

The Satoshi Overhang: Why the Bear Case is Bounded: Open Review

Karl T. Ulrich*

Reviewers: Reviewer A, Reviewer B, Reviewer C

Abstract. The final version of the paper “The Satoshi Overhang: Why the Bear Case is Bounded” can be found in Ledger Vol. 11 (2026) 31-55, DOI 10.5195/LEDGER.2026.614. There were three reviewers involved in the review process, none of whom has requested to waive their anonymity at present, and are thus listed as Reviewers A, B, and C. After initial review, Reviewers A and B asked for revisions and Reviewer C indicated that the journal should accept the submission. The submission was returned to the authors with the feedback of all reviewers (1A). The author resubmitted their work and responded to reviewer comments (1B). Reviewer B subsequently indicated that the journal should accept the paper, and while Reviewer A declined to read the submission a second time due to time constraints, the editor subsequently determined that the revisions and responses were sufficient to meet Reviewer A’s comments. The paper was then accepted for publication, thus ending the peer review process.

1A. First Round of Review

Reviewer A

Does this paper represent a novel contribution to cryptocurrency or blockchain scholarship?

Yes, important contribution(s)

Please briefly explain why you think the paper makes or does not make a novel contribution.

The synthesis of microstructure and behavioral arguments is interesting, but the quantitative result rest on heuristic inputs and the behavioural argument is informed judgment.

Is the research framed within its scholarly context and does the paper cite appropriate prior works?

* Karl T. Ulrich (ulrich@wharton.upenn.edu) is CIBC Endowed Professor at the Wharton School, University of Pennsylvania, Philadelphia, PA, USA.

Important references are missing

Please assess the article's level of academic rigor.

Unsatisfactory (better than poor but a long way from excellent)

Please assess the article's quality of presentation.

Good (not excellent but a long way from poor)

How does the quality of this paper compare to other papers in this field?

The paper has some value but it can easily be replaced by better scholarship in the field.

Please provide your free-form review for the author in this section.

Topic

The paper analyzes the “Satoshi overhang”, the approximately 1.148 million BTC attributed to Bitcoin’s pseudonymous creator via the “Patoshi mining pattern”, which has never moved in sixteen years. Conventional market wisdom treats this concentration as a tail risk: if sold, the volume and the identity of the seller could trigger a catastrophic price collapse. The paper challenges this framing on two tracks. The first derives an upper bound on the mechanical price impact of a patient liquidation, arguing that Bitcoin’s current market depth is sufficient to absorb the position with a cumulative impact in the single-to-double-digit percent range. The second maps the decision space of a holder whose preferences are inferred from sixteen years of observed behavior, arguing that the terminal states most consistent with that record, continued dormancy ending in cryptographic key destruction, or a voluntary burn, are neutral to slightly positive for Bitcoin’s effective supply. The paper concludes that the bear case is bounded and that the existential-tail-risk framing is not supported by either the market arithmetic or the behavioral evidence.

The topic is timely and relevant. The immediate occasion is a prominent April 2026 investigative report naming Adam Back as the leading Satoshi candidate. Whether the Patoshi position poses systemic risk is a live question for asset managers, ETF sponsors, and regulators with significant Bitcoin exposure.

Major Comments

Unanchored quantitative inputs

The central quantitative contribution is the cumulative price impact estimate in Appendix A, which the abstract and conclusion present as a key result. To estimate price impact, the paper applies the standard microeconomic definition for price elasticity E : a supply increase of X percent produces a price change of approximately $-X/E$ percent, which it

then adjusts by an additive execution friction term. Both X and E rest on unanchored inputs.

The supply shift X is constructed as the Patoshi position divided by effective float, where effective float deducts lost coins from nominal circulating supply. The paper uses an estimate of 3.7 million lost BTC (cited only as “Chainalysis-type adjustment” without source) while acknowledging that other analytics firms report substantially higher figures. No justification is given for preferring the low end. A larger lost-coins adjustment reduces effective float, increases X, and mechanically increases the estimated impact.

The elasticity range for E (0.3-1.5) is described as “heuristic” in the appendix. The central estimate of approximately 10 percent is simply the result of choosing a middle elasticity value of 0.7 within the stated range. The paper acknowledges the empirical literature on Bitcoin demand elasticity is ‘thin’ but proceeds to build the entire quantitative architecture on it regardless.

Execution friction estimates are attributed to unnamed “practitioner accounts.”

