# Authority Without Priority: The Revocation Trilemma and the Price of Stablecoin Freezes

Di Wu[1,2,†] Yuman Bai[1] Shoupeng Ren[1] Xinyu Zhang[3] Yiyue Cao[2]
Xuechao Wang[2,*] Wu Wen[1] Jian Liu[1]

[1] Zhejiang University [2] The Hong Kong University of Science and Technology (Guangzhou) [3] Sun Yat-Sen University

wu.di@zju.edu.cn, baiyuman@zju.edu.cn, spren@zju.edu.cn
zhangxy353@mail2.sysu.edu.cn, ycao948@connect.hkust-gz.edu.cn
xuechaowang@hkust-gz.edu.cn, wu.wen@intl.zju.edu.cn
liujian2411@zju.edu.cn

***Abstract*—Stablecoin contracts let issuers freeze an address, but a freeze is an ordinary transaction that block producers order like any other, so a transfer submitted after it can still execute first. We define the missing guarantee as *revocation priority*: once a restriction is submitted, no prohibited operation submitted later takes effect before it. We prove a trilemma: with revenue-maximizing block producers, no blockchain platform can provide revocation priority together with immediacy, which makes a transfer's effects usable upon inclusion, and neutrality, which rules out preferential inclusion or ordering for any sender. Ethereum keeps immediacy and neutrality and gives up revocation priority. Using USDT and USDC freezes on Ethereum from 2017 to 2026, we examine how issuers and their targets compete for execution order. In one case, a freeze had long been pending when a transfer submitted after it executed first. The competition has escalated from fee bidding to private submission to execution delegated through intent-based orders. Overall, the issuer pays more for order than its target. These patterns follow from the two sides' incentives, which we model as a Tullock contest. In equilibrium the issuer outbids its target but cannot buy certainty. Stricter freeze obligations thus strengthen its incentive to integrate into ordering infrastructure, at the expense of neutrality.**

## 1. Introduction

USDT and USDC grant their issuers freeze authority, and the GENIUS Act requires permitted issuers to have the technical capability to comply with lawful freeze orders [3]. Yet in the approximately $6.3 billion Binance theft on October 29, 2026, the attacker converted over $700 million in stablecoins into assets beyond issuers' freeze authority [26], despite Tether and Circle freezing approximately $2.6 billion [31]. A faster off-chain response lets an issuer submit a freeze sooner, but a question remains about execution. Once the freeze is submitted, can the issuer guarantee that it takes effect before any subsequent transfer attempt succeeds?

When a single component controls both access checks and request ordering, it can apply a revocation before processing subsequent requests [6], [27]. On Ethereum, these responsibilities are separated: contracts determine who may freeze an address, while block producers decide which transactions to include and in what order. Consequently, even if an issuer submits a freeze first, a subsequently submitted transfer can still take effect ahead of it.

Prior work on transaction ordering has extensively studied frontrunning and maximal extractable value (MEV) [13], [25], [30], alongside defenses that conceal pending transactions [8], [10] or order them by arrival time [18], [19], [39]. However, restrictive operations such as stablecoin freezes, system pauses, and permission revocations require a stronger guarantee than frontrunning protection, which we call *revocation priority*: once a restriction is issued, no subsequently submitted operation that it restricts may take effect before it. A restricted party can act without observing a pending restriction; neither concealing pending transactions nor ordering them by local arrival time restores this guarantee. This leads to our central question: **under what conditions can a blockchain guarantee revocation priority, and when it cannot, what happens in practice and what equilibrium emerges between the privileged and restricted parties in the ordering market?**

Because Ethereum thus leaves room for violating $P$, we measure USDT and USDC freezes in practice, analyzing 3,561 restricted address–token pairs, their fund flows, and public mempool records on Ethereum from November 2017 through November 2026. Although freezes and transfers are rarely pending concurrently in the public mempool, a later-appearing transfer can still execute first: a freeze bidding below the prevailing base fee had been pending for 22.6 minutes when a transfer appeared and was included 8 seconds later, hours before the freeze took effect. Across the observation period, submission and execution techniques in the ordering market escalated: from fee adjustments, to

∗. Xuechao Wang is the corresponding author.

†. This work was carried out jointly at Zhejiang University and The Hong Kong University of Science and Technology (Guangzhou).

private submission bypassing the public mempool, to intent-based orders delegating execution to specialized executors, through which 325 million USDT was sold during the Binance incident with 1,751.20 ETH paid to block producers.

When both sides pay for execution order, can an issuer buy reliable priority within the protocol? We model the two parties' bidding for candidate block selection as a Tullock contest. At equilibrium, whenever the restricted party bids, in-protocol fees cannot reduce the probability that its transfer takes effect first to zero. If a failed freeze is costly enough, the issuer stands to gain from securing priority outside the protocol, putting pressure on the blockchain platform's neutrality.

### 1.1. Our Contributions.

- **Theory (Sections 3–4).** We formalize revocation priority and its violation window, prove the PIN trilemma and a tight lower bound on that window under revenue maximization, and characterize the protocol conditions for restoring the guarantee and the costs of contract-level mitigation.
- **Measurement (Sections 5–6).** We analyze pre-freeze successful outflows and post-freeze reverted transfers across nearly nine years of Ethereum USDT and USDC freezes, identify concurrently pending freeze and transfer transactions, trace the technical evolution of the ordering market, show that private submission does not eliminate fee premiums, and analyze the fee premiums of USDC freezes.
- **Equilibrium (Section 7).** We characterize the unique form and equilibrium of the freeze contest and derive analytical bounds on both parties' bids, the probability of taking effect first, and the value of priority ordering outside the protocol.

## 2. Preliminaries

USDT and USDC give their issuers the authority to freeze an address, including in response to a lawful order. When that freeze takes effect depends on how the issuer authorizes it and where Ethereum's block producers place the enforcing transaction relative to transfers from the targeted address.

### 2.1. Ethereum Transaction Inclusion and Ordering

**Block production.** Since the Merge in Sept 2022 [14], Ethereum has used 12-second slots, with at most one block proposed by the designated proposer in each slot. Most proposers outsource block construction through MEV-Boost [15]: builders select and order transactions and submit bids, and the proposer signs the block with the highest bid [17]. The protocol guarantees neither when pending transactions are included nor their order.

**Submission.** Transactions broadcast to the public mempool are visible to all builders and network nodes before inclusion. Transactions submitted through private channels, such as Flashbots Protect, MEV-Share, and private RPCs [22], reach only selected builders and remain outside the public mempool before inclusion.

**Fees and payments.** Under EIP-1559 [9], [21], each transaction specifies caps on the total fee and priority fee per unit of gas. A transaction whose total fee cap is below the block's base fee cannot be included. The sender can replace a pending transaction with a higher-fee transaction using the same nonce. A transaction can pay the block's fee recipient in two ways: a priority fee per unit of gas consumed (paid up to the point of completion or revert); or an internal ETH transfer during execution, which takes effect only if execution succeeds and can thus be conditional on the execution outcome.

**Intent-based settlement.** In intent-based protocols such as 1inch Fusion, CoW Protocol, and UniswapX [4], [12], [32], users only sign off-chain orders specifying the token to sell and the minimum amount to receive. Specialized executor compete to fill these orders. The winning solver chooses the route, submission channel, and fees, then submits a settlement transaction that transfers the user's tokens using the user's prior approval of the protocol contract.

### 2.2. Sanctions and Statutory Freezing Obligations

**OFAC sanctions.** The U.S. Treasury's Office of Foreign Assets Control (OFAC) lists sanctioned parties and their associated cryptocurrency addresses on the Specially Designated Nationals and Blocked Persons (SDN) List [34]. Covered property is blocked upon designation (e.g., 31 CFR §§ 578.201, 578.304), with no enforcement grace period. Subsequent transfers in violation of these rules are null and void (§ 578.202), and relevant U.S. persons must deny all access to blocked virtual currency [33].

**GENIUS Act.** Enacted in July 2025, the GENIUS Act [3] subjects permitted payment stablecoin issuers to regulation as financial institutions under the Bank Secrecy Act (12 U.S.C. § 5903(a)(5)). It requires them to have the technical capability to comply with lawful orders, including seizing, freezing, or burning tokens, or preventing transfers (§§ 5901(16), 5903(a)(6)(B)). Foreign issuers whose stablecoins circulate in the United States must have the same capability (§ 5902(b)(2)).

### 2.3. On-Chain Freezes in USDT and USDC

**Transfers and freeze effects.** USDT and USDC are the two largest stablecoins by circulating supply [2], both implemented as ERC-20 tokens [35]. Holders can transfer tokens directly with `transfer` or authorize a third party with `approve` to transfer tokens on their behalf using `transferFrom`. Both contracts use a privileged call to add a target address $a$ to a blacklist; once this update takes effect, the contract rejects transfers out of $a$.

