Bitcoin Core: The Biggest Fallacies

Josh · Secure Sovereign

Published March 21, 2026 · Updated September 20, 2026

Inspired by Jude Nelson's 2014 analysis of systemd

Contents

There is a set of arguments that appears repeatedly in defense of Bitcoin Core's implementation monopoly. This post is meant to serve as a repository of those arguments and to explain why they are logically invalid. I will direct people here when they appear.

Core's engineering record is not the dispute. Review culture, release engineering, and reluctance to rush consensus changes are why the money rules have held. The question of whether Bitcoin Core is the right client for a given operator is separate from whether these eight defenses of the implementation monopoly survive scrutiny. The governance critique here is also separable from any opinion about Bitcoin's consensus rules. **This is not an argument that Bitcoin is broken.** For the numbered argument map, see *Bitcoin Governance: Argument Map*. For narrative evidence and funding maps, see *Who Controls Bitcoin*.

Fallacy 1: "Bitcoin Core has hundreds of contributors, therefore it is not controlled by a small group"

This is a non sequitur. **Contributor count and merge authority are independent measurements**, and conflating them is either naive or deliberate misdirection. Hundreds of people have opened pull requests. A handful of accounts land them.

Analysis of 25,122 pull requests, 430,613 IRC messages, and 17 years of data tells the merge-authority story. The top three merger accounts control 81.1% of all historical merges, with Wladimir J. van der Laan at 34.8%, Michael Ford at 25.8%, and Marco Falke at 20.5%. In the 2022+ window the top merger takes 50.2% and the top three take 82.2%. The authorship Gini is 0.851 historical / 0.834 recent. Power is not distributing over time. Top-10 authorship increased from 42.7% to 47.8%. Of 7,827 contributors across the project's history, 90.7% have exited with no activity in the past year of the dump. The contributor count is not the governance count, and one does not imply the other.

Fallacy 1.1: "The process has improved, therefore the structural problem is being addressed"

Process improvements are real. Zero-review merges dropped from 30.3% to 3.3%. First-review clocks are nearly equal. These are genuine workflow gains. The self-merge rate, however, is **25.5%** of 9,793 maintainer-authored merges (95% CI 24.7–26.4). The review weight bias between maintainers and non-maintainers has held at 5.5 to 1. Top-10 authorship increased. The authorship Gini is essentially unchanged. Non-maintainers wait 1.41× longer to merge. The oligarchy got more efficient at processing pull requests and the concentration of authority did not move. **Those are two different things.**

Fallacy 1.2: "Maintainers are just janitors, not decision-makers"

This framing appeared in mainstream Bitcoin commentary as early as 2018 and has been repeated since. The data does not support it. Maintainers self-merge **25.5%** of their own pull requests, and **45.1%** of those self-merges carry zero reviews (**11.6%** of all maintainer-merged PRs). One historical maintainer self-merged 77.1% of his own pull requests. Merge authority over the lineage that almost all economic nodes run **is not a janitorial function** by any honest definition of the term.

Fallacy 2: "Widespread adoption proves technical merit"

This is a self-serving bias. It assumes that the mechanism of adoption was active preference rather than the absence of alternatives, and that assumption does not hold up.

Other clients exist. A Core-derived fork that stays on Bitcoin is still a Core fork. Btcd abandoned active development and is not positioned to confront Bitcoin Core. Libbitcoin abandoned the UTXO set and mempool, requiring ecosystem retooling and accepting isolation. Running Bitcoin Core is not a positive selection on merit but the predictable outcome of path dependency, exchange requirements, and wallet vendor defaults. Adoption through switching cost is not adoption through technical merit, and treating it as such is begging the question.

Fallacy 2.1: "Alternative implementations exist, therefore there is no monopoly"

The existence of alternatives on paper is not the same as viable alternatives in practice. Forking Bitcoin Core inherits 300,000+ lines of monolithic C++, 17 years of technical debt, and the same capturable-architecture vectors under new management. Without a formal specification floating above the implementation, any alternative must reverse-engineer undocumented behavior from the reference client itself, which makes the reference client the de facto specification by default. The lack of a specification protects the incumbent by making independent implementation structurally more expensive than it would otherwise need to be. To date, no alternative implementation has successfully replicated Bitcoin Core's consensus implementation entirely unless it is directly forked from Bitcoin Core.

Fallacy 3: "Critics of Bitcoin Core just don't like change"

This is an ad hominem that routes around the substantive argument by attacking the character or motivation of the critic rather than the content of the criticism.

Jon Atack has been a Bitcoin Core contributor since 2019, a BIP editor, and a developer. In March 2026, he stated on record that maintainer selection involved favoritism and double standards, that preferred successors were picked and groomed explicitly, and that the record would speak for itself when laid out. He stated directly that the ship has sailed for someone like him to make change from within Core because the entrenched are entrenched. These are the words of a longtime insider with no competing implementation to promote, not the words of someone who dislikes change.

Beyond Atack, 31 Bitcoin Core developers signed an open letter against the OP_RETURN change and were overridden. Developers were banned from GitHub for raising a documented conflict of interest involving a contributor with a direct financial stake in the outcome. Luke Dashjr was removed from the Bitcoin Core security list after a decade of contribution with no documented process and no appeal. The pattern is insiders with documented grievances, not outsiders resistant to progress.

