The Price of Time: Deterrence as the Third Pillar of Payment Channel Security

Rong Qian[1⋆], Yu Cheng[1⋆⋆], Mengrun Chen[1⋆⋆⋆], Yuchang Zhang[2], and Zengli Guo[3]

[1] Department of Cyberspace Security, Beijing Electronic Science and Technology Institute, Beijing, China
rqian@besti.edu.cn, chengyu_cs@zjut.edu.cn, 13015310107@163.com

[2] Graduate Department of North China Institute of Computer Systems Engineering
zhangyuchang26@mails.ucas.ac.cn

[3] Faculty of Information Science and Engineering, Ocean University of China, Qingdao 266100, China
isgoodfishnew@gmail.com

Abstract. Safety and liveness are the two classical pillars of payment channel security. Both price attacks that merely occupy channel resources, such as jamming, at zero. We propose a third pillar, *deterrence*: harm carries a provable price. We give its theory for occupied shared-state resources. Within a holding-cost framework, we characterize occupation pricing for lock-resolve channels. Time-dependence in the occupier's liability is necessary for any deterrence: today's time-independent pricing admits unbounded griefing, with damage-to-cost ratios above 10^7. Linear time-proportional liability achieves the tight rate. Its two faces, deadline-proportional upfront fees and hold-proportional penalty bonds, differ exactly in what they charge honest traffic. Our main theorem maps the landscape of four desiderata: deterrence against colluding sinks, penalty-only honesty, zero intermediary lockup, and path privacy. Three of them are jointly unachievable, and every remaining combination is achieved by a matching construction; onion-message fee-swallowing gives a second enforcement instance. Single-channel enforcement is characterized exactly: penalty-only pricing is unilaterally enforceable if and only if scripts expose time. The proof is a three-case exhaustion, and it explains why pre-BIP65 Bitcoin could not support trustless occupation pricing. Our flagship fused-bond construction runs on Bitcoin today, with no covenants, oracles, or miner assumptions; it improves slot-jamming deterrence by eight orders of magnitude at zero net cost to punctual honest payments. For multi-hop payments we prove the forwarding trilemma, settling a 2020 conjecture of the Lightning engineering community, and we quantify the price of privacy, including a new $O(1/n)$ extreme-value side channel on the sender's path length.

⋆ AI tools were used to assist with code generation and debugging, and for language polishing of the manuscript.

⋆⋆ Corresponding author.

⋆⋆⋆ Corresponding author.

Keywords: Payment channels · Channel jamming · Griefing · Deterrence · Penalty bonds · Lightning Network · Path privacy.

1 Introduction

Safety and liveness are the two classical pillars of payment channel security: no one can steal funds, and the channel keeps operating. Both pillars price the same thing at zero: *occupation*. A payment channel holds a bounded number of contract slots and bounded capacity; an attacker that fills the slots with doomed contracts, declared for the maximum window and left to time out, paralyzes the channel while paying almost nothing. Under current Lightning parameters the damage-to-cost ratio of such slot jamming exceeds 2×10^7, and jamming the entire network of 32,674 channels for a month would cost about 10.4 BTC in forfeited bonds under the defense we propose, against effectively nothing today. Within eighteen months, the engineering community has converged three times, empirically and independently, on the shape of a fix: charging by occupation time. Law's three-fee design [14], the onion-message fee-swallowing observation [12], and Riard's conditional-message-transfer sketch [20], the last self-described as "prolegomena" to a formalization. None of these came with theorems; the question of what deterrence *requires*, what it *costs*, and what is *impossible* remained open. It was posed by Joost Jager on the Lightning mailing list in 2020 as a conjecture that a trustless scheme "may even be theoretically impossible" [10], and has stood for six years.

This paper answers it. We take *deterrence*, the requirement that harm carries a provable price, as the *third pillar* of channel security, and give its theory for held shared-state resources.

A layered theory. Within a small holding-cost model (§3), we prove three characterizations for a single channel. *Pricing:* time-dependence in the occupier's liability is necessary for any deterrence, and linear time-proportional liability achieves the tight rate $\lambda \leq \max(r/\alpha, \rho/\beta)$ (Theorem 1). *Enforcement:* penalty-only pricing (honest behavior free, violators forfeit) is unilaterally enforceable exactly when scripts expose time (Theorem 2), by a three-case exhaustion with a historical edge: pre-BIP65 Bitcoin could not support trustless occupation pricing. *Honest cost:* the deadline face (upfront fees) and the penalty face (bonds) meet the same deterrence bound and differ exactly in what they charge honest traffic (Proposition 1).

Multi-hop: an impossibility that explains six years. For payment chains we prove the *forwarding trilemma* (Theorem 4): per-hop enforceable deterrence, penalty-only honesty, and zero intermediary lockup are mutually incompatible; this is the 2020 conjecture, settled negatively. The escape routes of the trilemma are the forks the engineering record took (reputation abandoned pricing; Law's fees abandoned honesty; virtual channels pay locks). A matching construction, *chained bond forwarding* with decayed path bonds, attains the other face of the

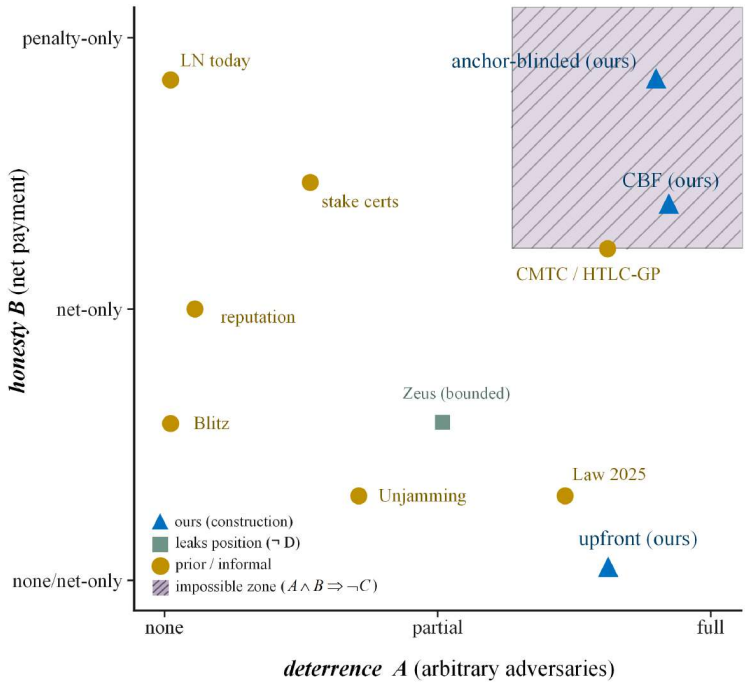

Fig. 1. The plot maps protocols across two axes: Deterrence (**A**) and Penalty-only Honesty (**B**). Every achievable combination of {**A** deterrence, **B** penalty-only honesty, **C** zero lockup, **D** path privacy} carries a price our theorems compute. Blue triangles represent our proposed constructions; gold circles denote prior or informal proposals; the teal square indicates protocols that leak path position (¬**D**); and the hatched purple band demarcates the impossible zone (**A** ∧ **B** ⇒ ¬**C**), where full deterrence, penalty-only honesty, and zero lockup cannot all hold.

characterization, with full per-hop deterrence and honest intermediaries at locks that provably cannot be improved: the lock *is* the price.

Privacy has a price schedule. Bond magnitudes leak position. Deterministic sink-resistant schedules leak exactly (Theorem 5(i)); anchor blinding restores complete per-payment privacy at bounded lockup while preserving sink-resistance exactly ((ii)); the sender's path length itself, however, leaks at rate $O(1/n)$ under repeated payments via extreme-value statistics, a side channel new to payment networks, removable by anchor reuse ((iii)).