The 10 percent central estimate is therefore algebra applied to three unanchored inputs and carries no more epistemic weight than the qualitative argument it purports to quantify. The paper would be more honest, and no less persuasive, presented as a structured qualitative argument with illustrative arithmetic rather than as a quantitative bound.

Unsupported market structure assertions

The paper dismisses the reflexivity concern, the argument that identity revelation would trigger a catastrophic price collapse independent of mechanical supply, which rests on two assertions about market structure that are stated without evidence.

(1) Bitcoin’s marginal price is now set by institutional flows into ETF products, corporate treasuries, and sovereign allocations, and that “none of these channels indexes meaningfully on founder news.” No empirical evidence is provided.

(2) The second is the high-frequency perpetual-futures overshoot story: that a first Patoshi movement would trigger a “10 to 15 percent” drawdown over “hours to days” followed by mean reversion as structural flows reassert themselves. This claim is particularly problematic because it converts what could be a permanent information shock into a transient noise event by assumption.

Together these two assertions allow the paper to dismiss the information channel as modest and transient. That conclusion may be correct, but as currently written it is not established.

Minor Comments

Style obscures the strength of claims

Throughout the paper, technical and academic language overstates the evidential basis of the claims and obscures what is actually being argued.

Phrases such as “terminal states most consistent with the observable record” and “reflexive-liquidation frame” present inferences and analogies as established findings. The paper would be more transparent, and more readable, if it used plain language and distinguished consistently between what is derived, what is inferred, and what is the

author's informed judgment. For example: "Section 2 develops the reflexive-liquidation concept and its antecedents in the blockholder and artist-estate literatures" means simply "Section 2 explains why people once feared that Satoshi selling bitcoin could crash the price." The plain version is clearer and more honest about the nature of the claim.

Recommendation

The central thesis is worth serious examination and a path to publication exists. However, the paper conflates informed judgment with derived results in ways that go beyond normal academic hedging. The technical language obscures rather than illuminates the actual strength of the claims, and the reviewer is left with the concern that academic language is being used to present opinions as findings. A major revision should either ground the quantitative claims empirically or reframe as a qualitative paper.

Reviewer B

Does this paper represent a novel contribution to cryptocurrency or blockchain scholarship?

Yes, incremental contribution(s)

Please briefly explain why you think the paper makes or does not make a novel contribution.

The individual ingredients of the thesis are already known, yet their synthesis is novel.

Is the research framed within its scholarly context and does the paper cite appropriate prior works?

Yes

Important references are missing

Please assess the article's level of academic rigor.

Good (not excellent but a long way from poor)

Please assess the article's quality of presentation.

Good (not excellent but a long way from poor)

How does the quality of this paper compare to other papers in this field?

This is a good or average paper.

Please provide your free-form review for the author in this section.

The paper offers a disciplined, well-scoped framework for a question usually argued by assertion, and its core is publishable. I recommend major rather than minor revision because of three issues that should be resolved before acceptance: one of internal consistency that bears on the headline result, one of novelty positioning, and one in the empirical anchoring. None is fatal, and the paper could settle at minor revision if the author rebuts the first point convincingly.

Summary

The paper addresses the "Satoshi overhang" — the ~1.148M BTC Satoshi cluster (~5.7% of nominal supply, ~7% of effective float), commonly framed as existential tail risk and freshly salient after Carreyrou (2026) named Adam Back as a leading candidate. It argues the bear case is bounded on two tracks. *Track 1 (mechanical)* bounds a pure wealth-maximizer's patient multi-year OTC liquidation at a cumulative ~5–25% price impact relative to counterfactual (central ~10%), via partial-equilibrium scenario arithmetic (Appendix A) bracketed by the German BKA (~15–20%) and Silk Road (~2–5%) episodes. *Track 2 (behavioral)* argues that sixteen years of dormancy across four bull cycles makes wealth-maximizing sale a weak prior, and maps a decision space whose most record-consistent terminal states (continued dormancy → cryptographic non-recovery; silent OP_RETURN burn) are neutral-to-positive for effective supply. Sections 5–7 add a comparative-cost case for cryptographic inheritance over conventional trusts, market implications, and a reflexive-liquidation taxonomy. Scope is explicitly conceptual: no identification, no market-implied probabilities, no trade advice.

