The Satoshi Overhang

Why the Bear Case is Bounded

Karl T. Ulrich
The Wharton School
University of Pennsylvania
ulrich@wharton.upenn.edu

Working paper, April 15, 2026

Abstract

Renewed public attention on the identity of Bitcoin's pseudonymous creator has sharpened focus on the Satoshi overhang, commonly framed as a tail risk for bitcoin. This paper argues that the mechanical downside of a disposition is bounded well below the existential-loss framing, and that the terminal states most consistent with sixteen years of holder behavior are non-bearish for bitcoin's effective supply. The approximately 1.148 million BTC Patoshi position is analyzed on two tracks. For a purely wealth-maximizing holder, a three-scenario quantitative analysis (Appendix A) shows that bitcoin's current market depth is sufficient to absorb a patient multi-year liquidation at a cumulative price impact in the mid-single-digit to mid-double-digit percent range relative to counterfactual, with the central scenario clustering near 10 percent. The paper maps a decision space rather than identifying a unique modal outcome, assuming a holder whose profile is consistent with the sixteen-year record. Preference sets consistent with the record, including ideological non-intervention, privacy above all, satisficing, and myth preservation, favor continued dormancy terminating in a cryptographically enforced non-recovery or destruction arrangement; preference sets favoring adversarial or wealth-maximizing action are possible but less supported. Across the plausible region of the decision space, the bear case is bounded and the terminal states most consistent with observed behavior are neutral to slightly positive for bitcoin's effective supply.

Keywords: Bitcoin; Satoshi Nakamoto; blockholder; market microstructure; price impact; reflexivity; cryptographic inheritance; effective supply.

JEL classifications: G12, G14, G32, E42, K34, D86.

1. Introduction and puzzle statement

Satoshi Nakamoto mined the first substantial fraction of the bitcoin monetary base during 2009 and the first half of 2010, and then, in April 2011, stopped communicating publicly (Nakamoto, 2008; Bradbury, 2014). The coins attributable to that early mining activity, approximately 1.148 million BTC according to Sergio Demian Lerner's identification of the Patoshi pattern in the extranonce field of the coinbase transactions of the first roughly 36,000 blocks, have never moved (Lerner, 2013; see also Lerner, 2020). At the time of writing, the bitcoin monetary base stands at approximately 20.01 million BTC (Blockchain.com, 2026). Estimates of permanently lost coins range from about 3.7 million, attributed to Chainalysis, to substantially higher figures from other on-chain analytics firms. The Patoshi holdings are therefore on the order of 5.7 percent of nominal circulating supply and approximately 7.0 percent of effective circulating supply once a conservative lost-coins adjustment is made.

The market has long treated this concentration as tail risk. The reasoning is intuitive. If the coins were sold, the volume could swamp normal market liquidity, and the identity of the seller would itself trigger panic. The standard conclusion is that Satoshi's holdings impose a persistent discount on bitcoin's fair value, an overhang that should resolve upward only on confirmation that the coins are permanently dormant and downward on any evidence that they are moving.

The immediate occasion for this paper is a recent wave of investigative journalism, most prominently Carreyrou (2026), naming Adam Back as the leading candidate for Satoshi. The market reflex on such attention has been to read the increased probability of a concrete and identifiable holder as elevated tail risk: if the holder can be named, the holder can be reached, coerced, taxed, or persuaded to sell. This paper argues the reflex overweights the downside. First, the mechanical fear that 1.148 million BTC is too large a position for the bitcoin market to absorb is not supported by the arithmetic of bitcoin's current depth. A patient multi-year liquidation would impose a bounded cumulative price impact, far below the existential loss the tail-risk framing implicitly prices (Appendix A). Second, the behavioral fear that the holder will in fact execute such a liquidation is difficult to reconcile with sixteen years of revealed preference. Across both tracks, the plausible dispositions of the Patoshi position are neutral to slightly positive for bitcoin. The bear case is bounded. The terminal states most consistent with the record are non-bearish.

This paper is a conceptual finance analysis rather than an empirical study. It combines stylized market-microstructure reasoning, a scenario-based quantitative bound on mechanical impact, event-based empirical anchors from prior large bitcoin sales, revealed-preference inference from sixteen years of holder behavior, and comparative institutional analysis of the cryptographic and legal primitives available for inheritance and destruction. It does not attempt to identify Satoshi, to measure market-implied probabilities of disposition, or to recommend trades. Its claim is about the

structure of the disposition problem and the ordering of likely terminal states.

The paper is organized as follows. Section 2 develops the reflexive-liquidation concept and its antecedents in the blockholder and artist-estate literatures, and notes why reflexivity, though conceptually apt, is no longer applicable in the bitcoin case. Section 3 presents Track 1, the disposition problem for a purely wealth-maximizing holder, and derives a quantitative upper bound on the bear case. Section 4 presents Track 2, the disposition problem for a holder whose profile matches the sixteen-year record, and maps the decision space rather than identifying a single modal outcome. Section 5 discusses estate planning and argues that cryptographic inheritance primitives may dominate conventional trust machinery along the cost dimensions most salient to this profile. Section 6 derives the market implications. Section 7 situates the problem within the broader class of reflexive-liquidation positions. Appendix A presents the quantitative scenarios underlying Track 1.

The closest prior observation in the practitioner literature, to our knowledge, is the 2014 *CoinDesk* piece by Danny Bradbury, "How Dangerous is Satoshi Nakamoto?", in which both Gavin Andresen and Sergio Lerner discussed the possibility that Satoshi might burn coins, with Lerner remarking that "if he did burn them, the market reaction would be terribly bullish" (Bradbury, 2014). That observation was casual, did not formalize the holder's optimization, did not bound the bear case quantitatively, did not address the estate-planning question, and did not engage the revealed-preference evidence. The present paper takes up those tasks.