A deployable construction. Fused-bond contracts (§5) execute the penalty face on Bitcoin today (taproot, no covenants, no oracles, no miner assumptions beyond default relay policy), with settlement extortion structurally eliminated (residual bounded by the slot-time price), zero net cost to honest payments, and slot-jamming deterrence improved by eight orders of magnitude.

Contributions. (1) a layered single-channel theory of deterrence (pricing, enforcement, and honest cost); (2) the forwarding trilemma and its matching construction (chained bond forwarding); (3) the privacy price schedule for sink-resistant bonds, including a new $O(1/n)$ extreme-value side channel and its repair; (4) fused-bond contracts deployable on Bitcoin today, with a complete nine-scenario walkthrough and two-floor economics (deterrence floor zero; compensation floor set by on-chain fees); (5) a classification of the proposed defenses onto the landscape (Table 1); (6) an evaluation grounding every quantity, from ρ calibration to network-wide attack costs.

Figure 1 summarizes this landscape.

2 Related Work

The two pillars. Safety and liveness as the classical pair goes back to Alpern and Schneider [2]; in the blockchain era, accountable safety/liveness for Casper [7] and the recent accountable-liveness line [15] state the two-pillar sentence explicitly. Layer-two surveys enumerate channel properties without pricing occupation. We are not aware of prior work treating *deterrence* (that harm carries a provable price) as a first-class security property of channels; "deterrence" appears in cryptography mainly as the capture-probability of covert MPC, an unrelated notion.

Jamming and its mitigations. Channel jamming was quantified by Mizrahi and Zohar [18] and surveyed in the jamming book [21]; Unjamming Lightning [22] proposed unconditional upfront fees plus local reputation (its famous "2%" is a simulation output, not a theorem, and its fee is value-proportional and hence, by Theorem 1(i), cannot deter slot jamming). HTLC-GP [17] introduced griefing-penalties with cumulative receiver-side bonds and a masking-based privacy argument; it assumes rational parties and an honest payee, and offers no characterization. Blitz [5] and successor multi-hop protocols mitigate griefing partially (payee-side griefing persists, as Zeus shows). Zeus [16] (preprint) proves that full griefing resistance for multi-relay lock-resolve payments is impossible (reduction to fair secret exchange) and constructs a bounded defense for *cost-sensitive* adversaries, at the price of payee on-chain deposits and fee-burning execution under miner-rationality. Our results are complementary: they characterize the *pricing structure* (necessary and sufficient), enforce penalty-only pricing without deposits or miner assumptions, and add the multi-hop trilemma and privacy layers; we compare on a common damage metric and note the adversary-class difference (Zeus bounds a k-cost-sensitive class; our bounds hold for arbitrary PPT adversaries). Informal engineering proposals converged on time-proportional pricing within eighteen months: Law's three-fee design [14] (March 2025: per-attempt upfront, delay-proportional hold fees, MAD burnable collateral), the onion-message fee-swallowing observation [12] (April 2026), and Riard's CMTC [20] (August 2026, self-described prolegomena). None came with theorems; Table 1 classifies these proposals and our constructions onto the landscape.

Table 1. Prior proposals and our constructions on the landscape {**A** deterrence, **B** penalty-only honesty, **C** zero lockup, **D** path privacy}. "net" = zero net payment; opportunity costs quantified separately.

Proposal	A	B	C	D
LN today (BOLT-2)	no ($2.1 \cdot 10^7$)	yes	yes	yes
Unjamming (2022)	bounded	no	yes	mostly
HTLC-GP (2020)	partial†	net yes	no	yes (masked)
Blitz (USENIX'21)	no	yes	improved	yes
Stake certs [19] (2020)	partial	yes	no	ZK only
Law three-fee (2025)	yes	no	no	conditional
Zeus (2025, preprint)	bounded‡	no	no	n/a
bolts#1280 reputation [13]	weak–mid	yes	yes	mostly
CMTC sketch (2026)	yes	net yes	no	claimed
Slot+value upfront (ours)	yes (0.0585)	no	yes	yes
CBF (ours)	yes	net yes	no	no
Anchor-blinded CBF (ours)	yes	net yes	no	graded

†honest-payee assumption; general-adversary λ unbounded (Zeus, Tab. 2 thereof).
‡for k-cost-sensitive adversaries only.

Pricing mechanisms. Our holding-cost lens continues a forty-year economics tradition: priority service [8] and efficient rationing [23] prove implementability iff monotonicity and the non-existence of a single anonymous tariff under stochastic demand, the economic mirror of our trilemma. In on-chain fee markets, rational harm-proofness [24] (the closest conceptual neighbor of penalty-only pricing) characterizes TFMs with a tetralemma; weighted-area pricing for parallel execution [1] prices time×resource within a single block (no adversary, no persistence, no iff); and multidimensional fee updates are essentially optimal [3] for worst-case adversarial demand. We bring this lens to off-chain shared state, where the enforcement substrate becomes theorem-relevant.

Channel economics and practice. Management-Science models of channel capital costs price value-time but not per-slot time; liquidity leases trade at an annualized anchor (α has a market price, β does not; our lower bound says the market is incomplete by necessity). Storage-rent systems price exactly $b(h)$'s shape, arriving empirically at the necessary direction of Theorem 1(i).

3 The Holding-Cost Model

3.1 Ledgers, Time, and Liveness

We work with a Bitcoin-style UTXO ledger $\mathcal{L}$ with delayed confirmations. The block height is the global clock; both absolute timelocks (`OP_CHECKLOCKTIMEVERIFY`, CLTV) and relative timelocks (`OP_CHECKSEQUENCEVERIFY`, CSV) are available in script. Scripts cannot constrain the spends of their own outputs: *no covenants* are available.[4]

[4] As of submission, CTV-style covenants remain unactivated; we treat their absence as a design constraint, not an assumption to be lifted.

Parties are probabilistic polynomial-time (PPT). Our liveness postulate is a *relay-policy* property, not a consensus rule, and is weaker than the miner-rationality assumptions used by fee-burning designs:

Definition 1 (Fee-bumped liveness). *Let t be a transaction authorized by an honest party P, together with a set $\{t_f\}$ of pre-signed fee-bumped variants constructed under TRUC/v3 package rules with anchor outputs (the default relay policy of Bitcoin Core $\geq$ 28.0). If P submits some member of $\{t_f\}$ at height h, then some member confirms by height $h + \delta_{\text{led}}$, except with negligible probability.*

This is justified by deployment: relay-level flooding attacks that flip HTLC outcomes on older clients (CVE-2024-55563 and the replacement-cycling family) are mitigated precisely by v3/TRUC, one-parent-one-child packages, and ephemeral anchors, which our enforcement paths use exclusively.

Time parameters. Contracts declare a *success deadline* T_s and a *refund deadline* $T_r = T_s + \delta$, both absolute heights. A *free-window margin* δ_{free} defines punctual off-chain resolution ($\leq T_s - \delta_{\text{free}}$); the maximum declared window is $\overline{T}$ (Lightning convention: $\approx$ one week). The minimum contract value is $v_{\min}$ (1 msat in LN).

3.2 Cryptographic Building Blocks

A two-party *adaptor signature* scheme $\mathsf{AS} = (\mathsf{KGen}, \mathsf{pSign}, \mathsf{pVrfy}, \mathsf{Adapt}, \mathsf{Extract})$ issues pre-signatures on a message m relative to a statement Y (a group element); any holder of a witness y with $Y = yG$ can adapt a pre-signature into a full signature, and any holder of both can extract y. We adopt the repaired definitional stack of *Foundations of Adaptor Signatures* [9] (Eurocrypt 2024): (AS.1) pre-signature existential unforgeability, (AS.2) pre-signature adaptability, and (AS.3) witness extractability. Schnorr adaptor signatures satisfy all three in the ROM, which we use for instantiation; hash-lock HTLCs are the degenerate, adaptor-free variant and suffice for our enforcement claims (§4). Throughout, Y_h denotes a payment-witness commitment with witness y_h.