Significance and novelty

The contribution is real but should be positioned more candidly. The genuinely novel elements are (i) the explicit quantitative upper bound on the mechanical bear case, (ii) the comparative-cost argument that cryptographic inheritance primitives dominate conventional trusts for this profile, and (iii) the structured mapping of Track 2's decision space. Other components — Lerner's cluster, burn-as-outcome (Bradbury 2014), the lost-coins/effective-supply adjustment — are, as §6.1 concedes, already in the practitioner domain. A revision should state plainly what is new versus synthesized and engage the adjacent academic literature (see Prior art), so the marginal contribution is locatable. My assessment is that the novel core clears a publication bar conditional on the revisions below.

Strengths worth preserving. The Track 1 / Track 2 split (how much vs. whether) is the paper's best idea and should remain front and center. The revealed-preference argument — silence across four bull cycles as evidence against dominant wealth-maximization — is economical and persuasive as *ordinal* guidance. Bracketing impact by execution regime

(German BKA vs. Silk Road) is practically useful, the §5 estate-planning argument is original, and §7's credible-destruction-option point is a clean generalization.

Major points (to be addressed for revision)

1. Reconcile the "already impounded" premium — it is currently invoked in opposite directions and never sized. The paper leans on the dormancy premium being *already impounded* to bound the downside of an aliveness reveal (§3.2, §6.1), yet predicts a *substantial* positive response to a confirmed burn or non-recovery (§6.2). These pull against each other: to the extent the permanent subtraction is already priced, confirming it cannot be a substantial *supply* re-rating; to the extent confirmation is substantial, the subtraction was not already priced, and the reveal downside is correspondingly larger than "modest." The same quantity also underwrites Track 1 — the $(1.07)^{-1/\epsilon}$ calculation is only coherent if the coins are currently priced as *out* of float (i.e., impounded). One unmeasured number is thus load-bearing in three places. This need not be a contradiction, but the bookkeeping must be made explicit: decompose the current price, even schematically, into (a) already-impounded permanent-removal and (b) a residual overhang/uncertainty premium, and show that the downside bound, the modest-reveal claim, and the substantial-burn claim are mutually consistent under one split. Because this is a logic issue, the "conceptual/non-empirical" scope does not excuse it.

2. Position novelty against the literature, and add the missing academic neighbors. Beyond the candor noted above, the paper omits its closest empirical and concentration-analysis neighbors (see Prior art). A short "Related literature" paragraph would strengthen, not weaken, the novelty claim by showing none of these makes the paper's integrated argument.

3. The mechanical anchors are confounded, and the central estimate leans on unaudited accounts. The German-BKA vs. Silk-Road gap is attributed chiefly to "execution quality," but the episodes also differ in size, venue, transparency, leverage, and macro regime (the paper concedes the last). The central Track-1 estimate then rests partly on practitioner claims of low OTC impact that §3.1 itself flags as unaudited. The "disciplined execution → Silk-Road-like" inference is plausible but under-evidenced and should be labeled as such, with the elasticity range (point 5) treated as illustrative rather than calibrated.

4. The information/demand channel (§3.2) is the paper's most fragile leg, and is asserted with more confidence than the evidence licenses. Two issues compound, and both point to the same fix.

(a) Demand and effective float are narrative-dependent, not structural. Track 1 holds demand elasticity and "effective float" fixed and adds supply against them. But for an asset the paper itself says has no cash-flow floor, a holder's willingness to sell is a function of expected upside: when conviction is high the offered float is thin, and when the upside narrative weakens that float thickens. A Patoshi disposition that compresses perceived upside would therefore shift demand inward and thicken supply *at the same time*, amplifying impact in precisely the states where a sale is most likely — so the downside is

not independent of the upside. Tellingly, Appendix A notes only the impact-**reducing** dynamic adjustment (market-makers pre-positioning around an announced program) and omits this impact-**amplifying** one. This also sharpens the practitioner point already implicit in the draft: a ~10–25% multi-year path plus a ~10–15% perpetual-futures overshoot is bounded but not small for leveraged funds and treasuries, whose mNAV and covenants are directly narrative-sensitive — a counterexample to §3.2's claim that "none of these channels indexes meaningfully on founder news," and in tension with the paper's own position (§2, §3.2) that bitcoin's reflexivity is now "substantially weaker than the reflexivity framing assumes." A short subsection distinguishing **existential** from **material** drawdown would align the thesis with how institutions actually size hedges.