2. The reflexive-liquidation frame

Standard market microstructure models (Kyle, 1985; Almgren and Chriss, 2001) describe the cost of liquidating a large position in terms of two components: a temporary impact arising from the immediate price concession required to find counterparties, and a permanent impact reflecting the information revealed by order flow. A sufficiently patient seller can, under normal market conditions, minimize the temporary component by spreading trades over time and can attenuate the permanent component by trading when information asymmetry is low. The Almgren and Chriss framework gives a closed-form efficient frontier between volatility risk and execution cost. In such models, the liquidation value of a large block is less than its marked-to-market value, but the loss is bounded and increases smoothly with trade size.

The Satoshi problem has conventionally been interpreted as exceeding this standard frame. The additional factor, in the conventional reading, is the price collapse caused by the identity of the seller becoming known. Bitcoin is not valued on cash flows or an intrinsic floor. Its price is a function of beliefs about future scarcity, future adoption, and the intentions of its largest holder. Any movement from a known Patoshi address would be detected within minutes, would be global news within hours, and, in the conventional reading, would discontinuously destroy whatever portion of bitcoin's current price reflects the assumption of permanent Patoshi dormancy.

The word "reflexive" applies to this dynamic in the sense of George Soros (1987, 2013). Reflexivity denotes the two-way causal loop between participants' expectations and the economic fundamentals those expectations bear upon. In Soros's framing, prices are not passive summaries of independent fundamentals; they influence the fundamentals they purport to reflect.

Two analogies from traditional finance carry the concept most directly to the Satoshi case. The first is the controlling founder whose shareholding exceeds the free float. The blockholder literature (Barclay and Holderness, 1989; Holderness, 2003) documents that in such positions the information content of founder trading often dominates its mechanical supply effect. The price response to founder liquidation is not simply proportional to free-float displacement; it reflects the market's inference about the founder's changed view of the firm. Standard mitigation strategies, combinations of charitable transfers and 10b5-1 structured secondary offerings (17 C.F.R. § 240.10b5-1), function as devices for separating the mechanical component of the sale from its informational component, preserving the former while neutralizing the latter. The Patoshi position shares this structural feature: mechanical absorption is manageable under any plausible depth assumption, and the analytical weight of the problem falls on the information channel.

The second analogy is the estate of a deceased artist holding an inventory whose forced simultaneous sale would depress per-work prices, a phenomenon formally recognized by U.S. tax law through the "blockage discount" sustained at 37 percent in *Estate of David Smith v. Commissioner*, 57 T.C. 650 (1972), affirmed 510 F.2d 479 (2d Cir. 1975), and at an effective 37 percent in *Estate of O'Keeffe v. Commissioner*, T.C. Memo 1992-210 (Center for Art Law, 2018). The estate case is closer to the Patoshi case than the controlling-founder case in one important respect: there is no ongoing production by the originator to calibrate the inventory against, and the market's estimate of scarcity must form without reference to the originator's continued activity. The standard mitigation, staged distribution through a foundation or estate over decades, is a direct analogue to Track 1's patient liquidation and illustrates that the reflexive problem is tractable where the holder (or the holder's successor) has the option to spread disposition over time.

These analogies ground the concept. They do not, on their own, establish the overhang as tail risk in the bitcoin case. In Sections 3 and 4 we show that, in the specific case of bitcoin in 2026, market depth has grown sufficient that the mechanical absorption of Satoshi's position is not a binding constraint, and the information channel is substantially weaker than the reflexivity framing assumes.

3. Track 1: The wealth-maximizer's problem

This section analyzes the disposition problem for a holder of the 1.148 million BTC Patoshi position whose utility function is personal wealth alone. The holder is otherwise rational, informed, and strategic, but has no ideological, legacy, or privacy preferences beyond those strictly necessary to

execute. Track 1 exists not to recommend wealth-maximizing liquidation but to bound the downside for bitcoin if the actual holder were in fact such a maximizer.

3.1 Market depth and absorption arithmetic

Satoshi's position is approximately 5.7 percent of nominal circulating supply and roughly 7.0 percent of effective float after the Chainalysis-type adjustment. At current prices on the order of 80,000 USD per BTC, the gross dollar value is about 92 billion. A patient holder executing a decade-long OTC program would therefore sell approximately 115,000 BTC per year, or roughly 315 BTC per day in a market that trades around the clock. Global bitcoin spot and derivatives volume runs in the tens of billions of dollars per day, with the non-wash-traded spot component plausibly 10 to 20 billion. A 25 million dollar daily OTC flow is 0.1 to 0.25 percent of real daily volume, the kind of flow a mid-sized institutional desk runs as background activity.

Applied against published and inferred estimates of bitcoin demand elasticity, which span a range from approximately 0.3 (highly inelastic) to approximately 1.5 (closer to equity-like), a full-supply shift of 7.0 percent of effective float implies a static partial-equilibrium cumulative price impact of approximately 4 to 20 percent relative to counterfactual, depending on the elasticity assumption. Appendix A presents three scenarios, conservative, base, and aggressive, with explicit assumptions on pace, participation rate, elasticity, execution quality, and demand growth. The central scenario clusters near 10 percent cumulative impact; the aggressive scenario reaches approximately 25 percent under the combination of low elasticity and mixed execution quality.

Empirical anchors bracket this range and sort by execution quality. In June and July 2024, the German Federal Criminal Police sold approximately 50,000 BTC in publicly tracked weekly tranches, and bitcoin declined roughly 15 to 20 percent over the episode. In the same window, Mt. Gox began distributing approximately 140,000 BTC to creditors, some of whom sold into the market. Both events were heavily announced, heavily anticipated, and executed in public venues with no effort at concealment; less-committed holders sold ahead of the supply hitting the market, and traders carrying leveraged long positions were forced to close them as price fell. Earlier U.S. Marshals auctions of Silk Road coins, which were less publicized and often sold directly to known institutional buyers, moved price much less, commonly 2 to 5 percent per tranche.

