BITCOIN GOVERNANCE FRAMEWORK

VOL. 01 · 2026

Consensus Change Standards

A Legal and Technical Framework for Bitcoin Protocol Governance

Asaf Fulks

California State Bar No. 343622
Solo Bitcoin Miner · Full Node Operator

THE FORUM PRESS
a Fulks, Inc. company

First Edition · April 2026
Revised May 2026
asaffulkslaw.com

Consensus Change Standards: A Legal and Technical Framework for Bitcoin Protocol Governance

First Edition — April 2026 (Revised May 2026)

Author: Asaf Fulks, J.D.
California State Bar No. 343622
Admitted, U.S. District Court, Central District of California
asaffulkslaw.com

Published by The Forum Press, a Fulks, Inc. company
California
theforumpress.com

Available at: asaffulkslaw.com

This document does not constitute legal advice. The legal analysis contained herein is a general discussion of legal principles and does not create an attorney-client relationship between the author and any reader. The views expressed are the author's own and do not represent the views of any employer, client, or affiliated organization.

Bitcoin is an experimental technology. This document does not recommend for or against any particular consensus change proposal, including BIP-110.

ABSTRACT

Bitcoin has no formal process for evaluating proposed changes to its consensus rules. The Bitcoin Improvement Proposal (BIP) system provides a mechanism for proposing changes, but establishes no minimum standards that a proposal must meet before the community considers activation. There are no required review periods, no mandatory code audit standards, no agreed-upon activation thresholds, no chain split risk assessment methodology, and no framework for evaluating the legal and economic consequences of a failed activation.

This absence of standards has produced predictable results. The block size wars of 2015–2017 consumed years of developer time, fractured the community, and produced a contentious hard fork. The SegWit activation depended on the credible threat of a novel and untested User Activated Soft Fork (UASF) — BIP-148 — to break a deadlock that had persisted for eight months after miner signaling stalled. More recently, BIP-110 — a proposed temporary soft fork to restrict arbitrary data in Bitcoin transactions — was released with a buggy activation client, a dangerously low 55% activation threshold, and no formal review period. Its activation client was made available alongside stable Knots releases on node management platforms, with no risk disclosure or visual differentiation to indicate that selecting it would alter the node's consensus behavior.

This paper proposes a comprehensive framework for evaluating Bitcoin consensus change proposals. It draws on the history of Bitcoin's prior consensus changes, established principles of software engineering governance, and legal analysis of the liabilities created by reckless activation. The framework is designed to be practical, concrete, and immediately applicable. It is not a BIP. It does not propose changes to Bitcoin's code. It proposes standards for the process by which such changes should be evaluated, debated, and either adopted or rejected.

The author is a practicing litigator, solo Bitcoin miner, full node operator, and computer scientist. This framework is written from the intersection of those disciplines because the problems it addresses — governance, liability, technical risk, and economic consequence — cannot be adequately analyzed from any single perspective.

SECTION 1: THE PROBLEM

1.1 The Absence of Standards

Bitcoin's consensus rules are the most consequential code in the financial world. They govern the creation, transfer, and validation of an asset with a market capitalization exceeding one trillion dollars. Changes to these rules affect every participant in the network: miners who invest capital in hardware, node operators who validate transactions, developers who build applications, businesses that accept payment, and individuals who store wealth.

Despite these stakes, there is no formal standard governing how changes to consensus rules should be proposed, evaluated, reviewed, tested, activated, or — critically — rolled back if they fail. The BIP process, established in BIP-1 and refined in BIP-2, provides a template for writing proposals and a loose taxonomy of proposal types. It does not establish minimum standards for activation safety, mandatory review periods, code quality requirements, or chain split risk assessment.

The result is an ad hoc system in which each consensus change proposal invents its own activation mechanism, sets its own threshold, defines its own timeline, and is evaluated by the community with no consistent framework. Some proposals receive years of careful review. Others are pushed to activation within weeks. The difference between these outcomes is determined not by any institutional process but by the personalities, politics, and persuasive abilities of the participants.

1.2 BIP-110 as Case Study

BIP-110 — the Reduced Data Temporary Softfork — illustrates every failure mode that a governance framework should prevent. Originally proposed as BIP-444 in late October 2025, the proposal sought to restrict methods of embedding arbitrary data in Bitcoin transactions. Its stated goal was to protect Bitcoin's function as monetary infrastructure by limiting what proponents characterized as "spam" uses of block space.

The proposal's technical merits are debatable. Reasonable people disagree about whether inscriptions, ordinals, and large OP_RETURN payloads represent legitimate uses of Bitcoin's base layer or parasitic exploitation of shared infrastructure. This paper takes no position on that question. The problems with BIP-110 are procedural, not substantive:

Activation threshold. BIP-110 specified a 55% miner signaling threshold for a User Activated Soft Fork. This is dramatically lower than historical precedent. SegWit's BIP-9 deployment required 95% miner signaling; when that stalled, BIP-91 created a parallel mechanism at 80%, and the BIP-148 UASF threatened to reject non-signaling blocks entirely. Taproot activated at 90% via Speedy Trial. A 55% signaling threshold provides no assurance that the share of hashrate enforcing the new rules at activation will exceed 55%; the remainder may continue producing blocks valid under the legacy rules but invalid under the new ones. Signaling at lock-in is not equivalent to enforcement at activation, and the divergence between the two is precisely the mechanism by which low-threshold soft forks produce persistent minority chains. This is not a theoretical risk — it is a recipe for a chain split.

Code quality. The activation client, released in late 2025 as a fork of Bitcoin Knots, was found to contain significant bugs. Multiple Bitcoin developers reported that the client could not reliably fork the network, and that users running the code might accidentally fork themselves off both chains. Public commentary from reviewers raised concerns about the code's structure and quality; whatever tools or methods were used to produce it, the activation client did not receive the level of independent review that consensus-critical software demands. An activation client for a consensus change affecting a trillion-dollar network was released without the review burden customary for consensus-critical software.