(b) The §3.2 magnitudes are point-forecasts of a news reaction, which the literature shows are unreliable. The section asserts a "modest" durable effect and a "10–15%" transient overshoot with mean reversion. But event studies find that only the **unanticipated** component of news moves prices (Andersen, Bollerslev, Diebold & Vega, 2003), that responses range from under-reaction and drift (post-earnings-announcement drift; Bernard & Thomas, 1989, 1990) to overreaction and reversal, and that time-sensitive news signals decay quickly after costs. For a singular, never-before-observed event like a Satoshi reveal, the market's prior — and hence the surprise — is unknowable ex ante, so confident magnitude, sign, and persistence forecasts are not supported. This epistemic limit is **symmetric**: it equally forbids the catastrophist's confident "existential" forecast, which is itself a point in the paper's favor. The constructive fix for both (a) and (b) is to widen the error bars on §3.2 and rest the conclusion on the two legs that do not require forecasting a reaction — the mechanical-absorption arithmetic and the revealed-preference base rates — which are the paper's genuinely robust contributions.

Minor points

Elasticity (Appendix A). The $\varepsilon \in [0.3, 1.5]$ range is, as the author says, a heuristic; the ~10–12% central case should be read as illustrative. Stress-testing against alternative "real volume" / liquidity definitions would help.

Track 2 ordinality (§4.4). The ordinal terminal-state ranking is intellectually honest; a natural next step is to translate the preference sets into scenario probability **ranges** for the pricing-model recommendation in §6, even if only as sensitivity bands.

Information channel is testable (§3.2, §6.1). The claim that much of the dormancy premium is already impounded is plausible but unquantified; event studies around prior "Satoshi" rumors or large dormant-wallet movements could sharpen it (and feed point 1).

Track 2 is somewhat over-determined (§4.5). Dormancy, key loss, death, group stalemate, and burn all yield the same non-bearish supply outcome, so "not bearish" is robust but close to structural. One sentence acknowledging this would pre-empt the obvious objection.

Terminology. Given that "Nakamoto" now also denotes a listed treasury vehicle, one clause distinguishing ticker from pseudonym would prevent confusion.

Optional suggestions (offered as directions, not conditions for acceptance)

A reflexive amplification channel via leveraged treasury vehicles. Following directly from major point 4: the sharpest extension of §3.2/§7 is that a Patoshi *move* (not sale) could pressure listed Digital-Asset-Treasury issuers — whose capital structures are narrative- and price-sensitive — into staggered forced selling, an amplification potentially larger than the direct Patoshi flow Track 1 bounds. This is a legitimate generalization of the reflexive-liquidation frame and a strong candidate for follow-on work. I offer it as a direction rather than a requirement, and note fairly that a Patoshi move is only one of several possible triggers of treasury-sector fragility, so it is a scope-expansion rather than a defect in the present, deliberately narrow, analysis.

Use the paper's own trigger as evidence. The Carreyrou reporting that names Back, weeks before this draft, is the most direct available test of the thesis. *How did the market actually react?* If the response was muted, it is the single strongest corroboration of "bounded"; if not, it deserves an answer.

Prior art (suggested citations — not in the current bibliography)

Citation | Why add it

Ante, L., & Fiedler, I. (2021). "Market reaction to large transfers on the Bitcoin blockchain – Do size and motive matter?" *Finance Research Letters*, 39, 101619. | Closest *empirical* neighbor to Track 1: event-study evidence on large on-chain BTC moves and prices (short-horizon; not Patoshi-specific). |

Makarov, I., & Schoar, A. (2021). "Blockchain Analysis of the Bitcoin Market." *NBER Working Paper* 29396 (also LSE FMG Discussion Paper). | Maps ownership concentration and intermediaries but *excludes* Satoshi/abandoned coins — leaves founder disposition open. *(Working paper; not published in the* *American Economic Review*.*)*

Ambrosia, M., Dorrell, J., & Stockwell, T. (2024). "Is active bitcoin supply decreasing? An empirical analysis." *Journal of Economics and Finance*, 48(4), 1166–1186. | Peer-reviewed *active/dormant supply* estimates; complements the paper's effective-float adjustment.

None makes the author's integrated argument; they position where this paper sits relative to adjacent work.

Methodological references cited in major point 4(b):

Andersen, T., Bollerslev, T., Diebold, F. X., & Vega, C. (2003). "Micro Effects of Macro Announcements: Real-Time Price Discovery in Foreign Exchange." *American Economic Review*, 93(1), 38–62.