The difference between these two classes of anchor is execution quality. State-actor sales are constrained by procurement and transparency rules that preclude disciplined execution. A rational wealth-maximizing holder faces no such constraint. Contemporary institutional-grade OTC execution uses dispersed desk relationships, time-weighted and volume-weighted algorithmic execution across a deep pool of counterparties, and partial settlement through venues that do not publicly report individual trades. Practitioner accounts describe institutional block execution in

bitcoin as materially lower impact than the public-venue state-actor sales that produced the German anchor, though we are not aware of audited per-block impact figures in the published literature. Applied to a Patoshi selldown program, this suggests the realized mechanical impact is closer to the Silk Road anchor than to the German anchor. The German episode should be read as an upper bound contaminated by sloppy execution, leverage unwinds, and concurrent macro effects, not as the central estimate.

A plausible point estimate for the mechanical cumulative impact of a patient Satoshi OTC selldown over 10 years, under disciplined execution and continued demand growth, is therefore in the mid-single-digit to low-double-digit percent range. Under highly inelastic demand and mixed execution quality, the impact could reach the low twenties. Even that upper bound is far below the fraction-of-pre-event-price framing implicit in the existential-tail reading.

3.2 The information channel

The reflexivity argument of Section 2 supposes that identity revelation, not mechanical volume, drives the catastrophic price response. In 2026, this argument is weaker than it was in earlier bitcoin cycles. Bitcoin's marginal price is now set by institutional flows into ETF products, corporate treasury allocations, sovereign interest, and late-cycle retail participation. None of these channels indexes meaningfully on founder news in its allocation decisions. The crypto-native cohort would treat an authenticated Satoshi action as a major event, but the trading volume of that cohort is no longer large enough on its own to determine bitcoin's price.

Two transient channels deserve explicit treatment.

The first is institutional compliance exposure. ETF sponsors, regulated custody providers, and corporate treasuries with bitcoin balance-sheet allocations operate under compliance frameworks that do not index on founder news for ordinary allocation decisions but do index on regulatory exposure. If a named Satoshi is associated with a jurisdiction under sanction, unresolved tax liability, or active litigation, compliance functions may trigger a pause on new allocations pending legal review. The effect is transient and symmetric: once review concludes, allocations resume. Its magnitude and duration depend on the specific identity and the specific exposure rather than on the fact of revelation.

The second is high-frequency reaction in the perpetual-futures market. The first observable movement from a cold Patoshi address will be detected within seconds by crypto-native high-frequency-trading firms and will trigger short positioning across perpetuals. This produces an initial overshoot on the order of hours to days, plausibly in the 10 to 15 percent range, followed by mean reversion as the news is digested and structural flows reassert themselves. The overshoot is noise around the multi-year mechanical bound, not a structural component of it.

Taken together, a signed Patoshi message announcing sales, or a first movement from a cold Patoshi address, would generate short-term volatility on the order of hours to days, probably including a double-digit percent drawdown followed by partial recovery as the content was digested. It would permanently destroy whatever component of the current bitcoin price genuinely reflects the assumption of permanent dormancy. We do not have a defensible estimate of that component, but sophisticated on-chain research has long applied an effective-supply adjustment that lumps Satoshi's coins together with lost coins, suggesting most of the dormancy premium is already impounded. The durable effect of an aliveness reveal is therefore probably modest, not catastrophic.

3.3 Optimal disposition under pure wealth maximization

Combining the absorption arithmetic of 3.1 and the information-channel analysis of 3.2, a wealth-maximizing holder's optimization reduces to a choice among variants of patient OTC liquidation. The holder might announce the program ex ante, preserving pricing credibility at the cost of the initial informational shock. The holder might execute covertly through a dispersed set of OTC desks, accepting the operational cost and the risk of eventual identification. The holder might constrain the pace of any announced program algorithmically through CLTV timelocks (BIP 65, activated December 2015; Todd, 2014) or CSV timelocks (BIP 112, activated May 2016; BtcDrak, Friedenbach, and Lombrozo, 2015), converting a signed pre-commitment from cheap talk into protocol-enforced constraint. The announced-and-constrained variant is a Rule 10b5-1 analogue in U.S. securities law, stronger than the 10b5-1 case because the signature from the Patoshi keys is unforgeable and the timelock is protocol-level rather than rule-based.

In any of these variants, the wealth-maximizing holder realizes proceeds on the order of 50 to 100 billion USD nominal over the selldown window, against approximately zero if he burned the position outright. For a holder whose sole objective is wealth, burning is strictly worse than patient sale on every horizon. The paper has nothing to recommend to a holder whose sole objective is personal wealth except that his optimal strategy is patient OTC, and the market should not price this outcome as existential tail risk.

4. Track 2: The observable Satoshi profile

This section analyzes the disposition problem for a holder whose identity is unknown but whose profile is partially identifiable from sixteen years of revealed preference. We remain agnostic as to whether the holder is Adam Back (discussed in Section 1 as the leading candidate in recent reporting), another early cypherpunk whose identity has not become the object of media attention, or a figure no outside observer has yet identified. The profile argument does not require us to name the holder. Unlike Track 1, which delivered a quantitative bound, Track 2 delivers a qualitative mapping of the decision space: the set of preference sets consistent with the record, and the terminal

states each would rationalize. We do not identify a single modal disposition.