**Freeze workflows.** Figure 1 shows the two workflows. USDT is administered by a multisignature wallet (*Multisig owners*). At $t_s$, one owner first submits a proposal transaction (`submitTransaction`)

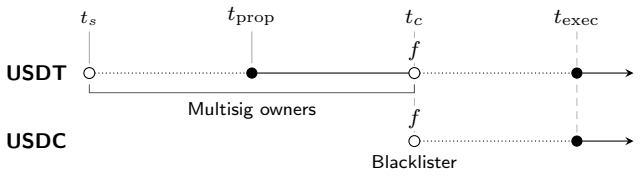


Figure 1. On-chain freezing process for USDT and USDC. Open circles mark transaction submission, dotted lines represent pending intervals before inclusion, and filled circles mark on-chain execution.

containing `addBlackList`$(a)$; the proposal is included on-chain at $t_{\text{prop}}$. The other owners then confirm it. The threshold-reaching confirmation transaction (`confirmTransaction`) is the *enforcing transaction* $f$: it is submitted at $t_c$ and executes at $t_{\text{exec}}$, adding $a$ to the blacklist. USDC instead uses a dedicated administrator account (*Blacklister*) to submit a single `blacklist`$(a)$ transaction $f$ at $t_c$; it executes at $t_{\text{exec}}$ and adds $a$ to the blacklist. For both tokens, the freeze does not restrict transfers out of $a$ until $f$ completes at $t_{\text{exec}}$.

## 3. Revocation Priority: Model and Definitions

A freeze takes effect only when the transaction that enacts it executes, and that transaction is ordered like any other. The contract thus stops transfers that execute after the freeze; the issuer needs to stop transfers submitted after the freeze is issued. This section formalizes that requirement as a security property, *revocation priority*, for any privileged operation that withdraws a permission.

### 3.1. Ledger Model

**Rounds and views.** Given the prior state, the transactions, and the block context, the execution result is deterministic. Ordering proceeds in rounds; the producer of round $k$ is the party that assembles the final block of that round. At the cutoff of the round, the pending transactions received by the producer constitute the view $M_k$.

The ordering rule $O_k$ selects and arranges transactions from $M_k$ to form at most one block that satisfies the protocol constraints. $O_k$ may involve randomness. A valid transaction in $M_k$ that is executable and meets the fee requirement is *includable* in round $k$.

The execution order of operations on the canonical chain is denoted $\prec$, including the order of execution within a single transaction. The moment a transaction is first submitted to any channel is denoted $t_{\text{sub}}$; an operation within a transaction inherits the submission time of the transaction that carries it.

An operation *takes effect* when its effect first becomes available to subsequent operations; the corresponding position is called the *effective position*. $t_{\text{exec}}$ is the timestamp of the block in which the restricting operation $p$ takes effect.

**Propagation.** A publicly broadcast transaction reaches every candidate producer. A transaction sent over a channel reaches each recipient with delay at most $\Delta$; delays are independent, and each falls in any open subinterval of $[0, \Delta]$ with positive probability.

### 3.2. Restrictive Privileges and Revocation Priority

**Definition 1** (Restrictive privileged operation)**.** An operation $p$ is *privileged* if it can only be made to take effect by authorized principals; these principals are called *authorities*. An operation $p$ is *restrictive* if the set $S(p)$ of principals that $p$ restricts contains no authority. The operations that are initiated by $S(p)$ or its delegates, that were previously permitted, and that are prohibited by $p$ constitute the *restricted operation set* $C(p)$.

Take stablecoins as an example: `mint` and `freeze` are both privileged operations, but only freezing is restrictive. When an address $a$ is frozen, $S(p) = \{a\}$, and $C(p)$ consists of the operations that transfer tokens out of $a$, initiated by $a$ or its delegates. These operations are legal before $p$ takes effect and are rejected afterwards, so their outcome depends on their order relative to $p$. Minting does not invalidate any previously legal operation, so $C(p)$ is empty.

**Issuance time and enforcing transaction.** The *issuance time* $t_c$ is the first moment from which $p$ can take effect with no further submission by the authority. The authority's strategy must make $t_c$ finite. A same-nonce replacement that keeps the operation content unchanged and only raises the fee is regarded as the same submission. The outer transaction that successfully makes $p$ take effect is called the *enforcing transaction*, denoted $f$.[1]

**Definition 2** (Revocation priority and violation window)**.** Let $\theta \geq 0$. If there exists an authority strategy such that, for any strategy of the restricted party, almost surely there is no $x \in C(p)$ satisfying both $t_{\text{sub}}(x) \geq t_c + \theta$ and that $p$ has not yet taken effect at the time $x$ successfully takes effect, then the system is said to provide *$\theta$-revocation priority* for $p$, denoted $P_\theta$. Write $P := P_0$. Since $P_\theta$ is upward closed in $\theta$, define the *violation window*

$$w := \inf\{\theta \geq 0 : \text{the system provides } P_\theta\},$$

with $w = \infty$ if no such $\theta$ exists.

**Contract guarantees and system requirements.** The contract checks restricted operations against the state at execution time: while the restriction is in force, restricted operations are rejected. Revocation priority further requires that restricted operations submitted from $t_c + \theta$ onward cannot execute successfully before the restriction takes effect. Therefore, even if the contract correctly performs every authorization check, the system may still violate revocation priority.

1. For USDT, $f$ is the outer transaction that successfully triggers `addBlackList`$(a)$; for USDC, $f$ is the outer transaction that successfully triggers `blacklist`$(a)$.

### 3.3. Blockchain Platform Properties

**Definition 3** (Immediacy (I))**.** If the successful, non-reverted operations in $C(p)$ take effect at their execution position, with their effects available to subsequent calls within the same transaction, the system is said to satisfy *immediacy* $I$.

For instance, USDT and USDC transfers on Ethereum satisfy $I$.

**Ordering information.** The public ordering features of a transaction include its bid, gas parameters, arrival channel, and arrival time. The producer can also identify the signature, the sender address, and the function called, and can obtain through simulation the revenue function $R_M$: for each feasible candidate block in the view $M$, it gives the total payment to the producer (including priority fees and direct payments made within transactions to the fee recipient).

**Definition 4** (Neutrality (N))**.**

(N1) Holding includability, public ordering features, and simulated revenue fixed, a transaction's identity label does not affect the distribution of block selection and ordering.

(N2) Every submission channel is open to all participants, and its latency distribution is independent of the sender's identity.

When $N$ holds, privilege is *in-band*: the authority's transactions share the submission and ordering mechanism with all other transactions, and the authorized identity confers no additional inclusion or ordering priority.

### 3.4. Threat Model

**Restricted party.** The adversary is $S(p)$ and its delegates, such as a spender holding an allowance and a contract controlled by $a$. It can observe the public mempool, on-chain multisig proposals, or off-chain announcements; choose its submission times and channels; and attach on-chain payments to the producer, including payments conditional on the execution outcome. The adversary does not control network latency.

**Authority.** The authority is honest, its keys are not leaked, and the contract correctly implements the prohibition on $C(p)$. The authority may choose the bid, payment structure, channel, and submission time of its transactions.

**Producer.** The producer is controlled by neither the restricted party nor the authority, does not collude with either, and selects blocks according to $O_k$. Profit-seeking behavior is captured by $O_k$; when the ordering rule is mandated by the protocol, the producer follows that rule.

### 3.5. Observability Boundary

The quantities in the definitions fall into three tiers according to observability.

- **Directly observable.** The execution results of on-chain operations, their effective positions, and the corresponding block timestamps can be determined from on-chain data.
- **Partially observable.** $t_c$ and $t_{\mathrm{sub}}(x)$ depend on submission records or mempool observations. Records from different vantage points may be incomplete, and private submissions may be invisible; the first observed time is not the first submission time.
- **Unobservable.** The authority's internal decision times, and off-chain arrangements among participants.

## 4. The PIN Trilemma: Bounds and Trade-offs

This section examines the trade-offs among revocation priority ($P$), immediacy ($I$), and neutrality ($N$). We establish the revenue cost of revocation priority and the resulting impossibility under revenue maximization, then give conditions and tight constructions for retaining each pair of properties. We also quantify how delayed effects and invalidation rewards increase the cost of sustained exclusion.

### 4.1. Revenue Cost and PIN Impossibility

Appendix A.1 lists the technical conventions used throughout this section.

**Theorem 4.1** (Revenue cost of revocation priority)**.** *Suppose the system satisfies $I$, and let $0 \le \theta < \Delta$. For every authority strategy and every revenue gap $G > 0$, there is a restricted-party strategy using only payments conditional on success such that the following event has positive probability. Some round's view admits a block $B^*$ that satisfies all protocol constraints except any authority-priority rule, makes a restricted operation $x$ submitted after $t_c + \theta$ take effect before $p$, and pays at least $G$ more than any block in that round in which $x$ does not succeed. Any ordering rule that guarantees $P_\theta$ must therefore forgo $B^*$ with positive probability.*

*Proof sketch.* Since $\theta < \Delta$, fix a submission time less than $\Delta$ before a round's cutoff such that, with positive probability, this time is after $t_c + \theta$ and $p$ has not taken effect in an earlier round. At that time, the restricted party submits a transaction carrying $x$ and a payment conditional on the success of $x$. The payment exceeds an upper bound on all previously available payments by $G$. With positive probability, this transaction arrives before the cutoff, while all messages sent afterwards, including fee increases and additional funding by the authority, arrive too late. A block that executes $x$ successfully before $p$ then pays at least $G$ more than any block in which $x$ does not succeed. By $I$, successful execution makes $x$ take effect before $p$. The full proof is in Appendix A.3. □

**Definition 5** (Revenue maximization)**.** A producer is *revenue-maximizing* if $O_k$ selects a block maximizing $R_{M_k}$ among blocks composed of includable transactions and satisfying the protocol constraints.