Bernard, V., & Thomas, J. (1989). "Post-Earnings-Announcement Drift: Delayed Price Response or Risk Premium?" *Journal of Accounting Research*, 27(Suppl.), 1–36.

Bernard, V., & Thomas, J. (1990). "Evidence That Stock Prices Do Not Fully Reflect the Implications of Current Earnings for Future Earnings." *Journal of Accounting and Economics*, 13(4), 305–340.

Reviewer C

Does this paper represent a novel contribution to cryptocurrency or blockchain scholarship?

Yes, incremental contribution(s)

Please briefly explain why you think the paper makes or does not make a novel contribution.

It is not a major contribution but rather an opinion paper. The discussion is interesting though.

Is the research framed within its scholarly context and does the paper cite appropriate prior works?

Yes

Please assess the article's level of academic rigor.

Good (not excellent but a long way from poor)

Please assess the article's quality of presentation.

Good (not excellent but a long way from poor)

How does the quality of this paper compare to other papers in this field?

The paper has some value but it can easily be replaced by better scholarship in the field.

Please provide your free-form review for the author in this section.

The manuscript “The Satoshi Overhang: Why the Bear Case is Bounded” presents an interesting and timely contribution to the discussion surrounding long-term Bitcoin supply

overhang and its potential market impact. The author tackles a widely debated concern—namely, the latent supply associated with early-mined coins—and argue that its downside implications are structurally limited. The paper succeeds in framing the discussion in a clear and accessible way.

The core argument is well articulated, and the paper provides a coherent narrative linking technical Bitcoin properties, behavioral assumptions, and market dynamics. The bounded bear case thesis is particularly compelling in that it reframes what is often treated as an unquantifiable tail risk into a more structured and analyzable component of Bitcoin’s supply profile. This is a valuable perspective for both researchers and investors.

Most of my detailed comments, including technical clarifications, suggestions for tightening certain arguments, and minor corrections, are provided directly in the annotated PDF of the manuscript.

At a higher level, my main recommendation is to further develop the section discussing transactions to OP_RETURN outputs. This part of the paper raises an important mechanism—not often very well known—that could influence the effective circulating supply. Expanding this section—both in terms of empirical evidence and conceptual framing—would strengthen the overall argument and improve the paper’s contribution.

Overall, the paper is well written and offers a perspective that is original. With some additional development in the areas noted above, it would make an interesting contribution to the literature.

1B. Author’s Responses

Summary and approach

We thank the editor and the three reviewers for unusually careful and constructive reports. We treated the reviews less as objections to rebut than as a guide to making the paper better. The most consequential changes are:

1. A consistency ledger (new Section 3.4, formalized in Appendix A.7) resolves Reviewer B’s central point. The downside of an unanticipated sale and the upside of a confirmed burn are complementary shares of a single bounded total that the absorption arithmetic fixes, so the paper’s separate claims are mutually consistent at every value of the one unobserved quantity, namely the share of the position the market already prices as permanently gone. We never need to estimate that share.
2. A claim-discipline pass distinguishes throughout the mechanical, supply-driven bear case, which we bound, from the total market reaction, which we do not forecast. Section 1 now states an explicit taxonomy of derived, inferred, and judgment claims and uses it

consistently. This addresses the concern, shared by Reviewers A and C, that academic language was presenting judgment as finding.

3. A rebuilt information-channel section (Section 3.2) withdraws all point forecasts, rests on the event-study literature (only the unanticipated component of news moves prices, and reactions can drift or reverse), and adds two natural experiments now on the record: the April 2026 identity reporting and the July 2025 Satoshi-era sale, both of which produced small, transient responses.

4. Empirical anchoring and sourcing of the Appendix A inputs. The lost-coins figure is sourced and stress-tested, the elasticity range is treated as an illustrative sensitivity range triangulated with an independent square-root-law cross-check (new Appendix A.8) and the episode anchors, and the unaudited practitioner claims have been removed.

5. Related-literature positioning (Section 1) and a candid new-versus-synthesized statement (Section 6.1) engage the closest academic neighbors. Every reviewer-suggested citation was verified against primary sources before inclusion, and DOIs were added where available.

One cross-cutting decision. Reviewer A recommended reframing the paper as qualitative; Reviewer B recommended keeping the quantitative bound but labeling it illustrative. We followed a path consistent with both. We retain the bound as a structured upper bound under explicitly stated, uncalibrated inputs, triangulated by three methods that share no parameters, and we label its epistemic status precisely. This preserves what Reviewer B identified as the paper’s most distinctive contribution while meeting Reviewer A’s demand for honesty about what the arithmetic can and cannot support.