3.3 Channels as Co-Signed Execution Domains

We instantiate the generalized-channel framework of Aumayr et al. [4] (Asiacrypt 2021): a channel is a funding output, a pair of revocable commitment transactions, and a state machine of sessions that add and remove off-chain contract outputs. Our axiom on the substrate is the following.

Axiom 1 (Execution domains) *An* execution domain *is a channel, or a shared-ledger domain (the parent ledger of a virtual channel, a factory's shared UTXO, an Ark round, a committee). Value enters an execution domain only through a transaction signed by one of its co-owners; no on-chain script can enforce a conditional transfer* across *domains.*

Axiom 1 is a structural consequence of UTXO semantics without covenants, and it yields an escape-route taxonomy that we reuse in §6: cross-domain conditionality is purchasable only by (i) *fronting* (an intermediary locks own capital: virtual channels and their compensation collaterals, where the Domino attack [6] is this cost made visible), (ii) *domain migration* (making transfers unconditional inside a larger domain, as within an Ark round), or (iii) *going on-chain* (splicing; under the current specifications, cross-channel atomic moves remain HTLC-only).

3.4 Held Shared-State Resources

To state our results at the level of the resource rather than the protocol, we use a small abstraction, in the lineage of Vickrey–Naor congestion tolls. A *bounded shared-state resource* has integer capacity μ (channel slots; batch queues and relay queues are further instances, §8). A *contract* $h = (v, T_s, T_r, Y_h)$, injected at height t_0 with declared window $\Delta = T_s - t_0$, occupies one unit of capacity and value v until resolution. The *holding-cost liability* of the injector is a function $b(h)$ of (v, Δ); our constructions use the bilinear family

$$b(h) \;=\; \alpha\, v\, \Delta \;+\; \beta\, \Delta, \qquad \text{with } \alpha\overline{T} \le 1, \tag{1}$$

where α prices value-time and β prices slot-time. Production storage-rent systems (Solana, RSK RSKIP-240, Ergo) charge exactly in this shape, and engineering arrived at the necessary direction of Theorem 1(i) before the theorem existed.

3.5 Security Definitions

Fix a channel between an *injector* A (the party adding contracts) and a *holder* B (the counterparty whose resource is occupied); either party may inject, and all definitions are symmetric.

Definition 2 (Occupancy and failure). *The* occupancy *of* h *is* $\mathrm{occ}(h) = \min(\mathrm{res}(h), T_r) - t_0$*, where* $\mathrm{res}(h)$ *is the resolution height. The* failure set F *contains* h *exactly when* h *is refunded at* T_r *(on-chain timeout or off-chain agreement to fail); success after* T_s *is unreachable, since the success branch expires at* T_s*.*

Definition 3 (Damage, cost, and the griefing factor). *For an execution, the* damage *and* adversary cost *are*

$$D \;=\; \sum_{h \in F} \big(r\, v_h + \rho\big)\, \mathrm{occ}(h), \tag{2}$$

$$C \;=\; \sum_{h} \big[\phi(h) + f_A(h)\big] \;+\; r \sum_{h} v_h\, \mathrm{occ}(h), \tag{3}$$

where both sums range over A-injected contracts (F for D; all for C), and r converts capital-time to satoshis (an annualized opportunity rate; we use the ACINQ figure of 1%), ρ is the slot-time value (sat/slot/day; calibrated in §7), φ(h) is the bond amount forfeited by A (whether confiscated by B or stranded), and f_A counts fees A itself pays to reclaim value. Both sums use the same occ *and F (*costing consistency*). The* griefing factor *is* $\lambda = \sup_{\mathcal{A}} D/C$*, and a protocol is* k-deterrent *iff* $\lambda \leq 1/k$*.*

Definition 4 (Balance safety). *No PPT adversary corrupting one party can, except with negligible probability, end an execution with the honest party's total on-chain takeaway below its initial funds plus legitimately earned payments plus confiscated bonds. Since bonds are funded exclusively by the injector, the honest party's balance is monotonically non-decreasing up to legitimate payments.*

Definition 5 (Penalty-only pricing). *Partition honest behavior into events: E_1 punctual success (off-chain, by $T_s - \delta_{\text{free}}$), E_2 agreed failure (off-chain), E_3 late success settled on-chain in $(T_s - \delta_{\text{free}}, T_s]$, and E_4 forced timeout settlement. A protocol is* penalty-only *iff the honest party's* net payment *is zero on E_1–E_3; E_4 occurs only when the counterparty is unavailable or non-cooperative at the deadline window (with standard availability, at the availability-failure rate).* Opportunity *costs of refundable locks are accounted separately (§7), not inside this definition.*

We state B in *net-payment* terms: it is the reading under which upfront fees (which transfer principal on every attempt) and penalty bonds (which lock refundable principal) separate exactly, which is the content of Proposition 1; opportunity costs of refundable locks are reported alongside but do not enter the definition.

Definition 6 (Lock-resolve contracts). *A channel-contract family is* lock-resolve *if accepting a contract binds the injector to a fee schedule fixed at injection; the contract occupies one slot and value v until resolution; resolution is by witness (success) or deadline (failure), possibly off-chain; and the injector's maximum cost for a contract resolved at any $\tau \leq T_r$ is a function $C_0(h)$ independent of both the realized holding time τ and the declared window Δ. The family is* hold-time-independent *if C_0 has this independence.*

All deployed and proposed designs we are aware of (LN's HTLCs and PTLCs, AMHL [4], Blitz, and the multi-hop protocol of [16]) are lock-resolve; LN today is hold-time-independent with $C_0 \approx 0$, and deadline-proportional upfront fees are lock-resolve but not hold-time-independent ($C_0 \propto \Delta$), a distinction Theorem 1 makes precise.

4 Single-Channel Characterization

This section establishes the theory of deterrence for a single channel in three layers: what pricing deterrence *requires* (Theorem 1), when penalty-only pricing is *enforceable* (Theorem 2), and what honesty *costs* under the two pricing faces (Proposition 1).

4.1 Layer 1: Time-Dependence Is Necessary

Theorem 1 (Pricing characterization). *Consider lock-resolve channel contracts over a path family with slot cap μ, minimum value $v_{\min}$, and maximum declared window $\overline{T}$.*

(i) Necessity. *For any hold-time-independent pricing C_0, there is an adversary strategy with griefing factor at least*

$$\lambda \;\geq\; \frac{\rho\,\overline{T}}{C_0 + r\,v_{\min}\overline{T}},$$

which diverges as $v_{\min} \to 0$ (and, when $C_0 > 0$, as $\overline{T} \to \infty$). In particular LN's current pricing ($C_0 \approx 0$) admits slot jamming with λ on the order of $\rho/(r\,v_{\min})$, which under our calibration exceeds 2×10^7.

(ii) Tight rate. *Conversely, if the injector's liability grows linearly in an effective time horizon h at rate $\alpha v h + \beta h$, where h is the declared window (deadline face) or the realized occupancy (penalty face), then $\lambda \leq \max(r/\alpha,\ \rho/\beta)$.*

Proof (Proof sketch). (i) The adversary fills all μ slots with minimum-value, doomed contracts declared for $\overline{T}$ and lets them time out; off-chain timeout resolution makes its marginal cost C_0 per slot plus capital lock-up $r\,v_{\min}\overline{T}$, while damage is $\mu\,(r\,v_{\min} + \rho)\,\overline{T}$. Dividing and letting $v_{\min} \to 0$ gives the claim; the accounting is the costing consistency of Definition 3. (ii) For each failed contract the numerator contributes $(rv + \rho)\,\mathrm{occ}$ and the denominator at least $(\alpha v + \beta)\,\mathrm{occ} + rv\,\mathrm{occ}$; the ratio of sums is bounded by the maximum component ratio (mediant inequality), giving $\max\big(r/(\alpha + r),\ \rho/\beta\big) \leq \max(r/\alpha, \rho/\beta)$.