4.1 Revealed preference: what sixteen years rule in and rule out

Sixteen years of Patoshi dormancy is a behavioral dataset. The strongest inference it supports is a negative one. A holder whose utility function rewarded additional realized wealth above all other considerations would not have been silent across four bitcoin bull cycles (2013, 2017, 2021, 2024), each of which offered attractive exit points. Even an unfavorably executed partial liquidation in any of those cycles would have realized dollar sums large enough to dominate almost any reasonable consumption or bequest function. The holder did none of these. The prior on pure wealth-maximization, consistent with Track 1's hypothetical, is therefore empirically weak.

The record also constrains the technical and operational profile. Satoshi's writing in the original whitepaper, the early BitcoinTalk forum posts, and the source code of the reference client displays skill in applied cryptography, distributed systems, and economic design that narrows the candidate pool to a small subset of the cypherpunk community active in the 2007 to 2010 window. The operational security maintained across pseudonymous communications for two years, the clean withdrawal in April 2011, and the sixteen-year silence since display a level of discipline that further constrains the profile. The ideological positioning of the whitepaper, its references to hard money, supply discipline, and disintermediation from centralized institutions, situates the holder within what is now called the bitcoin maximalist tradition: the position that bitcoin alone, with its fixed supply schedule, is the legitimate digital monetary asset, and that its supply discipline must not be compromised by alternative protocols, contentious forks, or accommodation of centralizing intermediaries. Back, as noted in Section 1, is a publicly discussed instantiation of this profile; the analysis that follows does not depend on the Back identification being correct, nor does it require any particular identity to be correct.

What the record rules in and rules out, on its own, is limited. It rules out wealth-maximizing exit as the dominant preference. It rules in a holder with cryptographic sophistication, operational discipline, and at least some ideological stake in bitcoin's character. It leaves open a range of possible preferences about disposition at horizon.

4.2 Alternative preference sets consistent with dormancy

Dormancy is consistent with many preference sets, not one. Before identifying likely terminal states, we make explicit the space of preferences that would rationalize the observed behavior.

Ideological non-intervention. The holder has strong views about bitcoin's supply discipline and the symbolic importance of the Patoshi position, and treats non-disposition as a constitutive commitment. This is the preference set closest to the maximalist profile, and is the one on which

most of this paper's analysis focuses.

Privacy above all. The holder's primary utility is pseudonymity preservation. Disposition in any form, including structured liquidation, creates detection risk that dominates the wealth gain. Dormancy is the lowest-risk equilibrium.

Satisficing and habit. The holder appears to have reached a consumption utility plateau long ago, through independent wealth, a modest lifestyle, or both, and has no marginal utility of additional wealth. Dormancy is the default in the absence of a positive reason to act.

Key loss or incapacity. The holder has lost access to the Patoshi keys, is cognitively or physically incapacitated, or is deceased without a successor mechanism. Dormancy is mechanical rather than chosen. A variant: Satoshi was more than one person, and group disagreement has precluded any coordinated action.

Myth preservation. The holder recognizes that the Patoshi position's mythological value as a permanent dormant reserve exceeds any realized consumption value and chooses to preserve the myth rather than cash it in.

Legal caution. The holder perceives the realized-wealth or identity-reveal scenarios as exposing him to tax, regulatory, or criminal risk (AML characterization, sanctions exposure, securities characterization) whose expected cost exceeds the realized gain.

These preference sets are not mutually exclusive. Ideological non-intervention and privacy likely co-occur in any profile consistent with the technical record. Key loss and incapacity are observationally indistinguishable from the others in the absence of positive evidence. The group-Satoshi stalemate variant under "key loss or incapacity" deserves explicit note: a Satoshi composed of several individuals whose current preferences diverge and who therefore cannot coordinate any action produces a terminal-state pattern observationally identical to individual ideological non-intervention, since inaction is the default in both cases. The subsequent analysis weights terminal dispositions consistent with ideological non-intervention and privacy preservation, but notes that satisficing, key loss, myth preservation, legal caution, and group stalemate are all consistent with the record and would yield similar near-term observations.

4.3 The adversarial dead-man's switch

One preference set not yet discussed is adversarial. A maximalist holder might interpret subsequent developments in bitcoin, particularly centralization of mining and custody, regulatory capture through KYC frameworks, or dilution of the original supply discipline through forks or protocol changes, as violations of the project's founding commitments. A programmed response, in the form of a dead-man's switch that liquidates or dumps the Patoshi position conditional on specified events, is cryptographically feasible and has been discussed informally in the cypherpunk community.

Three considerations argue against this preference set as dominant. First, the holder has had ample opportunity over sixteen years to signal adversarial intent through cheaper and more surgical tools, including public denunciation, funding of alternative-protocol development, or commissioning of a whitepaper on protocol correctness. None has occurred under any attributable channel. Second, a holder capable of engineering an adversarial switch would plausibly also engineer a neutral or constructive switch (donation to protocol-preservation infrastructure, endowment of a standards body, or destruction). The design space is symmetric; the choice of adversarial design reveals a preference not evident in the record. Third, a dormant adversarial switch is operationally fragile: the holder must maintain accurate views of protocol developments over decades, trigger conditions must be robust to adversarial manipulation, and the mechanism must not be accidentally triggered. The cleaner equivalent, destruction of the keys conditional on the holder's death or incapacity, removes all of these failure modes at the cost of the adversarial capability.

We do not exclude an adversarial switch from the outcome space. We weight it below the non-adversarial terminal states on the strength of the record.

4.4 Terminal dispositions consistent with the record

Three terminal dispositions are most consistent with the preference sets surveyed in 4.2 and 4.3. We list them in rough order of consistency with the record. The ordering is ordinal and is not a formal probability assignment.