Review period. From initial proposal to release of the activation client, BIP-110 moved through the pipeline in approximately six weeks — from the initial bitcoin-dev mailing list post on 26 October 2025 to release of the first activation client (v0.1rc1) on 10 December 2025. By comparison, SegWit was proposed in December 2015 and did not activate until August 2017 — a twenty-month process. Taproot was first proposed in January 2018 and activated in November 2021 — nearly four years. Six weeks is not a review period. It is a rush to deployment.

Activation client distribution. On at least one node management platform, the BIP-110 activation client was listed as a selectable version option in the same dropdown menu as stable Knots releases, with no warning label, risk disclosure, or visual differentiation. Selection was deliberate — a node operator had to affirmatively choose the BIP-110 version — but the presentation treated consensus-altering software identically to routine maintenance releases. A node operator who understood version management but not the implications of BIP-110 specifically could have activated consensus-changing code

believing it was a standard update. The absence of any risk disclosure at the point of selection is the governance failure, not the availability of the option itself.

Sunset mechanism. BIP-110 includes an automatic expiry at a defined block height, after which the new rules cease to be enforced. This is a meaningful improvement over proposals that lack any deactivation mechanism. However, there is no public evidence that the sunset mechanism was tested on testnet to confirm that the transition back to pre-activation rules would occur without consensus failures — a concern amplified by the significant bugs found in the activation client itself. A sunset clause that has not been demonstrably tested is a promise, not a guarantee. Furthermore, the one-year duration was chosen without empirical justification for why one year is the appropriate period, and the proposal contained no defined process for evaluation at the end of the enforcement period.

1.3 The Stakes

A failed consensus change activation is not a software bug that can be patched. It is a potential fracture of the monetary network. When a chain split occurs without replay protection, transactions valid on one chain may be valid on the other. Users can lose funds. Exchanges must choose which chain to list. Contracts denominated in Bitcoin become ambiguous. The economic damage is real, quantifiable, and potentially irreversible.

The most recent significant chain split without replay protection occurred in March 2013, when a database incompatibility between Bitcoin versions 0.7 and 0.8 caused a six-hour fork that included a successful double-spend attack (the technical post-mortem is BIP-50). The 2017 SegWit2x proposal came within days of producing another before being called off. The Bitcoin community has been fortunate. Fortune is not a governance strategy.

SECTION 2: HISTORICAL PRECEDENT

Bitcoin has undergone numerous consensus changes since its creation in 2009. The most significant of these provide instructive precedent for establishing governance standards.

2.1 P2SH (BIP-16) — 2012

Pay-to-Script-Hash was one of Bitcoin's first contentious soft forks. Competing proposals (BIP-16 and BIP-17) divided the developer community. Activation used a simple miner signaling threshold of 55% — the same threshold later adopted by BIP-110. The activation was messy, with miners signaling inconsistently and the community uncertain about which proposal would prevail.

Lesson: Low activation thresholds produce uncertainty even when the proposal itself has technical merit. P2SH ultimately succeeded because both competing proposals were small, low-risk changes. BIP-110's use of the same threshold for a far more consequential change ignores the increased risk.

2.2 The Block Size Wars (2015–2017)

The block size debate consumed more community energy, developer time, and political capital than any other event in Bitcoin's history. Multiple proposals competed: BIP-101 (8 MB blocks), BIP-102 (2 MB blocks), Bitcoin XT, Bitcoin Classic, Bitcoin Unlimited, and ultimately SegWit2x. The conflict produced the Bitcoin Cash hard fork in August 2017 and nearly produced a second split with SegWit2x in November 2017.

The block size wars demonstrated several principles that any governance framework must account for:

Miner signaling is unreliable as a measure of community consensus. Mining pools signaled support for proposals their users did not endorse. The SegWit2x "New York Agreement" secured signatures from companies representing over 80% of hashrate, yet the proposal collapsed when it became clear that node operators and users would not follow.

Economic nodes matter more than hashrate. Exchanges, payment processors, and major holders ultimately determine which chain carries economic value. A chain with 90% of the hashrate but no exchange listings and no merchant adoption is worthless. Governance frameworks must account for economic consensus, not merely miner signaling.

Hard forks are permanent and expensive. Bitcoin Cash continues to exist as a separate chain with a fraction of Bitcoin's value. Every hard fork fragments the ecosystem, confuses users, and creates legal ambiguity about which chain constitutes "Bitcoin" for contractual and regulatory purposes.

2.3 SegWit (BIP-141) — 2017

Segregated Witness was proposed in December 2015 and activated in August 2017 via BIP-9 version bits signaling with a 95% activation threshold. When miner signaling remained well below the 95% threshold through early 2017, the community developed BIP-148 — a User Activated Soft Fork that would have begun rejecting non-SegWit blocks on August 1, 2017, regardless of miner signaling.

The threat of BIP-148 — and the risk that it would cause a chain split — motivated miners to signal for SegWit. BIP-91 locked in on July 21, 2017, forcing miners to signal for BIP-141, and SegWit's lock-in was achieved before the August 1 UASF deadline. This episode established the UASF as a credible activation mechanism but also demonstrated its risks: had miners not capitulated, BIP-148 nodes would have split from the main chain.

Lesson: UASFs are a tool of last resort, not a standard activation mechanism. BIP-148 worked because SegWit had overwhelming community support and years of review. Applying the same mechanism to a proposal with weeks of review and marginal support — as BIP-110 attempted — is reckless.

2.4 Taproot (BIP-340/341/342) — 2021

Taproot is the gold standard for Bitcoin consensus change governance. Its design was first floated on the bitcoin-dev mailing list in January 2018 and formally specified in BIPs 340, 341, and 342 in 2020; it then underwent years of review, extensive formal analysis of its cryptographic primitives (Schnorr signatures), multiple rounds of community feedback, and a novel activation mechanism (Speedy Trial) that provided a defined three-month signaling window with a built-in timeout. The selection of Speedy Trial over BIP-8 with LOT=true and BIP-8 with LOT=false reflected an extended bitcoin-dev mailing list debate over the appropriate balance between activation speed, miner authority, and economic-node authority.