Part (ii) covers both faces of time-dependent pricing: fees proportional to the *declared* window (the deadline face) and bonds forfeited proportional to *realized* occupancy (the penalty face). The two faces are not separated by the theorem (they achieve the same bound) and are separated exactly in honest cost, which is Proposition 1. Deadline-proportional fees fall outside the class in (i): their liability does depend on time, and they are covered by (ii).

4.2 Layer 2: When Penalty-Only Pricing Is Enforceable

Layer 1 forces time into the price. Layer 2 asks when the *penalty face* (honest users pay nothing, violators forfeit) can be enforced unilaterally, without trusted parties.

Theorem 2 (Enforcement characterization). *On a ledger with revocable commitments, margin-anchored penalty-only enforcement is achievable iff the script language exposes time (CLTV or CSV). Moreover the construction of §5 requires no covenants, no oracles, and no assumptions about miner behavior beyond Definition 1.*

Proof. ($\Leftarrow$) The fused-bond construction (§5) and Lemma 1 below. ($\Rightarrow$) Suppose scripts cannot branch on time, and let O be any output earmarked as a bond for a contract. Consider the paths returning or confiscating O; since no branch can be delayed, each is spendable immediately upon its witness:

(a) A can unilaterally return O. Then A reclaims the bond immediately after injection and holds the contract to timeout: the bond is vacuous and the griefing factor is unbounded (deterrence fails).
(b) B can unilaterally confiscate O. Then a malicious holder robs on-time contracts (violating Definition 5 for the honest-injector role).
(c) returning requires B's cooperation or confiscating requires A's. Then an unresponsive B holds the honest injector's bond hostage: penalty-only pricing is not unilaterally enforceable.

Witness-conditioned paths collapse into these cases: a return conditioned on the resolution witness y_h does not exist for failed contracts (a variant of (c)); a confiscation conditioned on y_h is unavailable to B for failed contracts (deterrence fails); pre-signatures granted at injection amount to case (a). No other path types exist.

Theorem 2 gives a historical watershed: Bitcoin before BIP65 (December 2015), with no script-level time visibility, could not have supported trustless occupation pricing at all. The margin-anchored insight behind our construction is that *deterrence needs only coarse, chain-anchored time at the deadline margin*: a contract either resolves punctually (little damage) or reaches the chain, where CLTV is a native witness. Fine-grained per-time-unit settlement is needed only cooperatively, where the parties can simply recompute. This eliminates the time oracles of informal designs (§2).

4.3 Layer 3: What Honesty Costs

Proposition 1 (Separation). *Fix a deterrence target $\lambda \leq \max(r/\alpha, \rho/\beta)$. The deadline face (unconditional fees on every attempt, including honest ones) charges every honest payment a strictly positive net amount; the penalty face (Definition 5) charges honest payments zero net on E_1–E_3. Under equal deterrence bounds, the two faces differ exactly in what they charge honest traffic.*

With our calibration the upfront face costs an honest 100,000-satoshi, one-day payment 110 sat per attempt (a 0.92% annual tax on channel capacity at typical volume), while the penalty face costs 0 net and locks a refundable 110 sat for the flight. In the same deterrence target, the choice between faces is a statement about who pays for honesty: the folklore "upfront fees" debate, resolved as a theorem.

5 Fused-Bond Contracts

We now instantiate the penalty face on Bitcoin today. The design goal is a single contract output whose own spending logic enforces both the payment and the bond, so that neither party can touch the other's share except through behavior.

Table 2. Fused output C_h: four taproot leaves (payouts fixed at injection).

Leaf	Condition	Fixed payout
SUCCESS	B's signature, witness y_h, CLTV$\leq T_s$	$v \to B$; $b(h) \to A$
SUCCESS-WIT.	(A *or* B) signature, y_h, CLTV$\leq T_s$	$v \to B$; $b(h) \to A$
REFUND	A's signature, CLTV$\geq T_r$	$v \to A$; $b(h) \to B$
SWEEP	B's signature, CLTV$\geq T_r+\delta$	$v + b(h) \to B$

Table 3. Deviation walkthrough: scenarios, on-chain adjudication, bond destination.

#	Scenario	Adjudication	Bond
W1	punctual success (E_1)	off-chain update	$\to A$
W2	agreed failure (E_2)	off-chain removal	$\to A$
W3	late success (E_3)	SUCCESS on-chain	$\to A$ (same tx)
W4	injector abandons	SWEEP at $T_r+\delta$	$v+b \to B$
W5	injector reclaims at T_r	REFUND, A pays its own fee	$\to B$
W6	holder extorts (has y_h, refuses)	no path to b without paying v; waiting ends in W5	$\leq \beta\overline{T}$ net
W7	holder offline, payment succeeds	SUCCESS-WITNESS	$\to A$
W8	stale commitment	revocation+justice	—
W9	low-bond injection	rejected by acceptance policy (§5.3)	—

5.1 The Fused Output

A contract $h = (v, T_s, T_r, Y_h)$ is injected by locking $v + b(h)$ in one taproot output C_h with four script-path leaves (Table 2; δ is the safety margin): Off-chain, C_h lives as a pending output of the commitment transaction, exactly like an LN HTLC output (Fig. 2); revocation and justice transactions handle stale states wholesale, so bonds add no new commitment machinery. On resolution the parties update state (removing C_h); on-chain paths are needed only at disputes and deadlines. SUCCESS-WITNESS lets either party finalize a *succeeded* contract when the witness has become public on-chain (e.g., revealed downstream) and the counterparty is offline; it cannot harm either side, since its payouts coincide with SUCCESS. SWEEP provides liveness when the injector abandons a failed contract: past $T_r + \delta$ the holder takes everything, which only strengthens deterrence.

5.2 Walkthrough

Table 3 adjudicates the nine deviation/environment scenarios; it is the case base of the safety proof.

Lemma 1 (Enforcement). *Under Definition 1: (i) if $h \in F$ (timeout unresolved), B obtains $b(h)$ by $T_r + \delta_{\text{led}}$; indeed, the only way A recovers v is a* REFUND *whose fixed outputs pay $b(h)$ to B (or B* SWEEP*s everything); (ii) if $h \notin F$, A's bond returns in full (off-chain update, or the fixed outputs of* SUCCESS*/* SUCCESS-WITNESS*).*

Lemma 2 (Extortion bound). *B can never obtain the bond while denying A the payment. Delaying a successful settlement nets B at most $b(h) - v \leq \beta\overline{T}$*

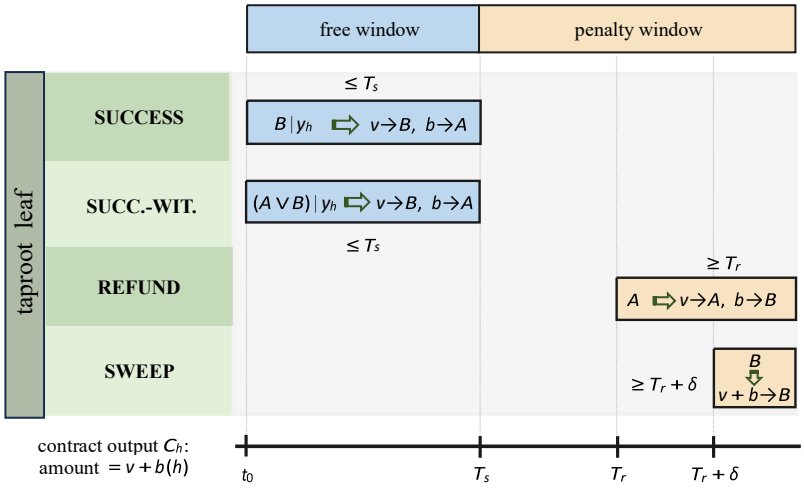

Fig. 2. The fused output C_h: both parties must, within the same transaction, respect each other's share. The bottom timeline shows when each path is available: beyond T_s, the contract enters the penalty window; after T_r, A may reclaim the principal v via the REFUND path; and after $T_r + \delta$, B may sweep all remaining funds via the SWEEP path as a final resolution. This last path has no expiry: it remains available indefinitely.