Two errors we caught and corrected. In the interest of credibility we disclose two corrections made during revision. First, the earlier Appendix A double-counted, including the Patoshi coins in the float while also adding them to it as supply; the consistent supply-shift factor is 1.0757, and the central permanent impact is therefore about negative 9.9 percent rather than negative 9.2 percent (conclusions unchanged). Second, the 2024 German sale was misattributed; the seller was the state of Saxony, not the BKA, and the manuscript is corrected accordingly.

Throughout, we use the labels from the reviewers’ reports. Each concern is restated concisely before our response, and section numbers refer to the revised manuscript.

Reviewer A

A.1. Unanchored quantitative inputs.

Concern: The Appendix A impact estimate applies algebra to three unanchored inputs (the lost-coins adjustment, the elasticity range, and an execution-friction term attributed to unnamed practitioners); present it as a structured qualitative argument with

illustrative arithmetic. Response: We addressed all three inputs and reframed the result. The lost-coins figure is now sourced (Chainalysis, via Roberts and Rapp 2017), with the choice of the high end of the 2.78 to 3.79 million range justified because losses accumulate; Appendix A.6 additionally reports the bound under an adverse dormancy-based envelope (Glassnode’s “Hodled or Lost Coins” metric, about 7.8 million). The point that choosing the low end was anti-conservative is conceded and absorbed: even the adverse envelope, the most inelastic demand, and adverse execution together bound the cumulative impact near negative 30 percent. The elasticity range is now explicitly a heuristic sensitivity range, not a calibration, and the central figure is triangulated with an independent square-root-law calculation (new Appendix A.8) and the episode anchors rather than resting on the midpoint. The unaudited execution-friction claims are removed; the temporary-impact allowance is anchored to observed episodes and to the peer-reviewed decay result of Donier and Bonart (2015). Per the remedy suggested, the bound is presented as a structured argument with illustrative arithmetic and an explicit epistemic label (Section 1; Appendix A introduction), not as a calibrated estimate.

A.2. Conflation of informed judgment with derived results.

Concern: The paper uses academic language to present opinion as finding, beyond normal hedging.

Response: Section 1 now states an explicit taxonomy and applies it throughout. Derived claims (the absorption bound and the ledger identity) are certain given stated assumptions; inferred claims (the weak prior on pure wealth-maximization) are drawn from the record with stated uncertainty; judgment claims (the Section 4 ordering of terminal states) are offered as a reasoned ranking, not a measurement. The Section 3.2 rebuild removes the largest instance of judgment presented as finding.

A.3. Missing references.

Concern: Important references are missing (specifics cross-referenced to Reviewer B).

Response: A related-literature paragraph (Section 1) now engages the closest empirical and concentration-analysis neighbors. See B.2; all citations verified.

A.4. Unsupported market-structure assertions.

Concern: Two claims are stated without evidence: that marginal price is set by institutional flows that do not index on founder news, and that a Patoshi move would produce a 10 to 15 percent transient overshoot with mean reversion, which converts a possible permanent shock into a transient one by assumption.

Response: Both assertions are withdrawn. Section 3.2 no longer claims that no channel indexes on founder news; it treats the institutional-compliance channel explicitly and concedes the leveraged and treasury-vehicle amplification channel with current evidence (Section 3.2, Section 7). It makes no point forecast of an overshoot. The conclusion now

rests only on the mechanical arithmetic and the revealed-preference base rates, the two legs that require no forecast of a reaction.

A.5. Style obscures the strength of claims.

Concern: Technical phrasing presents inferences and analogies as findings; use plain language and consistently separate derived, inferred, and judgment claims.

Response: Done in two parts. The categorization is now explicit (A.2), and a full plain-language pass shortened sentences and removed jargon. For example, the roadmap now says Section 2 “explains why people came to fear that a Satoshi sale could crash the price.”

Reviewer B

B.1. Reconcile the already-impounded premium (load-bearing).

Concern: The dormancy premium is invoked to bound the downside of a reveal, yet a confirmed burn is predicted to be substantially positive, and the Track 1 calculation is coherent only if the coins are priced as out of float; decompose the price and show the three claims are mutually consistent under one split.