Taproot activated in November 2021 with 90% miner signaling within the Speedy Trial window. There was no chain split, no community fracture, and no economic disruption.

Lesson: Patience works. Review works. High thresholds work. Defined timelines with built-in failure modes work. Every element that made Taproot's activation successful was absent from BIP-110.

2.5 Summary of Activation Parameters

Proposal	Threshold	Review Period	Mechanism	Outcome
P2SH (2012)	55%	~3 months	Coinbase signal + flag day	Activated (messy)
SegWit (2017)	95%	20 months	BIP-9 + UASF	Activated
SegWit2x (2017)	80% (NYA)	~6 months	Hard fork	Canceled
Taproot (2021)	90%	~4 years	Speedy Trial	Activated
BIP-110 (2025–26)	55%	~6 weeks	UASF (Knots)	Stalled (≤0.15%)

The pattern is clear: successful consensus changes correlate with high activation thresholds, long review periods, and broad community buy-in. Failed or stalled proposals correlate with low thresholds, rushed timelines, and narrow support. This is not coincidence. It is the predictable result of governance dynamics that any framework must formalize.

SECTION 3: THE FRAMEWORK

The following framework establishes minimum standards for Bitcoin consensus change proposals. These standards are not intended to be enforced by code — Bitcoin has no central authority capable of enforcement. They are intended to serve as a publicly available benchmark against which the community can evaluate proposals. A proposal that meets these standards deserves serious consideration. A proposal that fails to meet them should be treated with appropriate skepticism.

3.1 Proposal Submission Requirements

A consensus change proposal submitted for community consideration should include, at minimum:

A. Problem Statement. A clear, specific description of the problem the proposal addresses, supported by empirical data where available. “Bitcoin should do one thing and do it well” is a philosophy, not a problem statement. “The UTXO set has grown by X% in Y months due to Z transaction type, imposing quantifiable costs of $W per node operator per year” is a problem statement.

B. Technical Specification. A complete technical specification of the proposed change, including all modified consensus rules, their interaction with existing rules, and a formal analysis of edge cases. The specification should be detailed enough to permit independent implementation.

C. Backward Compatibility Analysis. A comprehensive analysis of the proposal’s impact on existing transactions, wallets, applications, and use cases. This analysis should identify every category of transaction or script that would become invalid under the new rules and quantify, to the extent possible, the number of affected users and the value at risk.

D. Activation Mechanism. A fully specified activation mechanism including signaling method, threshold, signaling window duration, timeout behavior, and defined failure mode. The activation mechanism should be described with sufficient precision to permit independent implementation and verification.

E. Rollback Procedure. A defined procedure for deactivating the change if it produces unintended consequences. For proposals described as “temporary,” this must include a

self-executing sunset mechanism that does not require further community coordination to trigger.

F. Reference Implementation. A complete, functional reference implementation against a current release of Bitcoin Core or a compatible client. The reference implementation must include a comprehensive test suite.

3.2 Minimum Review Period

No consensus change proposal should proceed to activation signaling until it has completed a minimum review period. The appropriate length of this period depends on the scope and risk of the proposed change:

Category 1: Low-Risk Changes

Changes that tighten existing policy without altering the consensus boundary — for example, reducing default mempool relay limits. These are not consensus changes and do not require this framework. Individual node operators can adopt or reject them at will.

Category 2: Moderate-Risk Consensus Changes

Soft forks that add new validation rules without invalidating any currently valid transaction type. Examples include Taproot and SegWit, which expanded the set of valid scripts without restricting existing scripts. Minimum review period: **twelve months** from publication of a complete technical specification and reference implementation.

Category 3: High-Risk Consensus Changes

Soft forks that invalidate currently valid transaction types, restrict existing functionality, or could cause loss of funds for users with pre-existing transactions or scripts. BIP-110 falls into this category: it would have invalidated transactions that are currently valid, potentially trapping funds in scripts that use restricted opcodes. Minimum review period: **twenty-four months** from publication of a complete technical specification and reference implementation.

Category 4: Hard Forks

Any change that old nodes would reject. Hard forks carry the highest risk of permanent chain splits and should be avoided except in extraordinary circumstances, such as a critical security vulnerability discovered in the deployed protocol. Where they are nevertheless pursued, this framework proposes the following minimum standards: (1) a thirty-six-month review period from publication of a complete technical specification and

reference implementation; (2) explicit replay protection or a published rationale for its absence; (3) demonstrated economic-node support, including affirmative commitments from major exchanges, custodians, and payment processors; and (4) a published chain-split contingency plan addressing user communication, exchange coordination, and the contractual question of which chain constitutes "Bitcoin" for outstanding obligations. The Bitcoin community's strong norm against hard forks — reinforced by the block size wars — remains the most effective deterrent, but the absence of a written standard has historically left proponents free to define their own.

3.3 Code Audit Requirements

The activation client for any consensus change must meet the following code quality standards before activation signaling begins:

A. Independent review. The reference implementation must be reviewed by a minimum of three developers who are not affiliated with the proposal's authors. Reviewers should have demonstrated competence in Bitcoin protocol development, as evidenced by prior contributions to Bitcoin Core, Bitcoin Knots, or another consensus-compatible implementation.

B. Test coverage. The reference implementation must include unit tests covering all new validation rules, integration tests demonstrating compatibility with existing valid transactions, and regression tests for all identified edge cases. Test results must be publicly reproducible.

C. Testnet deployment. The activation client must be deployed on Bitcoin's public testnet for a minimum of three months before mainnet activation signaling begins. The testnet deployment must demonstrate successful activation, enforcement of new rules, and — critically — successful deactivation if the proposal includes a sunset clause.

D. Fuzzing and adversarial testing. The reference implementation should be subjected to automated fuzzing and adversarial testing to identify vulnerabilities that could be exploited during or after activation. This is particularly important for proposals that restrict transaction types, as attackers may attempt to craft transactions that exploit ambiguities in the new rules.