**Theorem 4.2** (PIN impossibility)**.** *Suppose the system satisfies $I$ and $N$, and producers maximize revenue. For every $0 \le \theta < \Delta$, the system cannot provide $P_\theta$: against every*

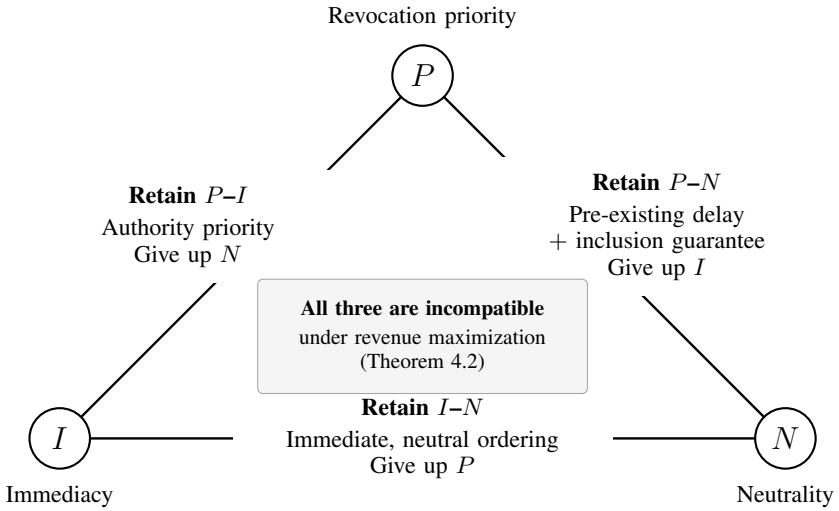


Figure 2. The PIN trilemma under revenue maximization. The $P$–$I$ edge uses authority priority; the $P$–$N$ edge requires a delay already in force and an inclusion guarantee; the $I$–$N$ edge gives up revocation priority.

*authority strategy, the restricted party can cause a restricted operation submitted after $t_c + \theta$ to take effect before $p$ with positive probability. Hence $w \geq \Delta$, and $P$, $I$, and $N$ cannot all hold.*

*Proof.* Under $N$, protocol constraints contain no authority-priority rule, so $B^*$ from Theorem 4.1 is feasible. On the event constructed there, a revenue-maximizing producer selects a block paying at least as much as $B^*$, and thus one in which $x$ succeeds. By $I$, $x$ takes effect before $p$. □

Under revenue maximization, retaining $I$ while providing $P$ requires relaxing $N$. The attack need not observe $f$: its round, submission time, and payment can be fixed in advance, so private submission of $f$ does not restore $P$. A preliminary restrictive operation, such as pausing transfers or temporarily enabling a delay, merely moves the same race to that operation.

The failure can also occur without revenue maximization. Consider a single enforcing transaction $f$ first submitted at $t_c$, with all execution prerequisites met. If $\theta < \Delta$, an operation $x$ submitted between $t_c + \theta$ and $t_c + \Delta$ can, with positive probability, reach every node before $f$. A rule that includes and orders transactions by arrival time may then execute $x$ first, making it take effect immediately.

## 4.2. Retaining Two Properties: Conditions and Costs

Figure 2 summarizes the three choices. Each edge retains two properties and gives up the third, subject to the conditions below.

**Retaining $P$–$I$: Authority Priority.** An authority-priority rule requires the producer to include an includable transaction that makes $p$ take effect and to execute it before the corresponding restricted operations.

For example, suppose each round lasts more than $\Delta$. At the start of a round, the authority publicly broadcasts $f$, with all execution prerequisites and fee requirements met. It reaches every candidate producer before the cutoff. A restricted operation submitted at or after $t_c$ can enter only that round or a later one and therefore cannot take effect before $p$. The system provides $P$ and $I$, gives up N1, and retains N2.

The rule excludes blocks that execute $x$ first, incurring the revenue cost in Theorem 4.1. Producers still maximize revenue over the remaining feasible blocks.

**Retaining $P$–$N$: Delayed Effects and Inclusion Guarantees.** When the transaction carrying a restricted operation is included, it registers a pending request. If $p$ takes effect while the request is pending, the request is invalidated. Otherwise, the operation takes effect at the end of the block in round $D$, where $D \geq 1$ and the registration round counts as round 1. Even if the request eventually succeeds, its effects are unavailable to subsequent operations in the original call, so $I$ is lost. We use $D = 0$ for immediate execution.

Let $\lambda$ bound the time between consecutive round cutoffs. If rounds correspond to Ethereum slots, their duration is 12 seconds.[2]

**Definition 6** ($d$-inclusion guarantee)**.** A platform provides a *$d$-inclusion guarantee* if, for every participant, any publicly broadcast valid transaction whose execution prerequisites and fee requirements are met is included no later than the $d$th round after broadcast. Rounds are counted by their cutoffs; a transaction broadcast exactly at a cutoff belongs to the next round.

2. Ethereum's slot duration is distinct from the interval between produced blocks, since slots can be missed. See https://ethereum.org/developers/docs/consensus-mechanisms/pos/.

**Theorem 4.3** (Effect delay and violation window)**.** *Suppose the platform uses the above effect rule, provides a $d$-inclusion guarantee, and has round lengths at most $\lambda$. Suppose the authority publicly broadcasts $f$ at $t_c$, with all execution prerequisites and fee requirements met. Then*

$$\theta \geq \max\{d - D, 0\}\lambda \quad \Longrightarrow \quad P_\theta.$$

*In particular, $D \geq d$ gives $P$, while immediate execution gives $w \leq d\lambda$.*

The proof is in Appendix A.3. Each additional round of delay reduces this upper bound by $\lambda$ until it reaches zero. The bound holds for any broadcast time; a suitable broadcast time can require less delay. If rounds last more than $\Delta$ and each block must include all includable transactions, $D = 1$ suffices for $P$ (Proposition 2). This still sacrifices immediate composability, including within the original transaction.

Both the effect rule and the inclusion guarantee are independent of identity, preserving $N$. The inclusion guarantee constrains feasible blocks, within which producers still maximize revenue. The costs are users' waiting time and the protocol's inclusion obligation. When $D \geq d$, observing $f$ does not help the restricted party, so the authority can broadcast it publicly. The delay must already be in force; enabling it on demand faces the race in Section 4.1. Delayed effects also disrupt existing interactions. For example, Uniswap V2 infers swap inputs from balances relative to reserves adjusted for outputs; a pending incoming transfer does not supply the required balance.[3]

**Retaining $I$–$N$: Immediate, Neutral Ordering.** Immediate execution and neutral ordering give up $P$ under revenue maximization. The violation window depends on whether inclusion is guaranteed.

**Ordering within a round.** Consider a stablecoin freeze. Suppose $f$ and the transaction carrying $x$ are both includable, and swapping their order leaves all other payments unchanged. If $f$ pays independently of order while $x$ pays more on success, transferring before freezing collects both payments. Raising the freeze fee increases revenue in both orders without changing their difference. The authority could instead offer an order-dependent reward, payable only if the balance remains in $a$ when the freeze executes. A later success-contingent payment from the restricted party can still make transfer-first ordering more profitable by exceeding that reward.

**Exclusion across rounds.** Without an inclusion guarantee, the restricted party can offer an *exclusion reward* conditional on the final state of a block [23], [37]: the producer receives the reward if $p$ has not taken effect by the end of that round, and later enforcement does not claw it back. Verification can use block hashes and account and storage proofs.[4]

**Proposition 1** (Sustained exclusion)**.** *Suppose the system satisfies $N$, producers maximize revenue including exclusion rewards, and every round permits omission of transactions that make $p$ take effect even in a block that registers a restricted operation. Under either immediate execution or any finite delay $D$, for every authority strategy and every finite $\theta \geq 0$ there is a restricted-party strategy that causes a restricted operation submitted after $t_c + \theta$ to take effect before $p$ with positive probability. Thus $w = \infty$.*

*Proof sketch.* Before each cutoff, the restricted party offers an exclusion reward exceeding an upper bound on existing payments. After $t_c + \theta$, it submits a restricted operation with a sufficient success-contingent payment. Under delayed effects, the payment instead rewards successful registration, and exclusion continues until the operation takes effect. Only finitely many rounds are involved. The joint event in which all offers arrive in time and new responses arrive too late has positive probability. The full proof is in Appendix A.3. □

With an inclusion guarantee, Theorem 4.3 gives the finite bound $w \leq d\lambda$. If inclusion is not guaranteed but occurs in each round with conditional probability at least a fixed $\rho_{\text{inc}} > 0$, the risk instead decays geometrically with the waiting period.