(using $\alpha\overline{T} \leq 1$), i.e. the slot-time price; for realistic values the bound is negative (B forfeits v to gain $b < v$).

Theorem 3 (Safety and deterrence). *Π_{fused} satisfies Definition 4 under DLP and AS.1 (with hash-lock instantiation, under discrete logarithms), and Definition 3 with $\lambda \leq \max(r/\alpha, \rho/\beta)$.*

Proof (Proof sketch). Safety: a stale commitment is punished by revocation (inherited from the generalized-channel ledger layer); any other balance gain requires a valid SUCCESS/SUCCESS-WITNESS/SWEEP spend, i.e., the witness y_h before T_s or a timelock at $T_r+\delta$; the former yields a pre-signature forgery (AS.1), the latter contradicts the ledger. Deterrence: for $h \in F$, Lemma 1 gives $\phi(h) = b(h) \geq (\alpha v + \beta)\,\mathrm{occ}(h)$ (full forfeiture exceeds proration); sum via costing consistency and apply the mediant bound of Theorem 1(ii).

5.3 Two Floors and the Acceptance Threshold

Binary forfeiture makes *deterrence* robust for any $\beta > 0$: even a stranded bond (nobody sweeps it) is a certain loss to the injector. *Compensation* of the holder, however, has a floor: sweeping one contract costs ≈ 93 vB (script-path input), so rational holders sweep only when $b(h)$ exceeds $93 \cdot$ fee (at 10 sat/vB: $\beta \geq 931$ for one-day windows, $\beta \geq 133$ for one-week). Holders therefore run an *acceptance threshold*: reject injections with $b(h)$ below the sweep cost. This closes the short-window bypass (attackers cannot inject sub-threshold bonds) and yields two

service tiers (strict: compensation guaranteed; lenient: deterrence only, β free). The fused design trades pooled amortization (a pooled bond output would cut the compensation floor by the number of contracts) for the structural elimination of settlement extortion (Lemma 2) and zero new commitment machinery; per-path on-chain costs are tabulated in §7.

6 Multi-Hop: Trilemma, Forwarding, and Privacy

A payment traverses a chain $P_0 \to P_1 \to \cdots \to P_L$ of channels. The injector at channel j is P_{j-1}; the holder is P_j. We ask which combinations of four desiderata are achievable:

- **A** *Deterrence (sink-resistant):* for every channel j, whenever damage D_j occurs there, an on-chain-enforceable mechanism transfers $\geq k\, D_j$ of value to P_j out of escrow inside channel j, and this survives the coalition $\{P_0, P_L\}$ (attacker plus payee).
- **B** *Penalty-only honesty:* every honest party (including intermediaries) pays zero net (Definition 5).
- **C** *Zero intermediary lockup:* intermediaries escrow no capital beyond in-flight forwarded value.
- **D** *Path privacy:* the Sphinx baseline, under which intermediaries learn nothing beyond their adjacent hops, in particular neither their position i nor the path length L [11].

6.1 The Forwarding Trilemma

Lemma 3 (Compensation locality). *Value enters an execution domain only through its co-owners (Axiom 1); hence any on-chain-enforceable compensation inside channel j is funded by escrow posted there by P_{j-1} or P_j. P_0's funds reach channel $j \geq 2$ only via $P_1, \ldots, P_{j-1}$, each hop of which is that intermediary's capital lock.*

Theorem 4 (Forwarding trilemma). *For $L \geq 2$, no lock-resolve forwarding protocol satisfies $\mathbf{A} \wedge \mathbf{B} \wedge \mathbf{C}$.*

Proof. Run the sink attack: P_0 and P_L collude, the payment is doomed, all contracts time out. Every intermediary behaves honestly (faithful forwarding; cooperative timeout resolution). Consider channel 2 and the escrow E backing P_2's compensation (it exists by **A** and Lemma 3). If E is refundable upon P_1's honest behavior, it is refunded here, and P_2 is left uncompensated, so **A** fails. If E is unconditional, the honest intermediary P_1 pays for an upstream-caused failure, contradicting **B**. If E is funded by pass-through of P_0's escrow, then since P_0's forfeiture lands at $T_r^{(1)} \geq T_r^{(2)} + \delta$ while E must be posted at injection time $t_0 < T_r^{(2)}$, party P_1 fronts own capital over $[t_0, T_r^{(1)}]$, violating **C** (the lock exists even if later reimbursed). Reputational or third-party funding satisfies no on-chain-enforceable $\geq kD_j$ transfer and is outside **A**; splicing-in funds still requires a co-owner signature (Lemma 3).

Theorem 4 settles, negatively, the conjecture voiced by Joost Jager on the Lightning mailing list in October 2020 [10], that a trustless compensation scheme "may even be theoretically impossible ... a proof of that would also be useful information". It explains the engineering record: reputation systems abandoned **A**; Law-style dual fees abandoned **B**; virtual channels pay **C** (the Domino attack is that price made visible). Read marginally, the trilemma *prices* deterrence: today's network has zero marginal lockup and zero deterrence; any protocol purchasing k-deterrence with penalty-only honesty must lock b per in-flight forward, and the theorem says this marginal cannot be zero.

6.2 Chained Bond Forwarding

The matching construction runs fused contracts per hop with two rules. *Deadline staggering:* $T_r^{(j)} = T_r^{(j+1)} + \delta$, so downstream failures resolve first and inform upstream settlements (the standard LN direction). *Decayed path bonds:*

$$b_j \;=\; k\sum_{i\geq j} D_i \;=\; \left(\alpha v\overline{T} + \beta\overline{T}\right)(L-j+1), \tag{4}$$

the second form showing the schedule needs no damage oracle (remaining hops $\times$ declared-window maximum). Uniform bonds fail: P_0 forfeits b to P_1 and the colluding sink recovers it, net cost 0; decay truncates the recapture at $b_L = kD_L$, exactly the sink's local damage, so the coalition pays $k\sum_{i<L} D_i$. Each honest intermediary nets $b_j - b_{j+1} = +kD_j$: fully compensated, and coalition self-dealing loses in aggregate (no arbitrage). Lockup per intermediary is b_{j+1} for the flight, exactly the amount that backs channel $j{+}1$'s deterrence at all: *the lock is the price.*

6.3 Privacy Has a Price Schedule

Decayed bonds leak: an intermediary observing b_j and knowing the family reads off its remaining-hop count (Fig. 3(a)). Three layers make this precise.

Theorem 5 (Deterrence–privacy).

(i) *Any* deterministic *sink-resistant schedule over a path family with two distinct lengths* leaks at the anchors*: the first-hop liability obeys* $b_1 \geq k\sum_{i<L} D_i + b_L$*, so its magnitude support shifts with* L *and a first-hop observer distinguishes path lengths, exactly for explicit linear families and by support separation in general. Interior positions* may *hide (flat-middle schedules are sink-resistant and hide their middle hops), but anchor positions remain distinguishable and the first hop learns* L*; since the Sphinx baseline grants intermediaries nothing beyond adjacency, this already violates* **D**.