E. No AI-generated code without disclosure and review. The use of AI coding assistants in developing activation clients is not inherently problematic, but it must be disclosed and the resulting code must receive heightened scrutiny. AI-generated code in

consensus-critical software is analogous to AI-generated legal filings — the tool can accelerate production, but the professional remains responsible for the output's correctness.

3.4 Activation Threshold Standards

The activation threshold for a consensus change should reflect both the risk of the change and the cost of a failed activation. The following thresholds are proposed as minimum standards:

Activation Mechanism Design Space

Threshold selection is downstream of mechanism selection. The activation mechanisms employed in Bitcoin's history form a small but instructive design space. Each represents a distinct trade-off between activation speed, fail-safe behavior, and the locus of signaling authority.

BIP-9 (Version Bits with Timeout and Delay). The original modern activation mechanism, used for CSV and SegWit. Miners signal readiness via version-bit flags; activation occurs when the configured threshold of blocks within a 2,016-block window signal readiness. BIP-9 includes a timeout: if the threshold is not met within a defined window, the deployment expires. BIP-9's principal weakness is that it grants miners effective veto power — a small mining pool coalition can block activation by refusing to signal, even where node operators, exchanges, and users overwhelmingly support the change. SegWit's eight-month stall demonstrated this failure mode in practice.

BIP-91 (Reduced Threshold MASF). A direct response to SegWit's stalled BIP-9 deployment. BIP-91 lowered the lock-in threshold to 80% and made signaling itself compulsory: BIP-91-enforcing miners would reject blocks that did not signal for SegWit. This created the coordination pressure that achieved SegWit's lock-in in July 2017. BIP-91 demonstrated that compulsory signaling — not just threshold reduction — can break a deadlock, at the cost of accepting some chain-split risk during the coordination period.

BIP-148 (User Activated Soft Fork). A flag-day mechanism: BIP-148-enforcing nodes would begin rejecting non-SegWit-signaling blocks on August 1, 2017, regardless of miner readiness. BIP-148 transferred activation authority from miners to economic nodes. Its credible threat is widely credited with motivating the miner capitulation that produced BIP-91. The mechanism is structurally riskier than miner-activated alternatives — if miners refuse to comply, the BIP-148 chain splits from the legacy chain — and is appropriate only where the underlying proposal has overwhelming economic support.

Speedy Trial. The activation mechanism used for Taproot in 2021. A bounded BIP-9 deployment with a short signaling window (approximately three months) and a high threshold (90%). If the threshold is reached, activation occurs after a defined delay; if not, the proposal expires without activation, freeing the deployment slot for revision or alternative proposals. Speedy Trial trades the certainty of activation for the certainty of a defined timeline. It is the most cautious of the modern mechanisms and is the closest existing match to the standards proposed in this framework.

LOT=true / LOT=false. A parameter that arose in BIP-8 during Taproot's activation debates. Lockin On Timeout (LOT) specifies whether a deployment that fails to achieve miner signaling within its window should nevertheless activate by user mandate at timeout. LOT=true converts a deployment into an effective UASF if miners do not cooperate; LOT=false makes it a clean Speedy Trial-style timeout. The LOT debate of early 2021 surfaced the central question of every activation mechanism: who holds final authority over consensus rule changes, and what happens when they disagree?

This framework treats activation mechanism selection as a deliberate design choice subject to the standards in this section. A proposal that meets the substantive standards but uses an inappropriate activation mechanism remains deficient. A proposal that uses an appropriate mechanism but fails the substantive standards is no more defensible. Mechanism and merit are both gating.

Miner-Activated Soft Fork (MASF)

Minimum threshold: **90% of hashrate** measured over a defined signaling period of at least two weeks (2,016 blocks). This threshold is consistent with Taproot's successful activation and ensures that the risk of a chain split is minimized. Thresholds below 80% should be considered presumptively dangerous. Thresholds below 60% are reckless and should be rejected regardless of the proposal's merits.

User-Activated Soft Fork (UASF)

UASFs should be reserved for situations in which a proposal has demonstrated overwhelming community support but miner signaling is blocked by a small number of mining pool operators acting against their users' interests. The original UASF rationale, articulated in BIP-148 and the contemporaneous bitcoin-dev mailing list discussions of early 2017, held that user-activated mechanisms transfer activation authority from miners to economic nodes when the two diverge. A UASF is an extraordinary measure. It should not be the default activation mechanism for any proposal.

A UASF should only proceed when: (1) the proposal has completed its full minimum review period; (2) the proposal has demonstrated broad support among economic nodes, exchanges, and major holders; (3) the UASF includes a defined activation date set at least six months in the future to provide time for preparation; and (4) the UASF proponents have published a detailed chain split contingency plan.

The 55% Problem

BIP-110's 55% threshold deserves specific discussion because it illustrates the danger of low thresholds. A 55% signaling threshold permits activation under conditions in which the hashrate actually enforcing the new rules at activation may not materially exceed 55%. The risk of a persistent minority chain depends on enforcement, not signaling — the two are not the same. If a meaningful share of post-activation hashrate produces blocks valid under the legacy rules but invalid under the new ones, those blocks will be rejected by activated nodes. The activated chain will fall behind the non-activated chain in cumulative proof of work, and activated nodes will be on a minority chain that, by Bitcoin's own rules, is not the valid chain. The 95% threshold used by SegWit and the 90% threshold used by Taproot exist precisely to drive that probability toward zero. A 55% threshold makes a persistent split foreseeable.

A 55% threshold does not safely activate a soft fork. It produces the conditions historically associated with persistent minority chains: weeks or months of competing tips, ambiguous economic status for transactions confirmed on either side, and unresolved questions about the meaning of "Bitcoin" in contracts and on exchanges. The fact that P2SH used the same threshold in 2012 is not persuasive precedent: P2SH was a narrow, low-risk change to a network with a fraction of today's value and user base. The stakes have changed. The standards must change with them.

Quantifying the Risk

The danger of a 55% threshold can be quantified. Model post-activation hashrate as enforcing (share E) and non-enforcing (share 1−E). Treat block production as a Bernoulli process: each block is, with probability E, valid under the new rules and built on the enforcing chain; with probability 1−E, it violates the new rules and extends a competing non-enforcing chain. Enforcing nodes follow the longest chain that is valid under the new rules; non-enforcing nodes follow the longest chain absolutely.