## 4.3. Tightness

**Proposition 2** (Tightness of the trilemma)**.** *Suppose the protocol requires each block to include all includable transactions in the producer's view (the* full-inclusion rule*), a block is produced every round, cutoffs are fixed in advance, and each round lasts more than $\Delta$.*
(a) *If the system satisfies $I$ and $N$, and producers maximize revenue subject to full inclusion, then $w = \Delta$, attaining the lower bound in Theorem 4.2.*
(b) *Under revenue maximization, any two of $P$, $I$, and $N$ can hold together. Authority priority gives $P$–$I$ (Section 4.2); full inclusion with delayed effects for $D \geq 1$ gives $P$–$N$; and the system in part (a) gives $I$–$N$.*

*Proof.* (a) Full inclusion imposes the same obligation for every participant and leaves the construction in Theorem 4.2 intact, so $w \geq \Delta$. Conversely, fix any $\theta > \Delta$. The authority publicly broadcasts $f$, with all execution prerequisites and fee requirements met, slightly more than $\Delta$ before a cutoff $T$, choosing $t_c$ so that $t_c + \theta > T$. Then $f$ reaches every candidate producer before $T$ and is included in that round. Any restricted operation submitted at or after $t_c + \theta$ can enter only a later round and is rejected.

(b) Again, the authority broadcasts $f$ slightly more than $\Delta$ before $T$, but after the preceding cutoff. A restricted operation submitted at or after $t_c$ can register no earlier than the round ending at $T$ and can take effect no earlier than that block's end. By then, $f$ in the same block has caused the operation to be rejected or its pending request to be invalidated. Both full inclusion and delayed effects are independent of identity and satisfy $N$. □

Any waiting needed to choose the broadcast time occurs before $t_c$ and is not counted in the violation window.

3. https://github.com/Uniswap/v2-core/blob/master/contracts/UniswapV2Pair.sol

4. https://eips.ethereum.org/EIPS/eip-1186

TABLE 1. VIOLATION WINDOWS UNDER NEUTRAL ORDERING AND REVENUE MAXIMIZATION.

| | Enforcement may be omitted every round | $d$-inclusion guarantee |
|---|---|---|
| Immediate execution | $w = \infty$ (Proposition 1). | $\Delta \le w \le d\lambda$ (Theorems 4.2 and 4.3); under full inclusion, $w = \Delta$ (Proposition 2). |
| Delayed effects, $D \ge 1$ | $w = \infty$ (Proposition 1); exclusion rewards total at least $\sum_{j=2}^{D} b_j$ (Proposition 3). | $w \le \max\{d - D, 0\}\lambda$; $D \ge d$ gives $P$ (Theorem 4.3). |

### 4.4. Contract-Level Mitigation: Delays and Invalidation Rewards

A contract can give up $I$ on its own but cannot guarantee $P$, because inclusion is decided outside the contract. Proposition 1 applies to every finite $D$. Under the conditions below, a delay with $D \ge 2$ requires sustained exclusion across rounds and raises the attack cost. An *invalidation reward* raises the price of exclusion in each round: the issuer commits sufficient funds in advance to pay the producer that includes $f$ when $f$ invalidates pending requests, once per freeze. The delay and reward logic can be implemented in the token contract without changing consensus or ordering rules.

**Proposition 3** (Cost of sustained exclusion)**.** *Suppose the system satisfies $N$, producers maximize revenue, effects are delayed with $D \ge 1$, a block is produced every round, and round lengths are at least $\Delta$. The authority publicly broadcasts $f$ at $t_c$ with all execution prerequisites met. Once received, $f$ remains includable in every round; including it does not reduce payments other than exclusion rewards and increases the producer's payment by at least $b_j > 0$ in round $j$, counting the registration round as round 1. Then any violation requires exclusion rewards totaling at least $\sum_{j=2}^{D} b_j$. In particular, when $D \ge 2$, no violation is possible without paying for exclusion. If $b_j = b$ for all $j$, the lower bound is $(D-1)b$.*

*Proof.* A violating operation is submitted no earlier than $t_c$, so its registration round ends after $t_c$; since rounds last at least $\Delta$, $f$ has arrived and is includable from round 2 on. Including $f$ in any of rounds $1, \dots, D$ rejects or invalidates the request. In each round $j \ge 2$, an exclusion reward below $b_j$ makes appending $f$ to the block strictly more profitable, so keeping $f$ out through round $D$ requires rewards totaling at least $\sum_{j=2}^{D} b_j$. □

With a constant payment increment, the lower bound grows linearly in $D$: each additional round of delay requires another round of paid exclusion. A fixed invalidation reward is paid at most once per freeze, in addition to transaction fees; when no request is invalidated, only those fees are paid. An attack whose gain is below the exclusion-cost bound is unprofitable.

Table 1 summarizes the bounds under neutral ordering. Each entry requires the assumptions of its cited result.

## 5. Data and Measurement Methods

To study freezes on the $I$–$N$ edge, we combine on-chain records of USDT and USDC freezes and outflows from restricted addresses on Ethereum with public mempool observations and off-chain data.

### 5.1. On-Chain Records

We extract USDT and USDC blacklist event logs (including USDT's `DestroyedBlackFunds`), internal call traces, block data (`timestamp`, `baseFeePerGas`, and `feeRecipient`), and transaction receipts (`effectiveGasPrice` and `gasUsed`) from Google BigQuery, covering the period from USDT's contract deployment on November 28, 2017, through November 10, 2026.

We track the freeze authorities for both tokens: the USDT multisig wallet's `owners` and confirmation threshold, and USDC's `blacklister` across proxy contract upgrades. We use these authority histories to reconstruct each freeze. For USDT, we link each proposal to all its confirmations and to the enforcing transaction $f$, the outer transaction that successfully triggers `addBlackList`$(a)$. We record the time the proposal is recorded on-chain, $t_{\text{prop}}$, and identify confirmations whose calls fail and subsequent execution retries. For USDC, $f$ is a successful `blacklist`$(a)$ transaction sent by the `blacklister` authorized at execution; we record reverted calls separately. Replacements with the same sender and nonce that only raise the fee are treated as one transaction. An address frozen for both tokens is counted once per token.

For each frozen address $a$, we also collect its complete incoming transfer, outgoing transfer, and approval histories for that token; rejected outgoing transfer calls; its balance in that token when the freeze takes effect; and internal ETH transfers within the transactions carrying its outflows.

### 5.2. Restricted Addresses and Outflows

We define *restricted addresses* as addresses whose restrictions have not been lifted via `removeBlackList` or `unBlacklist` by the end of the collection period.

Using the outer transaction as the unit of analysis, we extract restricted outflows $x \in C(p)$ from each restricted address $a$ on either side of the freeze's effective time $t_{\text{exec}}$:

- **Pre-freeze outflows ($x \prec f$):** Outflow operations that execute successfully before $f$.

- **Rejected outflows ($f \prec x$):** Outflow operations that revert because of the blacklist and for which the balance and allowance are sufficient in the state immediately before $f$.

To filter out dust and zero-value poisoning, descriptive distributions and bid statistics include only outflows of at least $10,000 (the reporting threshold under the U.S. Bank Secrecy Act).

We manually inspected 50 freeze transactions, 50 pre-freeze outflows, and 50 reverted outflows for both USDT and USDC (300 transactions in total) to verify their execution outcomes.

### 5.3. Mempool and Off-Chain Data

**Public mempool archives.** We combine Blocknative Mempool Archive (Jan 2019–Mar 2025), Flashbots Mempool Dumpster (Jul 2023–Nov 2026), and ethPandaOps Xatu (Mar 2023–Nov 2026). For each transaction, we use the earliest observation across these archives as its first-seen time, $t_{\text{first}}$; for a replace-by-fee (RBF) chain, we use the first attempt.

**Private submission channels.** To identify transactions that may bypass the public mempool, we collect auxiliary signals of private submission, including private-transaction flags from mempool sources(Archive, Dumpster, Xatu) and Dune data on Flashbots Protect and MEV-Share. We classify a transaction as privately submitted if any source flags it as private.

**Sanctions, judicial records, and entity labels.** We collect Ethereum addresses and their designation dates from the OFAC Specially Designated Nationals (SDN) list, along with publicly available federal court seizure warrants on CourtListener. We combine address labels from Etherscan, Dune, and security firms, assigning `unknown` when labels conflict or evidence is insufficient.

## 6. Empirical Study of On-chain Freezes

In this section, we ask which addresses issuers freeze, when outflows succeed or revert, and how the market of execution order has evolved.