(ii) Anchor-blinding: *let the sender draw* $j \leftarrow [0, \overline{L} - L]$ *uniformly per payment and set* $b_i = k(\overline{L} - j - i + 1)D$*. The coalition net* $b_1 - b_L = k(L{-}1)D$ *is* independent of j*: sink-resistance holds exactly. Per-payment position privacy*

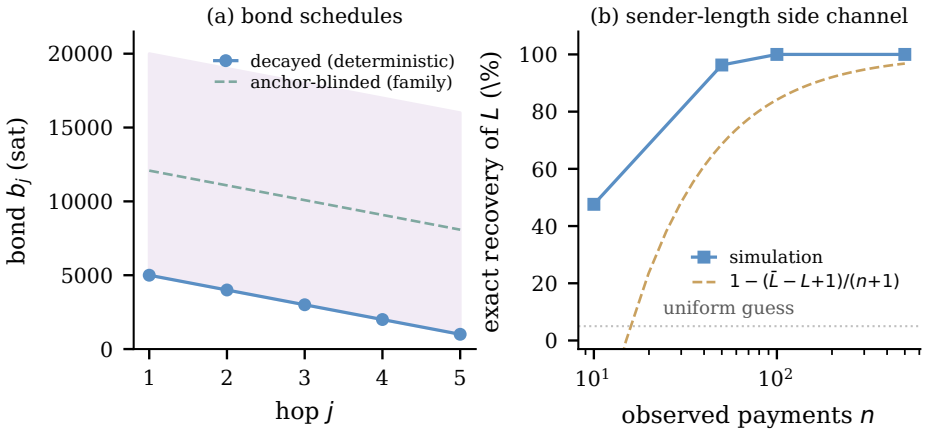

Fig. 3. (a) Path-bond schedules: decayed bonds encode position in their magnitudes; anchor blinding randomizes the start, hiding position per payment while the drop, and hence sink-resistance, is unchanged. The shaded band is the envelope of the blinded family across all anchor draws: its lower edge is the deterministic (unblinded) schedule, and every member has the same per-hop drop. (b) The sender-length side channel: exact recovery of L from minimum statistics over n payments.

is complete *(the posterior equals the prior; leakage zero bits) at lockup inflation* $(\overline{L}+L)/2L$*. However, the* sender's *path length leaks at rate* $O(1/n)$ *under repeated payments: an adversary adjacent to the sender observes* $m_1 = \overline{L} - j + 1$ *per payment and minimum statistics recover* L *(empirically* 96.3% *exact recovery at* n=50*; the bias* $(\overline{L} - L{+}1)/(n{+}1)$ *matches theory). This side channel is new to payment networks and is methodologically kin to multicast delay tomography and Tor latency structure attacks.*

(iii) Fixed anchor: *drawing* j *once per sender (bounded by its longest path) makes* m *constant across payments, with no extreme-value leak, while position remains hidden in the unknown* j*; inflation is at most* $\overline{L}/L_{\text{typ}}$.

We credit HTLC-GP (TrustCom 2020) as the first to treat bond-magnitude privacy: its $\varphi_0 t_0$ masking term is a per-payment secret additive offset, the same family as anchor blinding, and our Theorem 5(ii)/(iii) systematizes it with a sink-resistance-preservation proof and the leak hierarchy (including the $O(1/n)$ residual its masking does not address).

Together with the trilemma, the landscape of Figure 1 is complete: $\{\mathbf{A}, \mathbf{B}, \mathbf{C}\}$ is impossible; each achievable combination is realized by a matching construction, namely today's channels ($\neg\mathbf{A}$), slot-and-time upfront fees ($\neg\mathbf{B}$), chained bond forwarding ($\neg\mathbf{C}, \neg\mathbf{D}$), and anchor-blinded forwarding ($\neg\mathbf{C}$, graded $\mathbf{D}$). Finally, $\mathbf{A} \wedge \mathbf{D}$ carries the graded price schedule of Theorem 5 (Fig. 4).

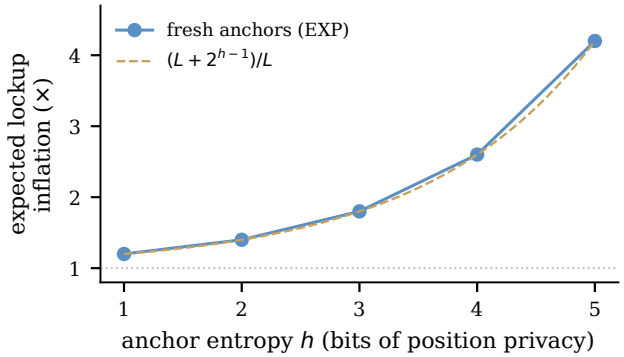

Fig. 4. The privacy price schedule: expected lockup inflation as a function of anchor entropy (Theorem 5(ii)).

Table 4. Headline numbers (v3 channels, calibrated ρ).

	LN today	Upfront (slot+value)	Fused bonds
Slot-jam $\lambda = D/C$	2.1×10^7	0.0585	0.058
Honest cost / payment	0	110 sat	0 net
Lockup (% capacity)	0	0	0.01 to 0.1 (refundable)
Network jam / month	≈ 0	10.4 BTC (fees)	10.4 BTC forfeited

7 Applications and Evaluation

Onion messages. On stateless relays, unconditional fees are swallowable: a relay that charges per message and drops it profits by the fee at zero cost, a deviation that does not exist in payment channels, where future success fees discipline the relay. We formalize this practitioner observation: under upfront pricing on stateless onion forwarding, fee-swallowing is *strictly* profitable; penalty-only pricing is immune wherever occupation is script-visible, which for onion messages requires first paying the statefulness price (measured in practice at a $0.5 \to 1.5$ round-trip overhead). This is the enforcement layer's second instance: the substrate determines what pricing can be enforced, and theory predicts the practitioners' intuition that payments and messages "are different domains."

Calibration. Unless stated: $r = 1\%$/yr (ACINQ), slots $\mu = 114$ (v3), $v_{\min} = 1$ msat, windows up to 7 days, average channel 12M sat, enforcement cost 93 vB script-path input at 10 sat/vB. Slot-time value ρ uses the conservative *router-damage* caliber: node revenue ($\approx$30k sat/month over $\approx$15 channels) divided by slot-days, $\rho \approx 0.585$ sat/slot/day; the network-externality caliber (throughput blocked per slot-day, $\approx$4,498 sat) is 7,700$\times$ larger and reported as an upper bracket. All numbers are reproducible from the released experiment scripts.

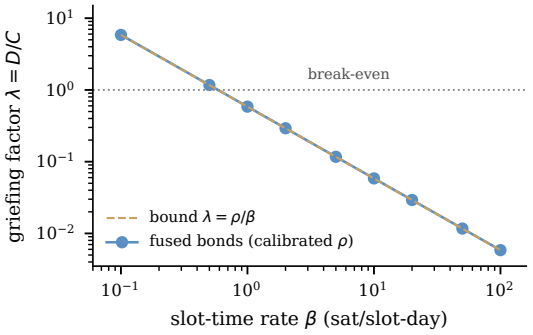

Fig. 5. Slot-jam griefing factor under fused bonds vs. the slot-time rate β (calibrated ρ; log–log). Dashed: theoretical bound $\lambda = \rho/\beta$; below the dotted break-even line ($\lambda = 1$) an attacker pays more than the damage. LN today and value-only upfront fees lie far above the plotted range (Table 4).

Deterrence. Table 4 and Fig. 5: the two time-proportional faces reach the theoretical bound $\max(r/\alpha, \rho/\beta)$ from opposite honesty sides; value-proportional upfront fees (the folklore "2%") fail on the slot axis ($\lambda > 10^5$), as Theorem 1(i) predicts. Under fused bonds, jamming the measured network (mempool.space snapshot, 2026-08-30, 32,674 channels) for a month forfeits ≈ 10.4 BTC.