Continued dormancy terminating in a cryptographically enforced non-recovery arrangement. Over a remaining life of some decades, the holder maintains silence and arranges for the Patoshi keys to become permanently unrecoverable at or near his death, through one of the cryptographic primitives discussed in Section 5. This disposition requires no announcement, no identity reveal, no pre-commitment apparatus, no legal infrastructure, and no behavioral change from the sixteen-year status quo. It is the terminal state most consistent with the ideological, privacy, and myth-preservation preference sets.

Silent unattributed burn, in whole or in part. At a moment of the holder's choosing in late life, a transaction from the Patoshi addresses to an OP_RETURN output (Bitcoin Wiki, 2024) removes the position, or the bulk of it, from effective float. The transaction is an action rather than an absence and is therefore somewhat more operationally exposed than key destruction, but pseudonymity can hold because interpretation of the event does not require identity. A practically important variant retains a small residual for option value: a transaction that sends approximately 99 percent of the Patoshi cluster (on the order of 1.136 million BTC) to an OP_RETURN output and approximately one percent (on the order of 11,500 BTC or roughly one billion dollars at current prices) to a fresh non-Patoshi address is at the market's scale of resolution indistinguishable from a full burn,

while preserving optionality for the holder. The bullish signal from a 99 percent subtraction is effectively unchanged from a 100 percent subtraction. This disposition is consistent with ideological non-intervention, myth preservation, and, through the retention variant, a residual satisficing or legal-caution preference.

Adversarial switch. Programmed liquidation or dumping conditional on protocol developments, as discussed in 4.3. Weighted lower than the two above on the strength of the record, but not excluded.

The ordering is ordinal and is based on the record. It is sensitive to the discovery of a preference set not captured in 4.2. It is consistent with, rather than disturbed by, evidence that the holder is in fact incapacitated, that the keys are in fact lost, or that Satoshi was in fact a group whose members have divergent current preferences: each of these reinforces continued dormancy as the realized terminal state. We note these sensitivities rather than attempting to price them.

4.5 Why the likely terminal states are non-bearish

The first two terminal dispositions converge to the same supply-side outcome for bitcoin: approximately 1.148 million BTC, or very nearly so in the retention variant, is permanently removed from effective float. They differ in timing, visibility, and the informational content they reveal about the holder's preferences. Their long-run implication for bitcoin's supply path is nearly identical. The market receives, through either channel, a confirmed permanent subtraction of approximately 5.7 percent of nominal circulating supply.

The third disposition, adversarial, would be bearish but is the least consistent with the record for the reasons given in 4.3. Under the observed-preference ordering, the terminal states most consistent with the record are neutral to slightly positive for bitcoin's effective supply, rather than the negative realization implicit in the tail-risk framing. This does not establish a bullish base case in the strong sense. It establishes that the preference sets most consistent with the record deliver terminal outcomes that are not bearish, and that the bearish scenario is the one requiring the strongest departure from the observed profile.

5. Estate planning: cryptographic primitives versus conventional trust machinery

A standard treatment of a very large position would turn at this point to trust structures, tax optimization, and intergenerational transfer mechanisms. For the profile developed in Section 4, that machinery imposes costs the revealed preferences do not bear, and the argument for departure is comparative-cost rather than psychological.

A conventional trust imposes three costs on a holder of this type. First, disclosure: the settlor's identity is disclosed to the trustee, recorded in trust instruments, and subject under the U.S. Cor-

porate Transparency Act and comparable international beneficial-ownership regimes to regulatory reporting. Second, counterparty exposure: the trust carries lifetime exposure to trustee discretion, to litigation by beneficiaries, and to administrative process that can be subpoenaed or disclosed through estate proceedings. Third, process cost: trust administration is slow and expensive relative to the cryptographic primitives the holder has used for sixteen years. For most large holders these costs are tolerable because the settlor's identity is already public and comparable alternative institutions are unavailable. For a holder who has maintained cryptographic self-sovereignty throughout, the costs are a departure from demonstrated preferences, and cryptographic substitutes preserve the pattern.

The direct alternative, a simple handover of private keys to adult heirs, is also a cost departure. Heirs are typically not trained in the operational security the holder has maintained for sixteen years, and once they have the keys, no ex ante mechanism binds them to the holder's preferences about dormancy or disclosure. A holder who has protected this position by trusting no one is unlikely to end the program by trusting his children with bare keys.

The natural substitutes for this profile are cryptographic inheritance primitives that encode the holder's preferences into the mechanism rather than delegating them to a human intermediary, and several are production-grade today. Shamir Secret Sharing (Shamir, 1979) splits a private key into m-of-n shards so that reconstruction requires a quorum; distributions across heirs, trustees, and geographically separated custodians impose delay, consensus, or conditional reconstruction by construction. Multisig wallets encumbered with CLTV (BIP 65) or CSV (BIP 112) timelocks prevent unilateral action until after events the holder has specified, including a fixed calendar date, an elapsed-time condition, or a protocol-observable trigger. Modern threshold-signature schemes such as MuSig2 (Nick, Ruffing, and Seurin, 2021) and FROST (Komlo and Goldberg, 2021) produce on-chain outputs indistinguishable from single-signature transactions and permit m-of-n reconstruction without revealing the threshold to observers. A dead-man's switch that publishes shards or executes a predetermined transaction after a long period without the holder's heartbeat signal is straightforward to engineer on top of these primitives. A deliberate key-destruction arrangement, in which the shards themselves are eliminated at the holder's death, renders the coins unrecoverable by anyone, including the heirs, and converts the bequest into a supply-discipline contribution to the project itself.

These mechanisms are instruments the holder understands natively and trusts by construction, and they preserve pseudonymity across generations in a way a conventional trust cannot. The argument is not that conventional trusts are inferior in the abstract but that for a holder with the revealed preferences of Section 4.1, cryptographic substitutes may dominate along the dimensions most salient to this profile: disclosure, counterparty, and process cost.