### 6.1. The Landscape of On-Chain Freezes

We collect all blacklist events for USDT and USDC, together with USDT's `DestroyedBlackFunds` events, on Ethereum mainnet from contract deployment (November 2017 for USDT and August 2018 for USDC) through November 2026, covering 3,206 USDT addresses and 895 USDC addresses. After excluding addresses whose restrictions were lifted via `removeBlackList` or `unBlacklist` by the end of the collection period, the restricted address set contains 2,863 USDT addresses and 698 USDC addresses: 3,561 address–token pairs representing 3,391 distinct addresses.

To identify the reasons for freezing, we cross-reference the restricted addresses with the OFAC Specially Designated Nationals (SDN) list, finding 127 listed addresses frozen for both USDT and USDC. A search of litigation records on CourtListener identifies another 245 USDT addresses and 38 USDC addresses. Tether also blacklisted 17 special addresses with at least 24 leading zeros, including the zero address (`0x00..00`) and precompiled contract addresses (e.g., `0x00..01`). Overall, public sanctions and U.S. litigation records account for only a small part of restricted addresses, leaving the vast majority of on-chain freezes without regulatory attribution.

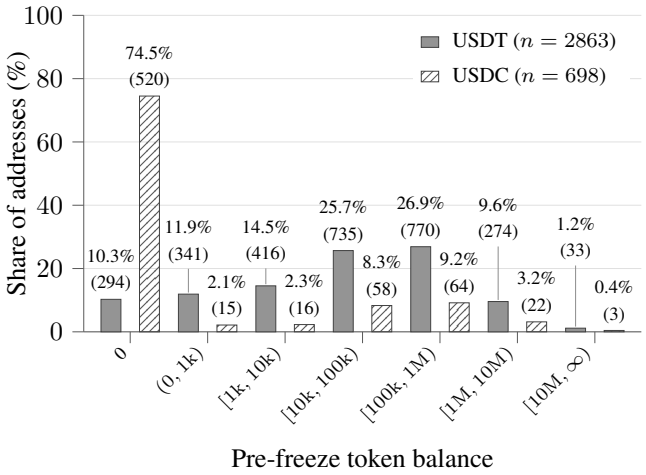


Figure 3. Distribution of pre-freeze token balances.

At the time their freezes took effect, the addresses held a total of $3.838 billion on-chain (USDT: $2.883 billion; USDC: $955 million; distribution in Fig. 3). Tether destroyed $860 million through calls to `destroyBlackFunds`. When their freezes took effect, 10.27% (294) of restricted USDT addresses and 74.50% (520) of restricted USDC addresses had zero balances. Their transfer histories show that, among these 814 zero-balance address–token pairs, 66.71% (543) had never received the corresponding token before the freeze. Among restricted USDC addresses with no prior incoming transfers, 18.48% (85) appear on the OFAC SDN list. Issuers thus preemptively freeze some SDN-listed addresses before they ever receive tokens.

> **Finding 1.** Stablecoin freezing includes preemptive blacklisting, while OFAC sanctions and U.S. litigation records account for only a small fraction of restricted addresses.

### 6.2. Pre- and Post-Freeze Outflows and Freeze Execution

**Successful pre-freeze outflows ($x \prec f$).** Among restricted addresses that had received funds, 40.47% (USDT) and 48.74% (USDC) had at least one outflow of $10,000 or more before $t_{\text{exec}}$. The per-address medians were 4 outflows ($390,000) for USDT and 2.5 outflows ($255,997.80)

for USDC, with totals of 544,353 outflows ($1.02 trillion) and 682,574 outflows ($696.249 billion), respectively. After aggregating both tokens by address, the two Binance hot wallets compromised on October 29, 2026[5] ranked first and second by both outflow count and value, jointly accounting for 96.79% of all qualifying outflows and 99.53% of their value across both tokens. Their respective outflow counts were 85.68 and 40.13 times that of the third-ranked address, which had only 9,440 outflows. The vast majority of outflows from these two addresses were historical hot-wallet activity before the compromise. Including them would dominate the transaction-level distributions, so all subsequent outflow statistics, including Section 6.3, exclude both addresses.

Among the remaining outflows, the vast majority occurred more than 24 hours before the last freeze $t_{\text{exec}}$ (97.68% for USDT and 99.01% for USDC). Within the final 24 hours, there were 673 USDT outflows from 272 addresses and 102 USDC outflows from 20 addresses. The median amounts were $69,951.93 and $29,919.95, respectively, compared with $68,277.17 and $15,696.86 for outflows more than 24 hours earlier; the median times to $t_{\text{exec}}$ were 8.57 and 0.85 hours. Some outflows occurred immediately before the freeze took effect: 25 were within the final 5 minutes (21 USDT and 4 USDC), including 8 less than 1 minute before (1–4 blocks). The shortest gaps were just 12 seconds for USDT and 36 seconds for USDC.

**Rejected post-freeze outflows ($f \prec x$).** After the freeze took effect, 108,158 positive-value USDT outflow attempts from 900 addresses and 660 USDC attempts from 80 addresses reverted because of the blacklist. Of these, 3,988 USDT attempts and 288 USDC attempts were initiated directly by the restricted address, requested at least $10,000, and did not exceed its balance immediately before $f$. The median times since $t_{\text{exec}}$ were 41.11 and 21.15 days, respectively. Eight attempts occurred within 5 minutes after the freeze took effect, all for USDT, including 5 less than 1 minute afterward. The shortest gap was just 12 seconds for USDT and 624 seconds for USDC.

**Freeze execution intervals.** On-chain time gaps relative to $t_{\text{exec}}$ alone do not establish submission order: if $f$ executes within seconds of issuance, an outflow included one minute earlier was still submitted before $t_c$; conversely, a reverted outflow may have been submitted before $t_c$ but overtaken by $f$. We therefore distinguish two stages, before and after $t_c$.

Before $t_c$, USDT freezes are initiated through an on-chain multisig: the proposal reveals the target address when it is recorded on-chain at $t_{\text{prop}}$, and the freeze takes effect only after the confirmation threshold is reached. USDC freezes use a single call and have no such interval. For USDT, the interval from $t_{\text{prop}}$ to $t_{\text{exec}}$ had a median duration of 3.87 hours and ranged from 0 seconds to 677 days. Across these intervals, 9 freeze proposals incurred 10 out-of-gas call failures; all freezes subsequently took effect. These intervals contained 157 successful outflows of at least $10,000 each, totaling $241 million.

After $t_c$, we use the first-seen time $\hat{t}_c$ of an enforcing transaction $f$ visible in the public mempool before execution as an upper bound on $t_c$ (1,517 USDT and 134 USDC transactions). From $\hat{t}_c$ to $t_{\text{exec}}$, the median waiting times were 7.12 seconds for USDT and 6.65 seconds for USDC. The respective minima were 0.02 and 0.48 seconds, P95 values were 43.72 and 13.85 seconds, and maxima were approximately 3.18 and 1.04 hours.

**Concurrently pending transactions.** We define a transaction's observed pending interval in the public mempool as the interval from its first-seen time to inclusion. Among 12,252 freeze–outflow pairs with first-seen records for both transactions, 2 had overlapping intervals, both for USDT. In both, the transaction first seen later executed first: the outflow in one case, the freeze in the other.

On November 9, 2023, enforcing transaction `0xb852...6b49` targeting `0xe86e...47c7` entered the public mempool with a bid of 50 gwei, below the base fee of 151.11 gwei at the time. Outflow `0xff73...3d63` first appeared 22.6 minutes later with a fee cap of 179.51 gwei. It was included 8 seconds later in block 18535340, transferring 100,000 USDT, while the freeze took effect in block 18536174 after a wait of 3.18 hours from its first-seen time.

In the other case, on October 26, 2020, outflow `0x1ed9...5a1c` from `0x18bc...80f6` (138,794 USDT; 61 gwei) entered the public mempool 12.78 seconds before enforcing transaction `0x9805...b685` (90 gwei), but the freeze was included first in block 11132910. The outflow was included 1 hour 42 minutes 17 seconds later and reverted.

> **Finding 2.** Concurrently pending freezes and outflows are rare in the public mempool; where they occur, execution order follows fees and inclusion rules, not authority.

## 6.3. The Market for Execution Order

We examine the freeze transactions considered in the preceding subsection, which take effect at $t_{\text{exec}}$, and outflow transactions of at least $10,000 within the preceding 24 hours.