Response: This is the revision’s lead change (new Section 3.4, formalized in Appendix A.7). We define the single unobserved quantity, decompose the current price into an already-impounded permanent-removal component and a residual, and show that the reveal-downside bound, the modest-reveal claim, and the burn-upside claim hold for every value of it. The burn upside and the unanticipated-sale downside are complementary shares of one total fixed by the absorption arithmetic, so they cannot both be large; the catastrophist and the burn-bull are competing for shares of the same bounded total. We never estimate the share. Building the ledger is what surfaced the double-counting error disclosed above, which we then corrected.

B.2. Novelty positioning and missing literature.

Concern: State plainly what is new versus synthesized, and add a related-literature paragraph engaging the closest academic neighbors (verify the suggested citations).

Response: Section 6.1 now states what is new (the quantitative upper bound, the comparative-cost inheritance argument, and the Track 2 decision-space map) and what is synthesized from the practitioner domain (Lerner’s cluster, burn-as-outcome, and the effective- supply adjustment). Section 1 adds a related-literature paragraph engaging Ante and Fiedler (2021), Makarov and Schoar (2021), and Ambrosia, Dorrell and Stockwell (2024), showing that none makes the paper’s integrated argument. All three were verified against primary sources; Makarov and Schoar is cited as an NBER working paper, per the reviewer’s own caution against attributing it to a journal.

B.3. Confounded mechanical anchors and an unaudited central estimate.

Concern: The German versus Silk Road gap is attributed chiefly to execution quality, though the episodes also differ in size, venue, transparency, leverage, and macro regime, and the central estimate leans on unaudited OTC claims; label the inference under-evidenced and treat the elasticity range as illustrative.

Response: Section 3.1 now reads the anchors correctly, as bounds on the temporary component under different execution regimes, with the permanent component addressed by the arithmetic, and it acknowledges the confounds explicitly. The disciplined-execution inference is labeled as such, and the elasticity range is illustrative. The unaudited practitioner claims are removed and replaced by the observed, sourced July 2025 OTC sale and the decay result of Donier and Bonart (2015). We also note (Appendix A.3) that the Saxony price effect substantially reverted, which is evidence of a large temporary and near-zero permanent component.

B.4. Information and demand channel overstated.

Concern: (a) Demand and effective float are narrative-dependent, so a disposition that compresses upside amplifies impact in exactly the states where a sale is most likely, and the draft included only the impact-reducing dynamic; a material-versus-existential distinction is needed. (b) The Section 3.2 magnitudes are point forecasts of a news reaction, which the literature shows are unreliable for a singular event.

Response: Section 3.2 is rebuilt. For (a), the impact-amplifying channel is added symmetrically (Appendix A.2 now presents the amplifying and dampening adjustments side by side), and the existential-versus-material distinction is added and carried into Section 6.2, including the leveraged and treasury-vehicle case. For (b), all point forecasts are withdrawn; the section rests on the event-study limit (Andersen, Bollerslev, Diebold and Vega 2003; Bernard and Thomas 1989, 1990), which is symmetric and equally forbids the catastrophist's confident forecast. The conclusion now rests on the mechanical arithmetic and the revealed-preference base rates.

B.5. Treat the elasticity central case as illustrative; stress-test.

Response: Done. The central case is labeled illustrative (Appendix A.1), stress-tested against alternative lost-coins and volume assumptions (Appendix A.6), and cross-checked by the square-root law (Appendix A.8).

B.6. Translate the Track 2 ordering into scenario probability ranges.

Response: Done. (Section 4.4), as explicitly reader-side sensitivity cases rather than the paper's own probabilities, and shown to leave the Section 6.2 recommendation unchanged because even the bearish branch is bounded by the absorption arithmetic.

B.7. The already-impounded claim is testable.

Response: Done (Section 6.1). We note the event-study design around large dormant-wallet movements and credible Satoshi rumors, with the July 2025 sale and the April 2026 reporting as the first data points, and Ante and Fiedler (2021) as a template, while declining to run the estimation here.

B.8. Track 2 is somewhat over-determined.

Response: Done (Section 4.5). One sentence now acknowledges the over-determination and pre-empts the by-construction objection, noting that the adversarial outcome is a genuine bearish state the framework admits, which the record, not the accounting, weights down.

B.9. Distinguish the “Nakamoto” ticker from the pseudonym.

Response: Done at first mention (Section 1), distinguishing the pseudonym and the Patoshi cluster from the listed treasury vehicle Nakamoto Inc. (Nasdaq: NAKA).