A structural observation connects this comparative-cost argument to the economics of trust law. Modern trust and estate practice exists to address the principal-agent problem that arises when a settlor cannot personally execute a multi-decade preference-consistent program after death. The trustee is an agent of the settlor, and the apparatus of trust law (fiduciary duty, accounting requirements, beneficiary standing) aligns the agent's behavior with the settlor's ex ante preferences. The cryptographic primitives described here change the character of that problem: a timelock, a quorum of shardholders bound only to a mechanical reconstruction rule, or a programmed destruction instruction does not involve an agent who must be incentivized or monitored. The settlor's preferences are encoded directly into the mechanism that executes them. The agency relationship is not merely reduced in cost; to the extent the primitive is well-designed, it is absent. This is a stronger substitution for the profile developed in Section 4 than conventional cost-minimization reasoning alone suggests.

One tax-law consequence is worth noting. Under U.S. Internal Revenue Code Section 1014, property held at death receives a stepped-up basis equal to its fair market value at death (26 U.S.C. § 1014; IRS Publication 551). Claiming the step-up, however, requires reporting the coins as part of the estate, which compromises the pseudonymity the profile has protected for sixteen years. A holder whose observable preferences place pseudonymity and supply discipline above tax optimization would decline the step-up, and the bequest to heirs becomes the preservation of the holder's participation in the bitcoin project rather than the coins as a consumable asset.

6. Market implications

The two-track analysis delivers two significant observations. The downside is bounded even under the adversarial Track 1 assumption that the holder is a pure wealth-maximizer. The terminal states most consistent with the observable record are neutral to slightly positive for bitcoin's effective supply.

6.1 What the market may or may not already price

The sophistication of crypto-native market participants is not in doubt, and none of this paper's components is entirely new. Lerner's identification of the Patoshi cluster has been in the public domain since 2013, OP_RETURN has been standardized since 2014, Bradbury's 2014 piece raised burn-as-outcome in informal terms, and the "lost coins" adjustment is routinely applied in on-chain research to arrive at an effective circulating supply. It is plausible that a substantial fraction of the probability mass on permanent-dormancy outcomes is reflected in the bitcoin price through some version of the effective-supply adjustment, and that the information channel we argued in Section 3.2 is weaker than the reflexivity framing implies is partially internalized by institutional flows whose allocation models do not index on founder news. We do not attempt to measure any of this.

Three components of the argument are less obviously present in current practitioner writing. The quantitative upper bound on the mechanical bear case at mid-single-digit to low-double-digit percent cumulative impact over a decade, under disciplined execution and continued demand growth, is not, as far as we can tell, articulated in published or practitioner sources. The comparative-cost argument against conventional trust machinery for this profile has not been developed in the estate-planning or asset-management literature. And the explicit mapping of Track 2's decision space, with its ordering of terminal states consistent with the record and its discussion of alternative preference sets, is to our knowledge not present in prior discussion. Whether any of these are reflected in current prices is a separate empirical question we do not undertake to settle.

6.2 Implications

The paper's claims are structural rather than claims about market mispricing. The bear-case bound of Section 3 rests on bitcoin's current market depth, and the terminal-state mapping of Section 4 rests on the sixteen-year record of holder behavior. Both are observable independent of prices, and the implications that follow are conditional on the structural account rather than on any particular assumption about what the market currently prices (see also Section 6.1).

First, any verifiable event consistent with one of Track 2's first two terminal dispositions, in particular a confirmed death of a profile-matching candidate followed by evidence of non-recovery, or a burn transaction of the position, would plausibly be interpreted as a positive signal consistent with a confirmed supply subtraction. The conditional probability of either event in any given near-term window is low, but the magnitude of the response, conditional on occurrence, would be substantial.

Second, practitioners who hedge "Satoshi risk" through options or structured products should recalibrate. The premium paid for such hedges should reflect the bounded Track 1 upper bound and the non-bearish plurality of Track 2 terminal states, not the existential-tail framing.

Third, pricing models that treat the Patoshi coins as "circulating supply discounted by probability of sale" should be replaced by models that treat them as effectively removed supply with a small residual probability of wealth-maximizing liquidation, itself bounded by the mechanical absorption arithmetic of Section 3.

7. The broader class of reflexive liquidation problems

The Satoshi case is an instance of a class of problems in financial economics: positions whose liquidation value is partially or wholly determined by the act of liquidation itself. The class is not new, but its members have tended to be studied in isolation rather than as a unified category. A preliminary taxonomy follows.

Controlling founders of listed companies. Where a founder holds a stake materially larger than the free float, a sale simultaneously adds supply and removes the signal of founder conviction on which the share price partly rests. Barclay and Holderness (1989) and the subsequent blockholder literature document the empirical price effects. The mitigation strategies observed in practice, combinations of charitable transfers and structured secondary offerings, function as partial analogues to Track 1's patient liquidation.

Estates of deceased artists. The blockage discount in U.S. tax law, sustained at 37 percent in *Estate of David Smith* (1972) and at an effective 37 percent in *Estate of O'Keeffe* (1992), formally recognizes that simultaneous sale of a large inventory of works by one artist depresses per-work prices. The standard mitigation is staged distribution managed by a foundation or estate, essentially a decades-long liquidation schedule.

Insider stakes in privately held companies. Where there is no public market, the attempt by an early holder to liquidate via secondary sales can signal to new investors that the entity is overvalued, depressing both the price of the stake and the valuation of subsequent financing rounds.