**6.3.1. Fees.** Fees are Ethereum's native means of purchasing an execution position. We define the *fee multiple* as a transaction's `effectiveGasPrice`, including the base fee and the actual priority fee, divided by the median `effectiveGasPrice` across all transactions in the same block. Overall, USDT freeze and outflow transactions remain close to the block median, with median multiples of 0.96 and 1.01, respectively. USDC freeze transactions pay a substantial premium (3.80), whereas outflow transactions remain near the median (0.91). As Figure 4 shows, USDC freeze transactions exhibit extreme premiums from H2 2023 through H1 2025, with the half-year median peaking at

5. Addresses `0x28c6...1d60` and `0xee7a...4055`.

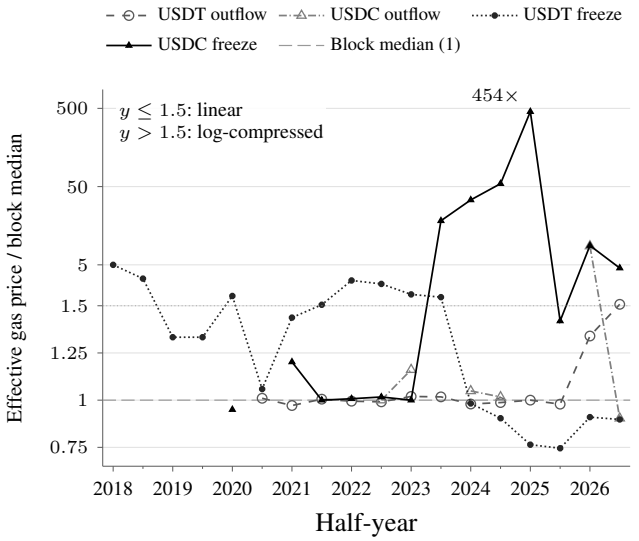


Figure 4. Relative gas prices of freeze and outflow transactions.

453.96× before declining. Among the 292 freezes with corresponding outflows, 69.86% have total execution fees above the median of their corresponding outflow transactions; the median paired fee ratio is 2.72×.

Both freeze and outflow transactions exhibit replace-by-fee (RBF), in which transactions are resubmitted with the same nonce at higher fees. All observed chains involve USDT: two are freeze-transaction replacement chains, one of which raises the bid to 7.5× its initial level within 6 minutes. Five are outflow-transaction replacement chains: four raise bids to 1.18–2.49× their initial levels, while the fifth adds only 1 wei.

Within the 10 minutes before a freeze, outflow value and fee multiple are significantly positively correlated (Spearman $\rho = 0.468$, $p = 0.0053$). The correlation weakens over a 1-hour window ($\rho = 0.279$, $p = 0.0012$) and disappears over a 5-hour window ($\rho = -0.009$, $p = 0.871$).

**6.3.2. Private Submission.** Bidding in the public mempool exposes pending transactions. Private submission instead sends transactions directly to builders, bypassing the public mempool before inclusion. The sample contains 220 private USDT outflow transactions (32.69%), first observed in March 2023, and 41 private USDC outflow transactions (40.20%), first observed in March 2026. Private freeze transactions total 567 for USDT and 106 for USDC, first appearing in March 2022 and May 2023, respectively. They occur in every month from July 2025 for USDT and July 2026 for USDC. In 2026 through November 10, private submissions account for 39.21% and 14.06% of all USDT and USDC freeze transactions, respectively.

Private submission does not eliminate fee premiums. The median fee multiples for private and public pre-freeze outflow transactions are 1.17 and 1.00 for USDT, and 2.63 and 0.91 for USDC. At the 90th percentile, the gap widens to 5.93 versus 1.49 for USDT and 16.93 versus 1.00 for USDC.

Even when $f$ is submitted privately, $x$ can still be included immediately before it. On March 18, 2026, the private outflow transaction $x$ (`0x44a4...cece`) from `0x11dc...70de` moved 100,014 USDT in block 24683662. The private freeze transaction $f$ (`0x326a...ffe3`) targeting that address took effect in the next block, 24683663, 12 seconds later.

**6.3.3. Specialized Execution.** Connecting to multiple builders and bidding in real time require infrastructure and expertise. Intent-based protocols delegate these tasks to specialized executors, such as resolvers. The first outflow executed through an intent-based protocol in the sample occurs on October 10, 2023. Of the 753 outflow transactions from EOAs, 50 (6.64%) are signed by third parties and executed through pre-approved contracts via `transferFrom`: 39 USDT transactions totaling $326 million and 11 USDC transactions totaling $6,369,712.09. The 1inch family, UniswapX, and ParaSwap/Velora Delta execute 32, 7, and 11 transactions, accounting for 15.25%, 82.84%, and 1.91% of the total value of these 50 outflows, respectively.

Of these transactions, 11 make internal transfers totaling 1,751.20 ETH to block fee recipients. All 11 involve USDT and together account for $325 million in outflows. The 1inch family accounts for seven transactions, moving $50.15 million (15.42%) and paying 581.73 ETH (33.22%) to fee recipients. UniswapX accounts for four transactions (36.36%), moving $275 million (84.58%) and paying 1,169.47 ETH (66.78%) to fee recipients. At contemporaneous ETH/USD exchange rates, the maximum payment-to-outflow ratio is 15.68%. Across these 11 transactions, 99.99% of the ETH paid to block fee recipients comes from recipient addresses associated with the 2026 Binance hack.

**Finding 3.** Private submission does not eliminate fee premiums, and outflows can still be included immediately before private freezes. Block producers collect both priority-fee premiums and internal transfers from specialized executors to fee recipients. The latter take effect only upon successful execution and arise almost entirely from a single incident.

The next section models the incentives in this market as a game, using equilibrium analysis to explain these observations and their implications for neutral ordering.

## 7. Equilibrium of the Freeze Contest

Section 4 allows arbitrarily large payments by the restricted party. Here we bound the stakes on the $I$–$N$ edge with revenue-maximizing producers: a freeze and a conflicting transfer then compete to take effect first in a *freeze contest*.

## 7.1. Setup

The issuer submits $f$ at $t_c$ with a sufficient fee, when $a$ holds $V > 0$. We say $f$ and $x$ *co-occur* if a transfer $x \in C(p)$ submitted after $t_c$ is pending alongside $f$.

The capabilities in Section 3.4 concern two layers: which view wins, and ordering within that view. Section 4 addresses the latter: when $f$ and $x$ share a view, $x$ takes effect first regardless of the freeze fee (Section 4.2). We address the former. Each constructs a revenue-maximizing block from its view, and the proposer selects the highest-value block; a transaction's bid thus raises the chance that a block containing it wins (Section 2.1). Freeze fees affect block selection but not ordering within a block. The *decisive round* is the first round whose winning view contains $f$ or $x$. If the winning view contains $x$, the transfer takes effect first; otherwise, the freeze takes effect and later transfers are rejected.

The issuer knows only the probability $q$ that the restricted party is present, and the restricted party knows only the probability $\mu$ that the freeze has been issued, with $q, \mu \in (0, 1]$. The issuer's bid $\beta_f \geq 0$ is paid on top of the fee required for inclusion, regardless of which transaction takes effect first. The restricted party's bid $\beta_x \in [0, V]$ is the success-contingent payment of Theorem 4.1, paid only if $x$ succeeds. Both contribute to block value.

**Assumption 1** (Decisive round)**.** *Conditional on co-occurrence, the decisive round's winning view contains both transactions with probability $\nu \in [0, 1]$. Otherwise, it contains only one of them. The probabilities of the $f$-only and $x$-only outcomes depend only on the bids and satisfy:*

(i) *exactly one outcome occurs;*

(ii) *their odds is the ratio of a function of the issuer's bid to a function of the restricted party's bid, both strictly increasing and zero at zero;*

(iii) *multiplying both bids by a positive constant leaves the odds unchanged;*

(iv) *swapping the bids swaps the outcome probabilities;*

(v) *each party's odds of winning are concave in its own bid.*

Channel overlap determines $\nu$. If both sides broadcast publicly, $\nu = 1$ and the model reduces to Section 4. Condition (i) defines the decisive round. Condition (ii) reflects that a bid raises only the value of blocks containing its transaction. Under condition (iii), scaling both bids up favors blocks containing $f$ or $x$ without changing the $f$-only/$x$-only odds. Condition (iv) imposes symmetry on the $f$-only/$x$-only contest, motivated by the neutrality principle in Section 3. Condition (v) ensures a pure-strategy equilibrium. We set $\sigma_x = 0$ when $\beta_x = 0$.

**Proposition 4** (Unique contest form)**.** *For positive bids, conditions (i)–(iv) of Assumption 1 hold if and only if there is a unique $r > 0$ such that, conditional on the winning view containing only one of the two transactions, it contains $f$ with probability $\beta_f^r/(\beta_f^r + \beta_x^r)$. Condition (v) is then equivalent to $r \leq 1$. Thus, under Assumption 1, for $\beta_x > 0$,*

$$\sigma_f = (1-\nu)\frac{\beta_f^r}{\beta_f^r + \beta_x^r}, \qquad \sigma_x = 1-\sigma_f, \quad 0 < r \leq 1. \tag{1}$$

See Appendix B.1 for the proof. Under Assumption 1, block selection when the winning view contains only one of the two transactions follows a Tullock contest, with $r$ measuring bid responsiveness. The parameter $\nu$ is the probability that Section 4 applies and accounts for the restricted party's within-view ordering advantage. Whenever $\beta_x > 0$, $\sigma_x > 0$, so the impossibility in Section 4 persists. The characterization draws on contest theory [11], [28]; we apply it to ordering two competing transactions. Our data cannot identify $r$ or $\nu$ (Section 7.3), so the results below cover all $0 < r \leq 1$ and $0 \leq \nu < 1$.