The difference in cumulative work between the chains is a random walk with drift 2E−1 per block. The probability that, at some point during activation, the non-enforcing chain

temporarily exceeds the enforcing chain by k blocks is well-approximated by $((1-E)/E)^k$. Substituting for k = 6, the depth at which most exchanges credit deposits as final:

At E = 0.55, the probability of a six-block deficit on the enforcing chain during activation is approximately $(0.45/0.55)^6 \approx 0.30$ — roughly thirty percent. At E = 0.90, the same probability is approximately $(0.10/0.90)^6 \approx 1.9\times10^{-6}$. At E = 0.95, it is approximately $(0.05/0.95)^6 \approx 2.1\times10^{-8}$.

Each such deficit represents an opportunity for a reorganization that orphans transactions previously confirmed on the enforcing chain. Exchanges that require six confirmations before crediting a deposit would, at E = 0.55, see roughly one in three confirmation chains exposed to a reorg event at some point during activation. At E = 0.90 or above, such events are statistically nonexistent over the relevant time horizons. The 55%, 90%, and 95% thresholds are not points on a linear spectrum of safety: they differ by five to seven orders of magnitude in expected reorganization exposure during activation.

This analysis simplifies several factors. It assumes non-enforcing miners produce blocks at the natural rate proportional to their hashrate; in practice, miners may defect to whichever chain becomes profitable, accelerating consolidation. It also assumes that every non-enforcing block contains a transaction that activated nodes would reject; for proposals like BIP-110 that target common transaction patterns, this assumption is largely realized, but for narrower changes the effective fork rate is lower. Both factors affect the absolute magnitudes; neither changes the qualitative conclusion that low-threshold activations are quantitatively distinct from high-threshold ones in their risk profile.

3.5 Chain Split Risk Assessment

Every consensus change proposal should include a formal chain split risk assessment addressing, at minimum:

A. Hashrate distribution analysis. What percentage of current hashrate is operated by pools or miners likely to adopt the change? What percentage is likely to reject it? Is there a credible path to the activation threshold, or is the proposal being pushed despite inadequate support?

B. Economic node analysis. Have major exchanges, payment processors, and infrastructure providers indicated support for the change? A consensus change that

activates without exchange support creates immediate economic disruption, as users cannot deposit or withdraw funds until exchanges upgrade.

C. Replay protection. If a chain split occurs, does the proposal include replay protection to prevent transactions from being valid on both chains? If not, what is the expected impact on users who are unaware of the split?

D. Contingency plan. What happens if the activation fails? What happens if the activation succeeds but produces a persistent minority chain? Who is responsible for communicating the split to users, and how?

3.6 Sunset and Reversibility Requirements

Proposals described as "temporary" must include a self-executing sunset clause. This means that the consensus rules imposed by the proposal must automatically expire at a defined block height or timestamp without requiring any further community action. The burden of continuation should fall on proponents of the change, not on opponents.

Specifically, a valid sunset clause must:

A. Define an exact block height or median time past (MTP) at which the new rules cease to be enforced.

B. Be implemented in the activation client such that nodes automatically revert to pre-activation consensus rules upon reaching the sunset trigger.

C. Be tested on testnet to confirm that deactivation works correctly and does not itself produce consensus failures.

D. Not require a subsequent soft fork, hard fork, or software update to effectuate deactivation.

A proposal that describes itself as temporary but requires active intervention to expire is not temporary. It is permanent with a stated aspiration.

SECTION 4: LEGAL ANALYSIS

The legal implications of Bitcoin consensus changes are largely unexplored. This is partly because Bitcoin's decentralized nature complicates the application of traditional legal frameworks, and partly because no chain split has yet produced litigation with reported opinions. But the absence of precedent does not mean the absence of liability. The following analysis applies established tort and contract principles to the specific risks created by reckless consensus change activation.

4.1 Negligence

Tort liability for negligence requires a duty of care, a breach of that duty, causation, and damages. The threshold question is whether the developers of a consensus change activation client owe a duty of care to node operators and users who run their software.

Under traditional tort principles, a person who creates a dangerous instrumentality and places it into the stream of commerce owes a duty of care to foreseeable users. An activation client for a Bitcoin consensus change is software that, if defective, can cause direct financial harm to its users. The analogy to products liability is imperfect — most activation clients are distributed as free, open-source software — but open-source licenses do not categorically eliminate tort liability, particularly where the developer actively encourages adoption and knows that defects could cause financial loss.

A second doctrinal obstacle warrants attention. California's economic loss rule generally bars recovery in tort for purely economic harm absent physical injury, property damage, or a special relationship giving rise to an independent duty. See Aas v. Superior Court, 24 Cal.4th 627 (2000); Robinson Helicopter Co. v. Dana Corp., 34 Cal.4th 979 (2004). The rule is a genuine impediment to a negligence theory against open-source developers whose users hold no contract with them and whose harm is financial rather than physical. Two routes around the rule remain viable. The multi-factor test of Biakanja v. Irving, 49 Cal.2d 647 (1958), permits a duty of care to non-contracting parties where the developer's conduct was intended to affect the user, the harm was foreseeable, and policy supports liability. And where the developer makes representations about the safety or readiness of an activation client on which users foreseeably rely, negligent misrepresentation under Restatement (Second) of Torts § 552 supplies a recognized cause of action without confronting the economic loss bar. Both routes require facts beyond the bare release of buggy software; both are available on facts of the kind BIP-110 presents.

BIP-110's activation client illustrates the potential for negligence liability. The client was released with known bugs. Developers publicly identified defects that could cause users to fork themselves off the network. Despite these warnings, the client was distributed and its adoption was promoted on social media. If a user had run this client and suffered a financial loss — for example, by mining blocks on a minority chain that were subsequently orphaned — a negligence claim against the client's developer would face challenging but not insurmountable hurdles.