The Satoshi case shares the structural features of this class but differs in three respects relevant to its 2026 character. First, bitcoin's market depth, unlike that of a single equity or a private company's secondary market, has grown large enough to absorb the full position through patient liquidation, which bounds Track 1's mechanical bear case in a way the traditional blockholder case is not. Second, the holder's identity remains unknown, so the information asymmetry is not about a known person's changed circumstances but about the prior question of whether the holder is alive, rational, and attentive, a different channel than anything in the blockholder literature. Third, bitcoin offers a uniquely credible destruction mechanism: a transaction to an OP_RETURN output is a voluntary, verifiable, and irreversible supply subtraction, and no other asset in the reflexive-liquidation class offers its holders this option.

The destruction option has a theoretical consequence. For an asset whose value partially rests on expected supply constraints, a large holder's credible destruction option is itself a positive contribution to the asset's price, independent of whether the option is ever exercised. A marginal buyer can reasonably assign some probability to the destruction event and price bitcoin's expected supply accordingly. The asymmetry works in favor of the asset: the destruction option adds supply discipline in expectation, the sale option adds supply in expectation, and the latter carries low probability for the holder who actually exists under the Section 4 profile.

The broader implication is that the study of illiquid and reflexive positions should incorporate voluntary-destruction mechanisms where available, and should treat the absence of such mechanisms in traditional assets as a binding constraint on the holders of those assets. For bitcoin specifically, the paper's claim stands: across both analytical tracks, the Satoshi overhang is bounded on the

downside and non-bearish across the terminal states most consistent with the record. It is not the existential tail risk the conventional framing implies.

Acknowledgements

I benefited from conversations with Shan Wang about the logic in the paper. As a matter of disclosure, my first real education about Bitcoin was in 2017 from a friend who had a single machine in his garage mining coin. I bought 1 BTC at that time just for fun, and I've not sold it. That's my only financial interest in Bitcoin.

References

Almgren, R., and N. Chriss (2001). "Optimal Execution of Portfolio Transactions." *Journal of Risk*, 3(2): 5-39.

Barclay, M. J., and C. G. Holderness (1989). "Private benefits from control of public corporations." *Journal of Financial Economics*, 25(2): 371-395.

Bitcoin Wiki (2024). "OP_RETURN." en.bitcoin.it/wiki/OP_RETURN. Accessed April 2026.

Blockchain.com (2026). "Total Circulating Bitcoin." blockchain.com/charts/total-bitcoins. Accessed April 2026.

Bradbury, D. (2014, November 23). "How Dangerous is Satoshi Nakamoto?" *CoinDesk*. coindesk.com/markets/2014/11/23/how-dangerous-is-satoshi-nakamoto.

BtcDrak, M. Friedenbach, and E. Lombrozo (2015). "BIP 112: CHECKSEQUENCEVERIFY." Bitcoin Improvement Proposals. github.com/bitcoin/bips/blob/master/bip-0112.mediawiki.

Carreyrou, J., with D. Freedman (2026, April 8). "My Quest to Solve Bitcoin's Great Mystery." *The New York Times*. nytimes.com/2026/04/08/business/bitcoin-satoshi-nakamoto-identity-adam-back.html.

Center for Art Law (2018, March 28). "Blockage Discounts and Artists' Estates: The De Kooning Post-Mortem." itsartlaw.org/case-review/blockage-discounts-and-artists-estates-the-de-kooning-post-mortem.

Estate of David Smith v. Commissioner, 57 T.C. 650 (1972), aff'd 510 F.2d 479 (2d Cir. 1975).

Estate of O'Keeffe v. Commissioner, T.C. Memo 1992-210.

Holderness, C. G. (2003). "A Survey of Blockholders and Corporate Control." *FRBNY Economic Policy Review*, 9(1): 51-64.

Internal Revenue Code § 1014, 26 U.S.C. § 1014 (Basis of property acquired from a decedent).

IRS Publication 551 (2025). "Basis of Assets." Internal Revenue Service.

Komlo, C., and I. Goldberg (2021). "FROST: Flexible Round-Optimized Schnorr Threshold Signatures." In *Selected Areas in Cryptography – SAC 2020*, Lecture Notes in Computer Science 12804, pp. 34-65. Springer.

Kyle, A. S. (1985). "Continuous Auctions and Insider Trading." *Econometrica*, 53(6): 1315-1335.

Lerner, S. D. (2013, April 17). "The Well Deserved Fortune of Satoshi Nakamoto, Bitcoin creator, Visionary and Genius." *Bitslog*. bitslog.com/2013/04/17/the-well-deserved-fortune-of-satoshi-nakamoto.

Lerner, S. D. (2020, August 31). "Protection Over Profit: What Early Mining Patterns Suggest About Bitcoin's Inventor." *CoinDesk*. coindesk.com/tech/2020/08/31/protection-over-profit-what-early-mining-patterns-suggest-about-bitcoins-inventor.

Nakamoto, S. (2008). "Bitcoin: A Peer-to-Peer Electronic Cash System." bitcoin.org/bitcoin.pdf.

Nick, J., T. Ruffing, and Y. Seurin (2021). "MuSig2: Simple Two-Round Schnorr Multi-Signatures." In *Advances in Cryptology – CRYPTO 2021*, Lecture Notes in Computer Science 12825, pp. 189-221. Springer.

Rule 10b5-1, 17 C.F.R. § 240.10b5-1 (Trading on the basis of material nonpublic information in insider trading cases). U.S. Securities and Exchange Commission.

Shamir, A. (1979). "How to Share a Secret." *Communications of the ACM*, 22(11): 612-613.

Soros, G. (1987). *The Alchemy of Finance*. Simon and Schuster.

Soros, G. (2013). "Fallibility, reflexivity, and the human uncertainty principle." *Journal of Economic Methodology*, 20(4): 309-329.

Spence, A. M. (1973). "Job Market Signaling." *Quarterly Journal of Economics*, 87(3): 355-374.