B.10. Reflexive amplification via leveraged treasury vehicles (offered as a direction).

Response: We incorporated this rather than deferring it. Section 7’s taxonomy now includes leveraged digital-asset-treasury vehicles as a reflexive-liquidation class in their own right, and Sections 3.2 and 6.2 identify them as the principal pathway by which a material drawdown could propagate, while noting a Patoshi move is only one of several possible triggers.

B.11. Use the Carreyrou and Back episode as evidence.

Response: Done (Section 3.2). Bitcoin was essentially flat on April 8, 2026, the day of publication, with steady open interest and no liquidation cascade. As the reviewer anticipated, a muted reaction is direct corroboration of the bounded thesis, and we present it as such while noting its limits as a single, unconfirmed-identification event.

Reviewer C

C.1. Expand the OP_RETURN section.

Concern: Develop the OP_RETURN burn mechanism further, in both empirical evidence and conceptual framing, since it bears on effective circulating supply and is not widely understood.

Response: Appendix B now details the mechanics, distinguishes three grades of destruction by verifiability (provable burn, burn address, key destruction), gives the largest historical burn for scale (the Counterparty proof-of-burn, slightly over 2,100 BTC), and addresses the 2025 relay-policy controversy, noting that it concerns the relay of data, not the consensus-level unspendability of value sent to an OP_RETURN output.

C.2. Opinion-paper framing.

Concern: The paper reads as an opinion piece; tightening the arguments should address this.

Response: The epistemic-status taxonomy (Section 1) and the claim-discipline pass convert assertions into labeled argument, and the substantive fixes (B.1, A.4 and B.4, and the sourcing in A.1) replace opinion with derivation and evidence where it was thin.

C.3. Is the dormancy discount real? Compare Elon Musk and Tesla. Concern: The claim that Satoshi’s holdings impose a persistent discount is arguable; how is it different from Musk and Tesla?

Response: Section 1 now engages this directly: markets apply no consensus discount to founder stakes of similar proportional size, with Tesla under Musk as the salient example, and large disclosed founder sales tend to produce bounded, partly transient drawdowns. Section 3.4 then gives the word “discount” precise content as the product of a residual sale probability and a bounded spread.

C.4. Almgren and Chriss is a simplification. Concern: Impact depends heavily on execution strategy and order-book shape; the Almgren and Chriss framework simplifies practice.

Response: Section 2 now uses the framework only for its conceptual decomposition of temporary and permanent impact and its volatility-cost frontier, states that none of the quantitative claims depends on its functional forms, and notes that empirical bitcoin impact follows the square-root law (Donier and Bonart 2015), which the cross-check in Appendix A.8 uses.

C.5. The liquidation menu omits the derivatives-first strategy.

Concern: A wealth-maximizer could short the perpetual market, then move and sell the coins, closing both legs together; the Section 3.3 optimization should account for this.

Response: Section 3.3 now treats this variant. Shorting first shifts the detection-triggered repricing onto counterparties, but the strategy is bounded by detectability (large perpetual shorts are identified within hours) and by market scale (a full hedge would require a short position several times the size of the entire bitcoin futures market). It strengthens the paper’s conclusion, since it moves the wealth-maximizer’s realized outcome closer to the mechanical bound; we take no view on its standing under derivatives anti-manipulation law.

C.6. “Satoshi risk” is one source of volatility, not a separately hedged risk.

Concern: Practitioners trade bitcoin volatility; Satoshi risk is one contributor among others.

Response: Section 6.2 adopts exactly this framing. Practitioners do not hedge a named “Satoshi risk”; they trade bitcoin volatility, to which a Satoshi event is one contributor among macro shocks, regulatory actions, exchange failures, and leverage unwinds, and it should be sized as a material-drawdown scenario rather than a separately hedged existential state.

Closing

The revision adds two natural experiments, an independent cross-check of the bound, a consistency ledger, a related-literature paragraph, and an explicit epistemic framework, while removing the unaudited inputs and the point forecasts that the reviewers rightly flagged. Every citation, including the reviewer-suggested additions and the methodology references, has been verified against primary sources, with DOIs added where available. We believe the paper is substantially stronger for the reviewers’ input, and we are grateful for it.



Ledger is published by Pitt Open Library Publishing, an imprint of the University Library System, University of Pittsburgh. Articles in the journal are licensed under a Creative Commons Attribution 4.0 License.