The strongest argument against liability is assumption of risk: users who run experimental software on a production network are arguably assuming the risk of loss. But assumption of risk is an affirmative defense, not a bar to the existence of a duty. And the defense is weaker when the software is presented alongside production releases in a platform's version management system, as BIP-110 was on at least one node management platform.

4.2 Tortious Interference

A reckless consensus change activation that causes a chain split could give rise to claims of tortious interference with contractual relations or business expectancy. Consider the following scenario: a business accepts Bitcoin as payment under a contract that specifies payment in "Bitcoin." A chain split occurs, and the payor delivers coins on the minority chain. The payee argues that "Bitcoin" means the majority chain. The resulting dispute was proximately caused by the chain split, which was proximately caused by the reckless activation.

The question of which chain constitutes "Bitcoin" after a split has no settled legal answer. During the Bitcoin Cash fork, exchanges and contracts generally treated the chain with the most accumulated proof of work and the greatest economic activity as "Bitcoin." But this convention is informal and could be challenged.

Proponents of consensus changes that carry a material risk of chain split should consider whether their actions could expose them to tortious interference claims. This is particularly relevant when the proponent is a company or public figure whose advocacy for the change is well-documented and whose economic interest in the change's success is apparent.

Two elements warrant emphasis. Tortious interference with contract requires intentional acts designed to disrupt performance, not merely conduct that has the foreseeable effect of doing so. See Pacific Gas & Electric Co. v. Bear Stearns & Co., 50 Cal.3d 1118 (1990).

Tortious interference with prospective economic advantage requires, in addition, an independently wrongful act — conduct unlawful for reasons other than the interference itself. See Della Penna v. Toyota Motor Sales, U.S.A., Inc., 11 Cal.4th 376 (1995); Korea Supply Co. v. Lockheed Martin Corp., 29 Cal.4th 1134 (2003). A developer who promotes an activation client out of ideological conviction will not typically meet either standard. The live claims are those in which the proponent of a consensus change has documented economic exposure to the outcome — a financial position whose value depends on which chain prevails — and the proponent's promotion can fairly be characterized as instrumental rather than principled. Such cases are not hypothetical. The framework's documentary requirements — problem statement, backward compatibility analysis, contingency plan — are themselves designed to create the record by which such inquiries can be conducted.

4.3 Fiduciary Duties

Some legal scholars have argued that Bitcoin developers owe fiduciary duties to Bitcoin holders, analogous to the duties owed by corporate directors to shareholders. The most developed version of this argument is Angela Walch, "In Code(rs) We Trust: Software Developers as Fiduciaries in Public Blockchains," in Regulating Blockchain: Techno-Social and Legal Challenges (Hacker, Lianos, Dimitropoulos & Eich eds., Oxford University Press 2019), which contends that core protocol developers exercise discretionary authority over property interests in a manner that triggers fiduciary obligations under traditional principles. This argument has been most fully developed in the Tulip Trading litigation in the United Kingdom. In Tulip Trading Ltd v van der Laan [2023] EWCA Civ 83, the Court of Appeal of England and Wales did not adjudicate whether Bitcoin Core developers owe a fiduciary duty to holders of lost Bitcoin. It held only that the claim was sufficiently arguable to survive a strike-out application and should proceed to trial. The case thus stands not for the proposition that developers owe such duties, but for the proposition that the question is justiciable on appropriate facts.

This paper does not take a position on whether developers owe fiduciary duties in general. However, the analysis is relevant to consensus change governance for the following reason: if a developer promotes a consensus change, distributes an activation client, and the change causes financial harm, the question of whether the developer's conduct constituted a breach of duty will be evaluated against the standard of care exercised in the process.

A developer who follows a rigorous governance framework — adequate review, thorough testing, conservative activation thresholds, and transparent communication of risks — has a strong defense against any claim of breach. A developer who releases a buggy client with a low activation threshold and no review period does not.

4.4 Mining and Node Operator Liability

Miners and node operators who adopt activation clients also face potential liability exposure. A mining pool that signals for a consensus change bears some responsibility for the consequences of that signaling, particularly if the pool operator has not communicated the risks to the pool's users (individual miners).

For solo miners — a category that includes the author — the liability analysis is simpler: a solo miner who runs an activation client is assuming the risk of their own operation. But a pool operator who signals on behalf of thousands of connected miners has a duty to those miners to exercise reasonable care in evaluating the consensus change. Mining pool operators who signal for poorly reviewed proposals with low activation thresholds are taking risks with other people's hashrate and, by extension, other people's money. The most likely cause of action is breach of contract: pool terms of service and service agreements typically obligate the operator to direct connected hashrate with reasonable care, and signaling for a consensus change that exposes connected miners to orphaned blocks is the kind of decision that brings such obligations into play. Where the pool agreement is silent on consensus signaling, the implied covenant of good faith and fair dealing may supply the duty. Negligence remains available as a backstop, subject to the economic-loss-rule limitations discussed above.

4.5 Regulatory Consequences

Beyond direct liability, a chain split creates regulatory uncertainty. Tokens on both chains may be treated as separate assets for tax purposes, requiring holders to determine cost basis allocation. Exchanges may be required to support both chains or face claims from customers who hold tokens on the unsupported chain. Custodians may face conflicting obligations to clients.

These regulatory consequences are not hypothetical — they occurred during the Bitcoin Cash fork in 2017 and required guidance from the IRS (Revenue Ruling 2019-24) to resolve. A governance framework that minimizes the risk of chain splits also minimizes regulatory disruption. This is a feature, not a byproduct.

SECTION 5: PROPOSED STANDARDS

The following checklist synthesizes the framework described in Sections 3 and 4 into a concrete set of standards. A consensus change proposal that meets all of these standards deserves serious community consideration. A proposal that fails to meet them should be treated with appropriate skepticism and, if it proceeds to activation without meeting them, should be actively resisted.

5.1 The Consensus Change Readiness Checklist

A. Proposal Quality

1. Does the proposal include a clear, empirically supported problem statement?

2. Does the proposal include a complete technical specification sufficient for independent implementation?

3. Does the proposal include a backward compatibility analysis identifying all affected transaction types, scripts, and use cases?

