is the service delivered when the block is produced; it protects the ordering of the ledger now and does not sit in the coin waiting to be used later.

Fees make the relationship visible because current activity pays them directly. The user pays for movement, the miner converts that payment into electricity and competitive work, and the network receives settlement.

If the user does not value moving enough to pay, the transaction does not enter this fee stream. If the fee is paid and the transaction settles, the purchasing power has reached the reward and becomes part of the miner’s revenue. It has not vanished between the two.

CHAPTER NINE

Nobody Schedules the Handover

Issuance and fees run on incompatible schedules. Their handoff unfolds across the clocks of users, miners and machines.

Issuance can fund security before users provide enough fees. That is how a new network buys time to develop a coin market and a settlement market.

The transition is not one moment. It is the point at which internal market heat begins to move the security budget from protocol distribution toward user-paid redistribution.

When issuance carries the block

In a cold part of the cycle, issuance is the clearest source of miner revenue. Fees may exist, but they are too small or irregular to enter a machine decision.

Miners sell enough of the reward to pay the bill. Price and hashrate find the level at which surviving machines continue. Weak hardware leaves, and the remaining work is divided among fewer operators.

Issuance then declines according to schedule while network use can rise or fall independently. The fee share has a direction even when fee activity is uneven. If users do not generate enough demand to pay for the work, the old level of security cannot be maintained indefinitely.

When fees carry part of the block

As activity becomes urgent, fees add revenue to the same work. The miner earns more for finding the block without spending more joules on that block.

At first the effect is local: a few transactions improve a few blocks. If the activity persists, the second stream becomes part of expected return.

This is the handover. The market’s internal heat, not a central decision, determines how much activity is worth settling and how much of that value reaches miners.

The same application can create demand for the coin it uses. The fee stream therefore need not only lower the burden on price; it can also increase the demand that price is measuring.

When the heat leaves

When the activity cools, the fee stream contracts. The monetary burden rises again toward what issuance alone must carry, and the joules supported per coin change with it. The rate at which existing coins pass through the security budget is set by the heat of the economy, not by the schedule.

The protocol does not decide what settlement is worth, so price and fees rise only while users continue to value moving more than waiting.

The corrections can disagree

Fee income changes the miner’s position in two directions. The miner may sell more because the additional return is valuable now. The same income may lower the effective cost of production and attract new machines.

Selling can happen as soon as the miner receives the fee. The hardware response begins with a decision and continues through manufacturing, shipping and installation. The two corrections can therefore point in opposite directions without contradicting each other.

If miners sell the extra revenue before machines arrive, pressure reaches the coin market first. If fees remain high until new capacity is deployed, the physical response arrives later with greater force.

What would settle the question

The unresolved question is not whether both forces exist. They clearly do. The question is which one dominates over a given transition.

The answer should be tested against timing and flows, not argued from a single price chart. Compare fee income with issuance, miner selling with retained rewards, machine orders with delivery dates, used-machine inventory with new-machine prices, and hashrate with the fee stream that preceded it.

The conclusion is conditional. A short fee event may mainly change the monetary burden and the miner’s immediate selling decision. A sustained fee market may also leave capacity that changes the next cold phase. The evidence is in what remains after the urgency has passed.

The next floor is built during the fall

New machines do not vanish when the fee event ends. Some continue operating, some enter the used market and some switch off while waiting for a better price.

A temporary fee stream can therefore lower the monetary floor while it lasts and leave more hardware behind than issuance alone would have justified. The next fall must absorb that capacity through lower machine prices, lower hashrate or a new demand wave.

The price and fee conditions can reset. The machines paid for during the transition remain available to compete in the next turn.

The transition is therefore a question of timing, not whether fees are simply good or bad. The local engine hands temporary revenue to the global ratchet, and the next cycle begins with the capacity that revenue left behind.

CHAPTER TEN

Knowing Is Not Moving

Two of the three ratchets are bounded by what they run out of. The third has been treated as though it only adds.

It does not. The informational ratchet is the most durable of the three and the least able to pay for anything, and those two facts have the same cause.

What it stores

The base keeps two things, and they do not decay at the same rate. One is the route: an interface someone can operate, a habit, a path between two parties that has been walked. The other is the proof that walking it was worth the trouble.

The route is specific and it fades. The proof is general and costs nothing to hold. Nobody maintains the knowledge that a thing can be done; it survives in people who are doing nothing at all.

That is why the ratchet does not reverse. It is also why what remains is the part that earns nothing. A route in use produces fees. A proof in storage produces none.

The ratchet advances by displacing

A route is rarely lost by being forgotten. It is lost by being replaced.

A better hub appears, a cheaper path opens, and the earlier route stops being walked while everyone who used it is still present and still knows what they wanted. The base has not shrunk. The particular thing it was doing has been destroyed by the thing that improved on it.

The replacement does not have to be better. A hub holding something reachable nowhere else can take the traffic while costing more and working less well, because the user is not choosing between two routes but between this one and nothing.

So each advance removes part of what the previous advance left. The level rises and the contents are exchanged. What accumulates is not routes. It is the proof that routes of that kind are worth building.

Two conditions keep a route walked

A route is used while two things hold. Using it has to cost little enough, in charge, in delay and in the effort of operating it. And the result has to be worth enough to someone.

Either one failing is sufficient. A route that becomes expensive is abandoned by users who still want the result. A route that stays cheap is abandoned when the result stops mattering. In both cases the proof remains and the traffic does not.