Fallacy 4: "Bitcoin Core is not monolithic"

Bitcoin Core is approximately 300,000 lines of C++ in which the wallet, networking layer, consensus engine, and RPC interface are entangled rather than independently composable. The fact that it has internal modules does not make it non-monolithic any more than the Linux kernel's loadable modules make it non-monolithic. Modularity and monolithism are independent properties and a codebase can be both simultaneously.

The deeper issue is that no formal specification exists above what Bitcoin Core ships. When almost all nodes run that lineage, the reference client and the protocol become functionally identical regardless of what anyone claims about their separability. The absence of a specification is not a neutral technical choice. It protects the incumbent by ensuring that any alternative implementation must reverse-engineer undocumented behavior rather than implement against an independent mathematical standard.

Fallacy 5: "Rough consensus governs Bitcoin Core decisions"

Rough consensus is defined and invoked by the same maintainers who close pull requests, which means there is no external arbiter, no defined threshold, and no appeal mechanism. Bitcoin Core's CONTRIBUTING.md states that the final arbiter of what constitutes sufficient consensus is the judgement of the maintainers. That is not a governance rule. That is unlimited discretion with a rule-shaped wrapper around it.

The stronger evidence is who can close the file. Pair-level reviewer cohesion exists, but the joint-appearance counts are too small to treat as a voting machine. What is not small is merge share: one account lands about half of 2022+ merges, three land 82.2%, and the implementation PR for package relay was titled `[NO MERGE]` until maintainers declared rough consensus on approach. If the people who define sufficient consensus are the same people who merge, invoking rough consensus as a legitimizing mechanism is circular. The outcome is not emerging from an external process. The process is ratifying a judgment the merge key already holds.

Fallacy 6: "You should judge Bitcoin Core on technical merits alone"

Judging software on technical merits alone requires a formal specification that defines precisely what the software does and does not do, and a formal proof that the implementation satisfies that specification. Bitcoin Core has neither. In the absence of both, developer trustworthiness, funding relationships, governance behavior, and conflicts of interest are not irrelevant considerations to be excluded from the analysis.

Their track record is not reassuring. A Bitcoin Core maintainer had their security compromised while holding merge access to infrastructure securing over $1 trillion in network value, and the community response was effectively silence. A wallet deletion bug shipped to production. The OP_RETURN change was merged over documented opposition that included a signed letter from 31 contributors, with at least one developer banned from the GitHub repository for pointing out that a contributor pushing the change was affiliated with a for-profit project that directly benefited from it. These are not external attacks on the project. They are governance failures that originated

inside the structure itself. For the same reputation-over-verification pattern applied to hardware wallets, see *Don't Trust, Verify*. For why irreversible stakes make a softer culture the wrong prescription while those failures continue, see *The Adversarial Default*.

Fallacy 7: "Bitcoin Core is conservative by design, and that's a feature"

The conservation of consensus rules is a genuine and important property of Bitcoin. What gets conflated with it, consistently and conveniently, is the conservation of the implementation itself. These are two distinct things and treating them as the same thing is either confused or dishonest.

Wallet and node separation has had universal agreement among contributors for approximately 12 years and has not shipped. UTXO commitments have been research-complete since 2014, would reduce the initial block download burden for new nodes by approximately 98%, and have not shipped. Every year that UTXO commitments do not ship, every new node operator pays the full cost of downloading 17+ years of chain history including all spam and bloat that UTXO commitments would have made irrelevant. That cost compounds with every passing year and every new participant: quantified in *Full Cost of Running a Bitcoin Node*. The governance structure cannot distinguish between a proposal that is contested on the merits and a proposal that has unanimous agreement, and the practical result is that nothing ships either way. For the evidence dossier on stalled improvements, see *What Bitcoin's Stalled Proposals Tell You*.

Fallacy 8: "Spreading reviewer resources across multiple implementations weakens Bitcoin"

This argument assumes that the global pool of qualified reviewers is fixed and that diverting any portion of it to alternative implementations reduces the total security of the network. Atack addressed this directly in March 2026, describing it as a collectivist framing that treats the pie as a fixed size to be divided rather than something that can grow. His point was that developers who cannot be productive in Bitcoin Core's social environment for whatever reason, whether personal situation, work style, or geography, could contribute productively in a different organizational context. Welcoming them into alternative implementations grows the total reviewer pool rather than redistributing a fixed one.

Bitcoin Core's power nexus sits in approximately three or four locations across the US, UK, and a handful of European cities. The developer population capable of contributing to Bitcoin infrastructure is global, and the constraint on participation is not a shortage of qualified people but a set of structural and social barriers that the current governance model imposes and benefits from maintaining.

Conclusion

This post will be updated as additional fallacies appear, and I will direct people here when these arguments come up in discussion.

Bitcoin was designed to eliminate the need for trusted intermediaries in money. **Its reference implementation is governed by a small, concentrated group of them.**

Sources

- Bitcoin Governance Research: merge concentration, Gini, and process metrics cited in Fallacies 1–2 and 5 (`findings/GOVERNANCE_FRAMES.md`, `findings/ARCHIVE_GEMS.md`)
- Jon Atack, Plan B Forum (Rumble): maintainer selection and governance dynamics
- Jude C. Nelson, "Systemd: The Biggest Fallacies": template for this article's structure

SECURESOVEREIGN

Markdown is the canonical source: https://secsov.com/articles/bitcoin-core-the-biggest-fallacies/index.md

Formatted snapshot of https://secsov.com/articles/bitcoin-core-the-biggest-fallacies

Josh · Secure Sovereign