4. Does the proposal include a fully specified activation mechanism with defined thresholds, timelines, and failure modes?

5. Does the proposal include a rollback procedure? If described as temporary, does it include a self-executing sunset clause?

B. Code Quality

6. Has the reference implementation been reviewed by at least three independent developers with demonstrated Bitcoin protocol expertise?

7. Does the reference implementation include comprehensive unit, integration, and regression tests?

8. Has the activation client been deployed on testnet for at least three months?

9. Has the code been subjected to fuzzing and adversarial testing?

10. If AI coding tools were used, has this been disclosed and has the code received heightened review?

C. Activation Safety

11. Is the activation threshold at or above 90% for a MASF?

12. If a UASF mechanism is proposed, has the proposal completed its full minimum review period and demonstrated broad support among economic nodes?

13. Has a chain split risk assessment been completed and published?

14. Does the proposal include replay protection or, where the proposal is a soft fork that cannot produce a chain split absent miner defection, a documented rationale for its absence?

15. Has the activation timeline been set at least six months from the publication of the final activation client?

D. Community Process

16. Has the proposal completed the minimum review period for its risk category (twelve months for moderate-risk, twenty-four months for high-risk)?

17. Has the proposal been discussed in public forums with participation from a diverse cross-section of the community (developers, miners, node operators, businesses, users)?

18. Have major exchanges and infrastructure providers been consulted regarding the proposal's impact on their operations?

19. Has the proposal's author published a chain split contingency plan?

20. Has the proposal been evaluated against this framework, with the results published?

5.2 Scoring

Each of the twenty criteria above receives a binary score: met or not met. Proposals are classified as follows:

20/20: Green. The proposal has met all minimum standards and is ready for activation signaling.

15–19/20: Yellow. The proposal has met most standards but has identified gaps that should be addressed before activation.

10–14/20: Orange. The proposal has significant deficiencies and should not proceed to activation signaling until they are resolved.

Below 10/20: Red. The proposal fails to meet minimum standards for serious consideration. Activation should be actively resisted through: (a) public documentation

of the deficiencies measured against this framework; (b) coordination among economic nodes — exchanges, custodians, payment processors, and major holders — to refuse to recognize the activated chain as "Bitcoin" for purposes of contracts, deposits, and withdrawals; (c) running non-signaling, non-activation client software; and (d) where activation proceeds despite these objections, publication of a chain-split contingency plan to ensure user safety and minimize the economic damage of the resulting fracture.

For reference, BIP-110 would score approximately 3/20 under this framework: credit for having a technical specification, a defined activation mechanism, and a self-executing sunset clause (though one whose deactivation was not demonstrably tested on testnet); no credit for adequate review period, code quality, independent review, activation safety, or community process. Taproot would score approximately 18/20: criteria 19 (a chain-split contingency plan in the form proposed here) and 20 (evaluation against this framework) postdate Taproot's activation and so are scored as not met, while every other criterion was met or exceeded.

SECTION 6: OBJECTIONS AND RESPONSES

6.1 “Bitcoin has no governance.”

Bitcoin has no centralized governance. It has governance. Every consensus change that has ever been adopted required coordination among developers, miners, node operators, and economic actors. The process by which this coordination occurs — however informal — is governance. This framework does not propose centralized governance. It proposes minimum standards for evaluating proposals within Bitcoin’s existing decentralized governance structure.

6.2 “Anyone can run whatever software they want.”

True. And this framework does not propose restricting that right. Node operators are free to run any software they choose. This framework proposes that the community develop shared standards for evaluating proposals, so that node operators can make informed decisions. A node operator who runs an activation client that fails every criterion in this framework is exercising their right. They are also assuming quantifiable risks that they may not fully understand. Providing a framework for understanding those risks serves the same function as securities disclosure: it does not restrict choice, it informs it.

6.3 “This framework would prevent necessary changes.”

This framework would slow down reckless changes. It would not prevent necessary ones. SegWit and Taproot both would have passed this framework with high scores. The changes this framework would impede are the ones that should be impeded: poorly reviewed, inadequately tested, rashly activated proposals that put the network at risk.

Bitcoin’s value proposition is stability, predictability, and resistance to arbitrary change. A framework that makes consensus changes harder to execute is aligned with that value proposition, not contrary to it.

6.4 “Who decides whether the standards are met?”

Everyone. And no one. This framework is a tool, not an authority. Any member of the community can evaluate a proposal against these criteria and publish the results. There is no certification body, no approval committee, and no veto power. The framework’s authority derives from its usefulness. If the community finds it useful, it will be adopted. If not, it will be ignored. That is how governance works in a decentralized system.

6.5 “The legal analysis is speculative.”

All legal analysis of novel situations is, to some degree, speculative. No court has ruled on the liability of a Bitcoin developer for a chain split caused by a reckless activation. But the absence of precedent does not mean the absence of risk. The legal principles applied in Section 4 — negligence, tortious interference, fiduciary duty — are well established. Their application to Bitcoin governance is novel but not unprecedented. Courts routinely apply existing legal frameworks to new technologies. The question is not whether these principles apply, but how. This paper offers an analysis, not a prediction.

SECTION 7: CONCLUSION

Bitcoin is the most consequential monetary experiment in human history. Its consensus rules govern the creation and transfer of value for millions of people and the storage of wealth measured in trillions of dollars. Changes to these rules should be evaluated with a rigor commensurate with their stakes.

The current system — in which proposals are evaluated ad hoc, activation mechanisms are invented on the fly, review periods range from weeks to years with no standard, and the community's only tools for evaluating proposals are Twitter threads and GitHub comments — is inadequate. It has produced near-catastrophic chain splits, wasted years of developer time on governance disputes, and created opportunities for reckless actors to push poorly considered changes to activation. BIP-110 stalled, but its failure to activate is a fact of community vigilance, not of structural protection. The next proposal of its kind will arrive on the same terms — no required review, no minimum code quality, no agreed-upon threshold — unless this gap is filled.