Honest cost and lockup. With availability failures at 0.1–5%, expected honest cost of fused bonds is 0.11–5.5 sat per payment (vs. 110–2,000 for the upfront faces); in-flight bond lockup is 0.009–0.104% of channel capacity, refundable, versus a 0.9–16.7% annual capacity tax for the upfront faces.

On-chain footprint. Per-contract dispute paths: SUCCESS 199.5 vB, REFUND 191 vB, SWEEP 148 vB. The compensation floor ($\beta \geq 93 \cdot \text{fee}/\Delta$) is the acceptance threshold of §5.3: 931/133 sat per slot-day at 10/1 sat/vB for one-day/one-week windows.

Adversarial fee markets. Pumping fees to push the acceptance threshold above honest bonds for six hours at 10 sat/vB costs ≈ 3.6 BTC of block-space purchases, four to six orders above any single-channel bond pool; fee-market griefing of the threshold is priced out.

Multi-hop and privacy. Chained bond forwarding (Eq. (4)) with $k = 10$, uniform per-hop damage: bonds decay $[5, 4, 3, 2, 1] \times kD$; every intermediary nets $+kD$ (attacks become the best-paying traffic, with no self-dealing arbitrage); the sink coalition pays $k(L-1)D$. Anchor blinding ($\overline{L} = 20$, $L = 5$) buys 4 bits of anchor entropy at 2.5× expected lockup inflation; the sender-length side channel reaches 96.3% exact recovery after 50 payments (100% after 100), matching the $(\overline{L} - L+1)/(n+1)$ bias law; the fixed-anchor variant removes it.

8 Discussion and Limitations

Scope discipline. All our statements are pinned to the lock-resolve class, PPT adversaries, and the costing of Definition 3; the trilemma holds under Axiom 1 (execution domains), whose escape routes we make explicit rather than assume away.

Enforceability landscape. The enforcement layer is substrate-dependent by content, not by accident: Bitcoin (CLTV/CSV, no covenants) admits fused bonds; account-based platforms read time natively but at gas prices; Solana's rent prices state size, not slot-time; rollups inherit sequencer latency constraints. A systematic comparison is future work; account-ledger non-custodial encumbrance results suggest the substrate boundary is real.

Deployment path. Injecting a fused contract maps to a `update_add_htlc` carrying two TLV extensions (the bond amount and schedule family); the four leaves are standard tapscript; watchtowers verify the same commitment surface as today, with Sweep giving them a positive-fee enforcement duty. Backward compatibility is opt-in per channel, exactly as v3 anchoring rolled out.

Limitations. (i) E_4: a party that is honest but has an unavailable counterparty at the deadline window forfeits the bond (the availability semantics of Definition 5); watchtowers reduce this to negligible in practice, and the expected cost (0.11–5.5 sat/payment at 0.1–5% unavailability) is three orders below upfront alternatives. (ii) Heterogeneous β across defenders changes attack routing incentives; we treat uniform parameters. (iii) Our privacy analysis covers bond-magnitude channels; composition with blinded paths is open. (iv) Dynamic pricing equilibria (competitive (α, β) under adversarial traffic) are beyond this paper; the non-existence of pure equilibria in hub pricing suggests richness. (v) Layer-1 generalization: the pricing and trilemma layers transport to any priced holdable resource; the enforcement layer is where substrates differ.

Outlook. The impossibility-and-characterization tradition this paper joins is recurring across domains: priority pricing in power systems, harm-proofness in fee markets, and, this year, latency pricing for LLM inference. In GPUs the argument runs *against* time pricing (the consumed resource is bandwidth); in channels it runs *for* it (the consumed resource is slot-time): one principle, opposite signs by resource structure.

9 Conclusion

We proposed deterrence as the third pillar of payment channel security and gave its theory: time-dependence is necessary and linear liability tight; penalty-only enforcement is possible exactly where scripts see time; the forwarding trilemma settles a 2020 conjecture, and its matching construction shows the lockup cannot

be improved; privacy carries an explicit price schedule, including a new extreme-value side channel; and the fused-bond construction runs on Bitcoin today, improving jamming deterrence by eight orders of magnitude at zero net cost to honest users. The engineering community converged on time-proportional pricing empirically; this paper gives the underlying theory.

Appendix A: Full Proof of Theorem 1 (Pricing)

(i) The adversary injects μ contracts of value $v_{\min}$ declared for $\overline{T}$ with hash images it never reveals, and resolves them off-chain at timeout (zero fees). Cost: $C = \mu(C_0 + r v_{\min}\overline{T})$; damage: $D = \mu(r v_{\min} + \rho)\overline{T}$. Hence $\lambda \geq \rho\overline{T}/(C_0 + r v_{\min}\overline{T})$, unbounded as $v_{\min} \to 0$ and, for $C_0 > 0$, as $\overline{T} \to \infty$. *(ii)* For each $h \in F$, $\phi(h) = b(h) \geq (\alpha v_h + \beta)\,\mathrm{occ}(h)$ by Lemma 1(i) (binary forfeiture dominates proration), and Definition 3 adds $r v_h \mathrm{occ}(h)$ to C for all injected h. Summing and applying the mediant inequality $\frac{x_1+x_2}{y_1+y_2} \leq \max(x_1/y_1, x_2/y_2)$ componentwise gives $\lambda \leq \max\big(r/(\alpha+r), \rho/\beta\big) \leq \max(r/\alpha, \rho/\beta)$. □

Appendix B: Full Proof of Theorem 2 (Enforcement)

($\Leftarrow$) Lemmas 1 and 2 with Table 3 verify Definition 5 event by event: E_1, E_2 settle off-chain (bond never leaves A's balance); E_3 spends SUCCESS/SUCCESS-WITNESS whose fixed outputs return b to A within the same transaction; E_4 is the penalty event. ($\Rightarrow$) As in the text: without time visibility no script can be delayed, so the bond output's paths are partitioned into (a) A-unilateral return (vacuity: inject, self-return, hold to timeout; cost 0, damage positive, λ unbounded); (b) B-unilateral confiscation (robbery of punctual contracts, violating Definition 5 for the injector role); (c) bilateral paths (hostage by an unresponsive holder). Witness-conditioned returns require y_h, absent for failed contracts (case (c) variant); witness-conditioned confiscation is unavailable for failed contracts (deterrence fails); pre-signed returns granted at injection are A-unilateral in effect (case (a)). Mixed witness/time conditions are impossible by assumption. □

Appendix C: Full Proof of Theorem 4 (Trilemma)

Lemma 3 first: an execution domain's outputs are created only by transactions its co-owners sign (funding, splices); absent covenants, no script constrains cross-domain spends; hence enforceable value inside channel j originates from P_{j-1} or P_j escrow, and P_0's value reaches $j \geq 2$ only through intermediary action. Theorem: run the sink attack (P_0, P_L collude; doomed payment; all timeouts; intermediaries honest). On channel 2, **A** requires enforceable escrow $E \geq kD_2$ inside the channel (Lemma 3); consider its funder. (a) P_1's own funds, refundable on honest behavior: refunded (the attack's failure is not P_1's fault, and P_1 met every honest-behavior condition). Contradiction with **A**. (b) P_1's own funds,

unconditional on damage: honest intermediary pays for upstream-caused failure, contradicting **B**. (c) pass-through of P_0's escrow: P_0's forfeiture realizes at $T_r^{(1)} \geq T_r^{(2)} + \delta$, but E must exist from injection $t_0 < T_r^{(2)}$, so P_1 fronts capital over $[t_0, T_r^{(1)}]$. Contradiction with **C**. Reputation/third-party funding provides no enforceable transfer (outside **A**); on-chain splicing into channel 2 requires a co-owner signature (Lemma 3), reducing to (a)–(c). □

Appendix D: Full Proof of Theorem 5 (Privacy)