**Payoffs.** The issuer values a timely freeze at $\Psi > 0$, its loss if the freeze fails; $\Psi$ need not equal $V$. The expected payoffs are

$$U_f = \Psi(1 - q\sigma_x) - \beta_f, \quad U_x = (1 - \mu\sigma_f)(V - \beta_x). \tag{2}$$

The restricted party is a value-maximizing adversary of the type allowed by Section 3.4.

**Lemma 7.1** (Equilibrium)**.** *For $\nu < 1$, the contest has a unique pure-strategy equilibrium, with both bids interior (Appendix B.2).*

Asterisks denote equilibrium values. Motivated by the observed bidding asymmetry and the issuer's regulatory obligations, our empirical interpretation concerns the regime $q\Psi/(\mu V) > 4/[4 - (1-\nu)\mu(2-r)]$, precisely when the issuer outbids its target at equilibrium.

## 7.2. Participation

The contest occurs only if the restricted party participates. Participation costs $c > 0$ to attach a success-contingent payment and reach producers through private channels. A nonparticipant sends an ordinary transfer with $\beta_x = 0$ and loses if a freeze has been issued. Define the *expected gain* as the additional balance retained through participation, net of the success-contingent payment but before $c$. The issuer cannot observe an address's participation; its belief $q$ is proportional to the participation rate among addresses with the same balance.

**Proposition 5** (Participation)**.** *Fix $\Psi$.*

(a) *A participant retains its balance ahead of a freeze with probability $\mu\sigma_x^*$. Its expected gain lies between $\mu\nu V$ and $\mu V$, increases strictly in $V$ and $\nu$, and decreases strictly in $q$.*

(b) *There is a unique threshold $\bar{V} > c/\mu$, with $\bar{V} < c/(\mu\nu)$ if $\nu > 0$. The equilibrium participation rate is zero for $V \leq c/\mu$ and one for $V \geq \bar{V}$. Between these thresholds, only a fraction participates; this fraction increases with balance, and participants' expected gain equals the cost.*

(c) *The threshold $\bar{V}$ increases with $c, \Psi$, and decreases with $\nu$.*

See Appendix B.3 for the proof. Participation can recover at most $\mu V$; the component $\mu\nu V$ requires no bidding competition. Hence neither threshold bound depends on the unobserved $\Psi, q, r$. For illustration, $c = \$10{,}000$, $\mu = 2\%$, and $\nu = 0.8$ place the threshold between \$500,000 and \$625,000. The model thus predicts that contested freezes concentrate at large balances, with frequency increasing in balance.

If delegated execution replaces the fixed cost $c$ with a proportional fee $\varphi \in (0, 1)$ of the expected gain, the net gain $(1 - \varphi)g$ remains positive at every balance $V > 0$, eliminating the participation threshold.

If delegated execution replaces the fixed cost $c$ with a fee equal to a fraction $\varphi \in (0, 1)$ of the expected gain, the participation threshold disappears. The gain after delegation is positive at every balance, eliminating the threshold.

### 7.3. Freeze Fees Cannot Buy Priority

Can the issuer buy priority by bidding more against a participating restricted party?

**Proposition 6** (Taking effect first under bounded stakes)**.** *Let $\nu < 1$. At equilibrium:*

(a) *conditional on the restricted party being present, its transfer takes effect first with probability*

$$\sigma_x^* \geq \nu + (1 - \nu)(1 + (q\Psi/(\mu V))^r)^{-1};$$

(b) *$\sigma_x^*$ increases strictly in $\nu$; as $\nu \to 1$, $\sigma_x^* \to 1$, both bids vanish, and the probability that the freeze takes effect first approaches $1 - q$;*

(c) *the bids satisfy*

$$\beta_f^* \leq (1 - \nu)qr\Psi/4, \qquad \beta_x^* < (1 - \nu)\mu rV.$$

See Appendix B.4 for the proof. The first term in (a) follows from the within-round ordering argument in Section 4.2: when both transactions share the winning view, the restricted party's transfer takes effect first regardless of the issuer's bid. The second term describes the contest between different views and depends only on the effective stake ratio $q\Psi/(\mu V)$. Part (b) recovers Section 4.2 in the limit $\nu \to 1$. Part (c) discounts bids by the probability of facing the other party, making them a small fraction of stakes when $q$ and $\mu$ are small. For example, with $V = \$1$ million, $\Psi = 2V$, $r = 1$, $q = 2.5\%$, $\mu = 2\%$, and $\nu = 0.8$, the bound in (a) is approximately $85.7\%$. The issuer bids about \$2,043 above the inclusion requirement, versus a success-contingent payment of \$819 by the restricted party, a bid ratio of approximately $2.49$. Since co-occurrence is rare, the overall probability that the freeze takes effect first remains about $97.9\%$. We cannot estimate $r$ or $\nu$: same-block pairs are sparse and censored, since revert protection can keep a transfer off-chain when the freeze executes first.

### 7.4. Buying Priority Outside the Protocol

An issuer–producer agreement privately applies authority priority (Section 4.2), relaxing N1 without changing the protocol. Fix $q$ and $\beta_x^*$. The agreement only converts transfer-first outcomes into freeze-first outcomes, paying at most $\beta_x^*$ per interception; other outcomes and payments are unchanged. At the baseline bids, let $\rho \in [0, 1]$ be the fraction of otherwise successful transfers intercepted, conditional on co-occurrence. The baseline freeze-failure probability becomes $(1 - \rho)q\sigma_x^*$.

**Proposition 7** (Value of ordering control)**.** *For any ordering rule satisfying N1, let $\beta_f^*$ maximize the issuer's baseline payoff against $\beta_x^* \in [0, V]$, with conditional transfer-success probability $\sigma_x^* > 0$. If $\Psi > \beta_x^*$, define $\Gamma := q\sigma_x^*(\Psi - \beta_x^*)$. For any agreement $A$ with $\rho > 0$, the issuer's expected gain $G_A$ after bid reoptimization, before maintenance costs, satisfies*

$$\rho\Gamma \leq G_A \leq q\Psi.$$

*Maintenance costs below $\rho\Gamma$ thus leave a positive net gain. Compensation per interception is at most $V$. These bounds do not require Assumption 1.*

See Appendix B.5 for the proof. The lower bound uses the minimum net gain $\Psi - \beta_x^*$ per interception. In Section 7.3, $\Gamma \approx 0.043V$: at $\rho = 1$, the gain lies between approximately \$42,800 and \$50,000 per freeze, over twenty times the freeze-fee increment.

**Proposition 8** (Stronger incentives for ordering control)**.** *Under Assumption 1, $\Psi \geq \mu rV$ implies $\Gamma > 0$, and:*

(a) *$\Gamma$ increases strictly in $\nu$ and approaches $q\Psi$ as $\nu \to 1$;*

(b) *if $\Psi \geq 2\mu rV$, $\Gamma$ increases strictly in $\Psi$ and $q$.*

See Appendix B.6 for the proof; $\Psi \geq 2V$ satisfies both conditions. Part (a) shows that greater view overlap strengthens the incentive to buy ordering control: bids cannot change ordering within a shared view, whereas an agreement can. Part (b) shows that the incentive grows with the issuer's value of a timely freeze and the restricted party's participation frequency. If adopting private channels and delegated execution increases $q$, specialization on the restricted party's side strengthens the incentive for integration. Under the stated cost condition, these incentives favor private issuer–producer integration that can move the $I$–$N$ edge toward authority priority without protocol changes or visibility in the public mempool.

## 8. Discussion

**Protocol-Level Mitigations.** Protocol upgrades that preserve immediacy and neutrality can only narrow the violation window; they cannot restore revocation priority. Encrypted mempools [8], [10] hide transaction contents and conditional payments before ordering. Yet the restricted party can submit an outflow without observing a freeze, and the outflow can still take effect first while the freeze is propagating. Fair-ordering consensus [18], [19] bases ordering on receipt order rather than revenue maximization, but cannot eliminate races during propagation or frontrunning within a batch [24].

Ethereum has scheduled fork-choice enforced inclusion lists (FOCIL, EIP-7805 [29]) for its Hegotá upgrade. A block must include transactions on timely received lists if they remain valid at the end of the block and enough gas remains to include them. For a publicly broadcast freeze, FOCIL sets no deadline for listing and therefore provides no fixed $d$ for the $d\lambda$ bound in Theorem 4.3. Even when the freeze is included, the builder can place a later-submitted transfer before it in the same block, so revocation priority still fails, as in Theorem 4.2.

**Compliance Pressure and Neutrality.** If issuers are held accountable for whether funds leave after a freeze is issued, those unable to secure guaranteed execution priority through public or private bidding will bypass the public market and privately purchase priority from producers. Ethereum's block construction is concentrated among a few builders [17], [38]. A small number of off-chain agreements can therefore cover most block production and shift the ledger toward authority priority without changing the protocol or leaving a public record, eroding public-chain neutrality through arrangements outside the consensus protocol [36].