The fee market sees only the traffic.

A route can outlive its reason

The usual order is that the proof outlasts the route. The reverse happens too.

Infrastructure stays in place, interfaces keep working, habits persist, and the reason that made any of it urgent has moved elsewhere. What is left is a base fully able to move with nothing it needs to move for.

That condition is hard to see from the inside. The participants are there. The connections are there. The hubs are still running. Nothing in the base reports that the urgency has left, because the base was never measuring urgency.

Being able does not fund the work

A fee is what a user pays rather than wait. A base that can act and has no reason to act pays nothing, and its size does not change that.

So the third ratchet reaches its limit the same way as the other two, by running short of what it was converting. The mining-side ratchet runs short of margin. The coin-side ratchet runs short of sellers. The

informational ratchet runs short of urgency, and it can do so while every user, connection and hub it ever accumulated is still in place.

Everything it keeps is a precondition. None of it is a payment.

CLOSING

What Remains

A cycle can return to where it began without returning the system to what it was.

Price can fall back. Fees can disappear. Machines can be switched off. Coins can leave the market and wait in storage.

The surface cools.

Underneath it, three ratchets continue.

The informational ratchet leaves behind people who know how the network works, connections that did not exist before, and hubs that give those connections somewhere to meet. The routes fade when nobody walks them, though what they proved fades more slowly.

The coin-side ratchet leaves more coins held and fewer available for the next movement.

The mining-side ratchet leaves behind something harder to see. It leaves machines, manufacturing knowledge, deployed infrastructure and the memory of what miners were willing to build when the reward justified it. Much of the hardware may become unprofitable. Some of it may change hands. Some of it may never run again.

But the system has learned how to produce it.

That is why a cold phase is not a reset. It is a rearrangement of what the previous heat made possible.

The fee market adds another layer to the rearrangement.

A user pays when waiting costs more than moving. The payment may come from fear of missing an opportunity, fear of not getting out, the need to settle a trade or the simple value of certainty. The network does not need to know which reason is present. It only receives the revealed price of time.

The reasons do not last equally, and the indifference is not costless. What a fee prices is urgency, and urgency leaves first. The least durable demand in the system is what pays for its most durable equipment. The

short-term chase lays down the long-term cost, and that cost is what the coin is later held against.

Either way, the payment reaches miners as coins.

The coins may be newly issued or redistributed from someone who already held them. The machine does not care about the accounting origin. It cares whether the purchasing power arriving at the miner can buy electricity, hardware and continued operation.

The market cares about the origin because the two streams leave different traces.

Issuance expands the stock according to a schedule. Fees move existing coins through the security budget. The monetary floor can therefore move before the physical work changes. A larger fee stream can fund the same level of mining at a lower coin price, all else equal.

But if the fee stream remains, the lower monetary burden becomes a signal.

Miners may keep older machines online. Manufacturers may receive new orders. Capital may move into more efficient hardware. Delivery may arrive after the urgency that created it has cooled. The monetary floor moves on the user’s clock. The mining-side ratchet moves on the factory’s clock.

The mismatch is the engine.

The network does not decide how much redistribution is correct, how much security is sufficient or how much hardware should survive. Users reveal what movement is worth. Miners reveal what work is worth continuing. Manufacturers reveal what future revenue is worth building toward.

The system responds to the flow that reaches it.

That flow can lower the burden on the present, create competition for the future and leave behind conditions that neither side intended. The next cycle begins inside those conditions.

This is bounded self-regulation. Each correction works while it has room to work, then reaches a limit.

The loop returns heat to cold, price to cost, expansion to contraction. The ratchet carries information and physical capacity forward. The states repeat. The system does not.

Satoshi’s Engine is therefore not a machine that produces a final equilibrium. It is a machine that keeps converting pressure into residue.

Some pressure becomes held coins.

Some becomes used machines.

Some becomes new machines and more work.

Some becomes fees, paid by the participants who value resolution more than waiting.

Some becomes knowledge that does not have to be rediscovered.

Think of the engine as a sealed vessel. Take the heat away and everything condenses. The fee falls, the queue clears, the machines go quiet. The vessel returns to the state it started in.

It does not return to the contents it started with. A machine switched off is not a machine undone. It is capacity standing idle, waiting for a reason to run. Each cycle leaves more of it in the vessel than the last one did, and reheating all of it costs more than it cost before.

Every bound on value in this book is a bound on the same thing: what someone is willing to pay.

The fee cannot exceed what settlement is worth to the user. The coin price cannot exceed what a coin is worth to the buyer.

The protocol wrote both scarcities: how much can settle, how many coins arrive, how fast they stop arriving. It wrote neither value.

The coin market runs because supply is finite; the fee market runs because settlement is finite, and it sits inside the first.

A fee is paid in coins and priced by what a coin is worth, so the first engine feeds the second. A coin is worth holding because there is something worth settling, so the second feeds the first.

Each is the other’s fuel, and the protocol owns neither.

Relieving a bound does not relieve what the bound was measuring. Pressure is visible only where something is scarce.

The question is never whether the pressure disappears. The question is where it settles, how long it remains there and what it makes possible when the next reason to move arrives.

The three ratchets meet in the residue. Machines that can still run. Coins that are still held. Users who still know what to do with both. None of the three resets with price.

The local engine cools.

The global ratchet remains.