(i) Sink-resistance requires $b_1 - b_L \geq k \sum_{i<L} D_i$ for every path in the family, hence $b_1 \geq k \sum_{i<L} D_i + b_L$: the first-hop magnitude is bounded below by a quantity strictly increasing in L over any family with two distinct lengths. Consequently the supports of b_1 under different lengths are separated (deterministic schedule; explicit linear families invert as $i = \overline{L} + 1 - m/(kD)$), and a first-hop observer recovers L exactly. Interior hiding does not rescue **D**: a flat-middle construction (large b_1, constant middle, small b_L) is sink-resistant and hides middle indices, yet its first-hop magnitudes separate perfectly across lengths and its anchors are distinguishable from its middle, verified numerically, including noisy-anchor inference at 91.6–100% top-1 accuracy against a 25% chance baseline. (ii) With $j \leftarrow [0, \overline{L} - L]$ uniform and fresh per payment, the observed $m = \overline{L} - j - i + 1$ is uniform on a range independent of (i, L) given feasibility, so the posterior equals the prior (zero leakage); the coalition net $b_1 - b_L = k(L-1)D$ is independent of j. For the sender-length channel: $m_1 = \overline{L} - j + 1$ is uniform on $[L+1, \overline{L}+1]$; the estimator $\hat{L} = \min_n m_1 - 1$ has bias $(\overline{L} - L+1)/(n+1)$ and its exact-recovery probability converges as simulated (96.3% at n=50). (iii) With j fixed, m is constant across the sender's payments: no order statistics available; the per-payment posterior argument of (ii) is unaffected since j remains unknown to observers. □

Appendix E: Fee-Swallowing Proposition (Onion Messages)

Under unconditional per-message fees on stateless relays, the swallow deviation (collect fee, drop message) yields payoff equal to the fee at zero cost and zero future consequence (statelessness), hence is strictly profitable whenever fees are positive. In payment channels the same deviation forgoes the future success-fee revenue that conditions the relay's compensation on delivery, which is the discipline practitioners refer to; penalty-only pricing restores it wherever holding is script-visible. □

References

1. Acilan, B., Constantinescu, A., Heimbach, L., Wattenhofer, R.: Transaction fee market design for parallel execution. In: AFT 2025 (2025). https://doi.org/10.4230/LIPIcs.AFT.2025.23, lIPIcs 354, 23:1–23:25

2. Alpern, B., Schneider, F.B.: Defining liveness. Information Processing Letters **21**(4), 179–185 (1985). https://doi.org/10.1016/0020-0190(85)90056-0
3. Angeris, G., Diamandis, T., Moallemi, C.C.: Multidimensional blockchain fees are (essentially) optimal. In: AFT 2025 (2025). https://doi.org/10.4230/LIPIcs.AFT.2025.24, lIPIcs 354, 24:1–24:23
4. Aumayr, L., Ersoy, O., Erwig, A., Faust, S., Hostakova, K., Maffei, M., Moreno-Sanchez, P., Riahi, S.: Generalized channels from limited blockchain scripts and adaptor signatures. In: ASIACRYPT 2021 (2021). https://doi.org/10.1007/978-3-030-92075-3_22, ePrint 2020/476
5. Aumayr, L., Moreno-Sanchez, P., Kate, A., Maffei, M.: Blitz: Secure multi-hop payments without two-phase commits. In: USENIX Security 2021 (2021), ePrint 2021/176
6. Aumayr, L., Moreno-Sanchez, P., Kate, A., Maffei, M.: Breaking and fixing virtual channels: Domino attack and donner. In: NDSS 2023 (2023). https://doi.org/10.14722/ndss.2023.24370, ePrint 2021/855
7. Buterin, V., Griffith, D.: Casper the friendly finality gadget. arXiv:1710.09437 (2017)
8. Chao, H.P., Wilson, R.: Priority service: Pricing, investment, and market organization. American Economic Review **77**(5), 899–916 (1987)
9. Gerhart, P., Schröder, D., Soni, P., Thyagarajan, S.A.: Foundations of adaptor signatures. In: EUROCRYPT 2024 (2024). https://doi.org/10.1007/978-3-031-58723-8_6, ePrint 2024/1809
10. Jager, J.: Hold fees: 402 payment required. lightning-dev mailing list, Oct 2020 (2020), archived: https://diyhpl.us/~bryan/irc/bitcoin/bitcoin-dev/linuxfoundation-pipermail/lightning-dev/2020-October/002849.txt; "Maybe it is even theoretically impossible... (A proof of that would also be useful information to have.)"
11. Kappos, G., Yousaf, H., Piotrowska, A., Kanjalkar, S., Delgado-Segura, S., Miller, A., Meiklejohn, S.: An empirical analysis of privacy in the lightning network. In: FC 2021. pp. 167–186 (2021). https://doi.org/10.1007/978-3-662-64322-8_8
12. Kirk-Cohen, C.: Onion message jamming in the lightning network. Comment on Delving Bitcoin thread (opened by E. Cestari), April 2026, https://delvingbitcoin.org/t/onion-message-jamming-in-the-lightning-network/2414/9
13. Kirk-Cohen, C., Shikhelman, C.: Outgoing reputation and htlc accountability. lightning/bolts #1280 (2025), https://github.com/lightning/bolts/pull/1280
14. Law, J.: Fee-based spam prevention for lightning. Delving Bitcoin, March 2025, https://delvingbitcoin.org/t/fee-based-spam-prevention-for-lightning/1524, covered in Bitcoin Optech Newsletter #347
15. Lewis-Pye, A., Neu, J., Roughgarden, T., Zanolini, L.: Accountable liveness. In: CCS 2025 (2025). https://doi.org/10.1145/3719027.3765032, ePrint 2025/693
16. Liu, J., Xue, Y., Wu, D., Liu, J., Wang, X.: Zeus: Defending against fee stealing and griefing attacks in multi-hop payments. Cryptology ePrint Archive, Paper 2025/1070 (2025), https://eprint.iacr.org/2025/1070, preprint; last revised 2025-10-21
17. Mazumdar, S., Banerjee, P., Ruj, S.: Time is money: Countering griefing attack in lightning network. In: 2020 IEEE 19th International Conference on Trust, Security and Privacy in Computing and Communications (TrustCom) (2020). https://doi.org/10.1109/TrustCom50675.2020.00138, extended version: arXiv:2005.09327
18. Mizrahi, A., Zohar, A.: Congestion attacks in payment channel networks. In: FC 2021. pp. 170–188 (2021). https://doi.org/10.1007/978-3-662-64331-0_9

19. Naumenko, G., Riard, A.: Mitigating channel jamming with stake certificates. lightning-dev mailing list, Nov 2020, https://www.mail-archive.com/lightning-dev@lists.linuxfoundation.org/msg02138.html, covered in Bitcoin Optech Newsletter #126
20. Riard, A.: Conditional message transfer contract to solve jamming. Delving Bitcoin, August 2026, https://delvingbitcoin.org/t/conditional-message-transfer-contract-to-solve-jamming/2772, self-described "prolegomenas for a formalization"; covered in Bitcoin Optech Newsletter #418
21. Riard, A., Naumenko, G.: Advances in channel jamming research (the jamming book). https://jamming-dev.github.io/book/ (2022)
22. Shikhelman, C., Tikhomirov, S.: Unjamming lightning: A systematic approach. Cryptology ePrint Archive, Paper 2022/1454 (2022), https://eprint.iacr.org/2022/1454
23. Wilson, R.: Efficient and competitive rationing. Econometrica **57**(1), 1–40 (1989). https://doi.org/10.2307/1912571
24. Zhang, F., Park, E., Wu, K.: Beyond incentive compatibility: Rational harm-proof transaction fee mechanisms. Cryptology ePrint Archive, Paper 2025/2247 (2025), https://eprint.iacr.org/2025/2247