This framework does not solve the fundamental challenge of decentralized governance. No framework can. What it provides is a common vocabulary, a shared set of criteria, and a concrete checklist against which proposals can be evaluated. It shifts the burden of proof onto proponents of change — where it belongs — and provides the community with a structured way to say: this proposal is not ready.

 It is available on GitHub for community review and amendment. If it is useful, it will be used. If it can be improved, it should be improved. That is how Bitcoin works. That is how Bitcoin's governance should work too.

Asaf Fulks
Asaf Fulks Law | asaffulkslaw.com
California State Bar No. 343622
Published by The Forum Press, a Fulks, Inc. company | April 2026 · Revised May 2026

REFERENCES

Bitcoin Improvement Proposals

BIP-1: BIP Purpose and Guidelines. Amir Taaki, 2011. github.com/bitcoin/bips/blob/master/bip-0001.mediawiki

BIP-2: BIP Process, Revised. Luke Dashjr, 2016. github.com/bitcoin/bips/blob/master/bip-0002.mediawiki

BIP-8: Version Bits with Lock-in by Height. Shaolin Fry, Luke Dashjr, 2017. github.com/bitcoin/bips/blob/master/bip-0008.mediawiki

BIP-9: Version Bits with Timeout and Delay. Pieter Wuille, Peter Todd, Greg Maxwell, Rusty Russell, 2015. github.com/bitcoin/bips/blob/master/bip-0009.mediawiki

BIP-16: Pay to Script Hash. Gavin Andresen, 2012. github.com/bitcoin/bips/blob/master/bip-0016.mediawiki

BIP-17: OP_CHECKHASHVERIFY (CHV). Luke Dashjr, 2012. github.com/bitcoin/bips/blob/master/bip-0017.mediawiki

BIP-50: March 2013 Chain Fork Post-Mortem. Gavin Andresen, 2013. github.com/bitcoin/bips/blob/master/bip-0050.mediawiki

BIP-91: Reduced threshold Segwit MASF. James Hilliard, 2017. github.com/bitcoin/bips/blob/master/bip-0091.mediawiki

BIP-101: Increase Maximum Block Size. Gavin Andresen, 2015. github.com/bitcoin/bips/blob/master/bip-0101.mediawiki

BIP-102: Block Size Increase to 2MB. Jeff Garzik, 2015. github.com/bitcoin/bips/blob/master/bip-0102.mediawiki

BIP-110: Reduced Data Temporary Softfork (originally proposed as BIP-444). Dathon Ohm, 2025. github.com/bitcoin/bips/blob/master/bip-0110.mediawiki

BIP-141: Segregated Witness (Consensus Layer). Eric Lombrozo, Johnson Lau, Pieter Wuille, 2015. github.com/bitcoin/bips/blob/master/bip-0141.mediawiki

BIP-148: Mandatory Activation of Segwit Deployment. Shaolin Fry, 2017. github.com/bitcoin/bips/blob/master/bip-0148.mediawiki

BIP-340: Schnorr Signatures for secp256k1. Pieter Wuille, Jonas Nick, Tim Ruffing, 2020. github.com/bitcoin/bips/blob/master/bip-0340.mediawiki

BIP-341: Taproot: SegWit Version 1 Spending Rules. Pieter Wuille, Jonas Nick, Anthony Towns, 2020. github.com/bitcoin/bips/blob/master/bip-0341.mediawiki

BIP-342: Validation of Taproot Scripts. Pieter Wuille, Jonas Nick, Anthony Towns, 2020. github.com/bitcoin/bips/blob/master/bip-0342.mediawiki

Legal Authorities

Aas v. Superior Court, 24 Cal.4th 627 (2000).

Biakanja v. Irving, 49 Cal.2d 647 (1958).

Della Penna v. Toyota Motor Sales, U.S.A., Inc., 11 Cal.4th 376 (1995).

Korea Supply Co. v. Lockheed Martin Corp., 29 Cal.4th 1134 (2003).

Pacific Gas & Electric Co. v. Bear Stearns & Co., 50 Cal.3d 1118 (1990).

Robinson Helicopter Co. v. Dana Corp., 34 Cal.4th 979 (2004).

Tulip Trading Ltd v van der Laan [2023] EWCA Civ 83, Court of Appeal of England and Wales (3 February 2023). judiciary.uk/wp-content/uploads/2023/02/Tulip-v-Van-Der-Laan-judgment-030223.pdf

Rev. Rul. 2019-24, 2019-44 I.R.B. 1004 (Oct. 9, 2019) (tax treatment of cryptocurrency hard forks). irs.gov/pub/irs-drop/rr-19-24.pdf

Restatement (Second) of Torts § 552 (Am. Law Inst. 1977).

Walch, Angela. “In Code(rs) We Trust: Software Developers as Fiduciaries in Public Blockchains.” In Regulating Blockchain: Techno-Social and Legal Challenges, edited by Philipp Hacker, Ioannis Lianos, Georgios Dimitropoulos & Stefan Eich, 58–81. Oxford: Oxford University Press, 2019.

Software Releases

Bitcoin Core v30.2 Release Notes. Bitcoin Core Project, 10 January 2026. bitcoincore.org/en/releases/30.2/

Bitcoin Knots v29.3.knots20260210. Bitcoin Knots Project, 2026. github.com/bitcoinknots/bitcoin/releases

License

DISCLAIMER

This document is provided for informational and educational purposes only. It does not constitute legal advice. The legal analysis contained in Section 4 is a general discussion of legal principles and does not create an attorney-client relationship between the author and any reader. Readers should consult with qualified legal counsel regarding the application of these principles to their specific circumstances.

The author is a practicing attorney licensed in the State of California (State Bar No. 343622) and admitted to the United States District Court for the Central District of California. The views expressed in this document are the author's own and do not represent the views of any employer, client, or affiliated organization. This document is published by The Forum Press, a Fulks, Inc. company.

Bitcoin is an experimental technology. Running a Bitcoin node, mining Bitcoin, and participating in consensus changes all carry financial and technical risks. This document does not recommend any particular course of action and specifically does not recommend for or against any particular consensus change proposal, including BIP-110.