Todd, P. (2014). "BIP 65: OP_CHECKLOCKTIMEVERIFY." Bitcoin Improvement Proposals. github.com/bitcoin/bips/blob/master/bip-0065.mediawiki.

Appendix A. Scenario arithmetic for the Track 1 absorption bound

This appendix presents three scenarios bounding the cumulative price impact of a patient OTC liquidation of the 1.148 million BTC Patoshi position relative to counterfactual. The analysis is stylized and partial-equilibrium. It is intended to discipline the qualitative claim of Section 3 that the mechanical bear case is bounded and to make the sensitivity to component assumptions explicit. The scenarios are not point forecasts.

A.1 Common parameters

Position size: 1.148 million BTC, consistent with the Patoshi cluster identified in Lerner (2013, 2020).

Effective float: approximately 16.3 million BTC, constructed as 20.01 million total mined (Blockchain.com, 2026) less approximately 3.7 million lost coins (Chainalysis-type adjustment).

Position as share of effective float at t=0: approximately 7.0 percent.

Reference price: 80,000 USD per BTC. Gross nominal position value: approximately 92 billion USD.

Demand elasticity ε_D is defined such that a 1 percent permanent supply increase, holding demand constant, produces a price change of approximately $(1+0.01)^{-1/\varepsilon_D} - 1$, which for small shifts is approximately $-1/\varepsilon_D$ percent. The published empirical literature on bitcoin demand elasticity is thin, and we do not treat any narrow range as established. The interval from $\varepsilon_D = 0.3$ (highly inelastic) to $\varepsilon_D = 1.5$ (closer to equity-like) is used here as a heuristic sensitivity range intended to span the plausible calibrations that have been proposed in practitioner and academic discussion, not as a confidence interval. Reported sensitivities should be read accordingly.

A.2 Partial-equilibrium impact by elasticity

In a partial-equilibrium model with constant demand elasticity, a full liquidation of the Patoshi position returns 7.0 percent of effective float to the market, producing a price change relative to counterfactual of approximately $(1.07)^{-1/\varepsilon_D} - 1$. The table below reports this impact for three elasticity values spanning the heuristic range described in A.1.

Elasticity ε_D	Cumulative partial-equilibrium impact
1.5 (equity-like)	approximately -4.4 percent
0.7 (central)	approximately -9.2 percent
0.3 (highly inelastic)	approximately -20.2 percent

Elasticity is the dominant source of variation. Halving ε_D from 0.7 to 0.3 approximately doubles the impact; doubling ε_D from 0.7 to 1.5 approximately halves it. Demand growth over the disposition period affects the absolute price path (bitcoin's realized price at horizon is higher under higher demand growth in both the counterfactual and the disposition cases) but under the constant-elasticity log-linear form used here the relative impact between the two cases is invariant to the demand-growth trajectory. The analysis is static partial-equilibrium: it holds elasticity constant and treats market-makers as passive. A dynamic equilibrium in which strategic market-makers

internalize the announced program over a multi-year horizon would likely produce smaller realized impact than reported here, because the anticipated supply would be priced into positioning ahead of the flow rather than absorbed on contact.

A.3 Execution-friction adjustment

The partial-equilibrium impact above is the permanent component attributable to the supply shift. Execution through public markets introduces an additional temporary component whose magnitude depends on participation rate and execution quality. Empirical anchors bracket this component. Practitioner accounts describe disciplined OTC execution as producing small, non-accumulating per-block impact, well below the impact of public-venue state-actor sales. Sloppy or highly public execution contributes more: the German BKA 2024 episode moved approximately 50,000 BTC in weekly public tranches and coincided with a 15 to 20 percent decline, attributable in part to leverage unwinds, macro effects, and execution visibility rather than permanent supply shift.

A.4 Three scenarios

Scenario A (conservative). $\varepsilon_D = 1.5$, disciplined OTC execution, 12-year horizon, participation rate approximately 0.14 percent of real daily spot volume, pace approximately 95,600 BTC per year. Permanent impact approximately -4.4 percent. Execution friction approximately 1 to 2 percent. Total cumulative impact relative to counterfactual: approximately -5 to -6 percent.

Scenario B (base). $\varepsilon_D = 0.7$, disciplined OTC execution, 10-year horizon, participation rate approximately 0.17 percent of real daily spot volume, pace approximately 114,800 BTC per year. Permanent impact approximately -9.2 percent. Execution friction approximately 2 to 3 percent. Total cumulative impact: approximately -11 to -12 percent.

Scenario C (aggressive). $\varepsilon_D = 0.3$, mixed execution quality (partial public venue), 5-year horizon, participation rate approximately 0.34 percent of real daily spot volume, pace approximately 229,600 BTC per year. Permanent impact approximately -20.2 percent. Execution friction approximately 3 to 5 percent. Total cumulative impact: approximately -23 to -25 percent.

A.5 Sensitivity and anchors

The dominant sensitivity is to elasticity. Execution quality shifts the realized impact by a factor of roughly 3 to 5 relative to the permanent-component calculation, consistent with the observed gap between the German BKA anchor (approximately -15 to -20 percent under public-venue execution) and the Silk Road Marshals-auction anchors (approximately -2 to -5 percent per tranche under institutional execution). Scenario C approaches the German anchor under its adverse assumption

stack; Scenario A approaches the Silk Road anchor under its favorable stack; Scenario B lies between them.

A.6 Takeaway

Across plausible calibrations of elasticity, execution quality, selldown pace, and participation rate, the cumulative mechanical price impact of a patient Satoshi OTC liquidation falls in the range from approximately 5 percent (conservative) to approximately 25 percent (aggressive) relative to counterfactual. The central scenario clusters near 10 to 12 percent. No plausible calibration supports the existential-tail framing that implicitly values the position at a fraction of pre-event price. The bear case is bounded by the arithmetic.