**Across Stablecoins and Chains.** On Ethereum, PYUSD, FDUSD, and RLUSD submit freezes in a single transaction through authorized externally owned accounts or multisig wallets that collect signatures off-chain, such as Safe. This avoids exposing on-chain proposals, but the enforcing transactions still enter the same ordering market once submitted. Across chains, the game varies with the concentration of ordering power. Public chains with block auctions, such as Solana, likewise exhibit public bidding and private bundle bidding. Single-sequencer L2s such as Arbitrum, Optimism, and Base cannot prevent later submissions from arriving first during propagation delays, whether they order by fees or arrival time; switching to authority priority requires only a change in sequencer policy. Public chains with concentrated validator sets and no public auctions, such as Tron, shift the contest to off-chain coordination with a few validators, while dedicated settlement chains such as Circle's Arc directly adopt regulated, permissioned validator sets.

**Privileged Revocation Beyond Stablecoins.** The same dilemma extends to other forms of on-demand revocation on-chain. For tokenized assets such as PAXG and ERC-3643, the competing operations are address-level freezes and ordinary transfers; for lending protocols such as Aave V3, they are an independent emergency administrator's pause and harmful borrowing or liquidation when an external price feed is compromised. The same conflict arises between a security role's mainnet pause and withdrawal settlement once a problematic withdrawal becomes finalizable in optimistic rollup bridges such as the OP Stack, and between an administrator's role revocation and calls made with a compromised private key in access-control contracts such as OpenZeppelin AccessControl. Subsequent revocation does not roll back operations that have already taken effect.

## 9. Related Work

**Sanctions and stablecoin freezes.** Studies of blockchain sanctions trace the Bitcoin activity of designated entities [41], examine continued use of Tornado Cash under uneven enforcement [20], and measure inclusion delays for Ethereum transactions affected by OFAC sanctions [36]. For USDT, BlockSec [1] measures outflows before blacklisting takes effect. We pair on-chain records with public mempool observations to determine whether a freeze was already visible when an sucessful/reverted outflow that executed ahead of it first appeared.

**Economic censorship and inclusion guarantees.** McCorry et al. [23] show that an in-band contract can pay miners for a censoring mining strategy. Winzer et al. [37] consider one account over a finite window, where a per-block bribe sustains exclusion at a cost linear in the window's duration. Fox et al. [16] define censorship resistance as the cost of excluding a transaction for a fixed interval as a function of its tip. They show that conditional tips to concurrent proposers raise this cost. Berger et al. [7] size the challenge period that lets an honest fraud-proof party include its moves against such bribes. That period is a confirmation delay the application already imposes. These works study whether a transaction can be included by a deadline, while we ask whether a freeze takes effect before a later-submitted transfer from the targeted address.

**MEV and its mitigation.** Daian et al. [13] term the revenue obtainable from inclusion and ordering miner extractable value (MEV). Prior strategies aim to mitigate it [8], [10], [18], [19], [39]. Zhang et al. [40] study MEV even in batch auctions, where execution does not depend on order. This paper focuses on revocation priority, the requirement that a submitted restriction take effect before later prohibited operations. As discussed in Section 8, none of those mitigation strategies can fundamentally restore it.

## 10. Conclusion

We prove that under immediacy, neutrality, and revenue maximization, no issuer bid can guarantee that a freeze blocks a later-submitted outflow, and nearly nine years of freeze records show that execution order already has a price. During the October 2026 Binance hack, delegated outflows paid block producers only on success, so a producer facing both a freeze and an outflow would forgo that payment by executing the freeze first. If a failed freeze costs the issuer more than that payment, it has an incentive to buy priority from a few builders off-chain: eroding ordering neutrality in practice even while the consensus protocol remains nominally unchanged.

## References

[1] Following the Frozen: An On-Chain Analysis of USDT Blacklisting and Its Links to Terrorist Financing, 2025-07-11. https://blocksec.com/blog/following-the-frozen-an-on-chain-analysis-of-usdt-blacklisting-and-its-links-to-terrorist-financing.

[2] DefiLlama Stablecoins Data, 2026. https://defillama.com/stablecoins.

[3] 119th United States Congress. Guiding and establishing national innovation for U.S. stablecoins act (GENIUS act). https://www.congress.gov/bill/119th-congress/senate-bill/1582, 2025. S.1582, signed into law July 18, 2025.

[4] 1inch. 1inch intent-based swap (Fusion mode). https://business.1inch.com/portal/documentation/apis/swap/intent-swap/introduction. Accessed: 2026-11-16.

[5] J. Aczél. *Lectures on Functional Equations and Their Applications*, volume 19 of *Mathematics in Science and Engineering*. Academic Press, 1966.

[6] James P. Anderson. Computer security technology planning study, volume II. Technical Report ESD-TR-73-51, Vol. II, Electronic Systems Division, Air Force Systems Command, Hanscom Field, Bedford, MA, October 1972.

[7] Ben Berger, Edward W. Felten, Akaki Mamageishvili, and Benny Sudakov. Economic censorship games in fraud proofs. In *Proceedings of the 26th ACM Conference on Economics and Computation*, EC '25, page 670–687, New York, NY, USA, 2025. Association for Computing Machinery.

[8] Jan Bormet, Sebastian Faust, Hussien Othman, and Ziyan Qu. {BEAT-MEV}: Epochless approach to batched threshold encryption for {MEV} prevention. In *34th USENIX Security Symposium (USENIX Security 25)*, pages 3457–3476, 2025.

[9] Vitalik Buterin, Eric Conner, Rick Dudley, Matthew Slipper, Ian Norden, and Abdelhamid Bakhta. EIP-1559: Fee market change for ETH 1.0 chain. https://eips.ethereum.org/EIPS/eip-1559, 2019. Ethereum Improvement Proposal. Accessed: 2026-11-16.

[10] Arka Rai Choudhuri, Sanjam Garg, Guru Vamsi Policharla, and Mingyuan Wang. Practical mempool privacy via one-time setup batched threshold encryption. In *34th USENIX Security Symposium (USENIX Security 25)*, pages 3477–3495, 2025.

[11] Derek J. Clark and Christian Riis. Contest success functions: An extension. *Economic Theory*, 11:201–204, 1998.

[12] CoW Protocol. CoW Protocol: Solvers. https://docs.cow.fi/cow-protocol/concepts/introduction/solvers. Accessed: 2026-11-16.

[13] Philip Daian, Steven Goldfeder, Tyler Kell, Yunqi Li, Xueyuan Zhao, Iddo Bentov, Lorenz Breidenbach, and Ari Juels. Flash boys 2.0: Frontrunning in decentralized exchanges, miner extractable value, and consensus instability. In *2020 IEEE symposium on security and privacy (SP)*, pages 910–927. IEEE, 2020.

[14] Ethereum Foundation. The merge. https://ethereum.org/en/roadmap/merge/, 2022. Accessed: 2026-10-30.

[15] Flashbots. MEV-Boost: Overview. https://docs.flashbots.net/flashbots-mev-boost/introduction, 2025. Last updated: 2026-05-06. Accessed: 2026-11-16.

[16] Elijah Fox, Mallesh Pai, and Max Resnick. Censorship resistance in on-chain auctions. *arXiv preprint arXiv:2301.13321*, 2023.

[17] Lioba Heimbach, Lucianna Kiffer, Christof Ferreira Torres, and Roger Wattenhofer. Ethereum's proposer-builder separation: Promises and realities. In *Proceedings of the 2023 ACM on Internet Measurement Conference*, pages 406–420, New York, NY, USA, 2023. Association for Computing Machinery.

[18] Mahimna Kelkar, Soubhik Deb, Sishan Long, Ari Juels, and Sreeram Kannan. Themis: Fast, strong order-fairness in byzantine consensus. In *Proceedings of the 2023 acm sigsac conference on computer and communications security*, pages 475–489, 2023.

[19] Mahimna Kelkar, Fan Zhang, Steven Goldfeder, and Ari Juels. Order-fairness for byzantine consensus. In *Annual International Cryptology Conference*, pages 451–480. Springer, 2020.

[20] Endong Liu, Mark Ryan, Liyi Zhou, and Pascal Berrang. Evasion under blockchain sanctions. In *Proceedings of the ACM Web Conference 2026*, pages 3507–3518, 2026.

[21] Yulin Liu, Yuxuan Lu, Kartik Nayak, Fan Zhang, Luyao Zhang, and Yinhong Zhao. Empirical analysis of EIP-1559: Transaction fees, waiting times, and consensus security. In *Proceedings of the 2022 ACM SIGSAC Conference on Computer and Communications Security*, pages 2099–2113. ACM, 2022.

[22] Xingyu Lyu, Mengya Zhang, Xiaokuan Zhang, Jianyu Niu, Yinqian Zhang, and Zhiqiang Lin. An empirical study on ethereum private transactions and the security implications. *arXiv preprint arXiv:2208.02858*, 2022.

[23] Patrick McCorry, Alexander Hicks, and Sarah Meiklejohn. Smart contracts for bribing miners. In *Financial Cryptography and Data Security: FC 2018 International Workshops, BITCOIN, VOTING, and WTSC, Revised Selected Papers*, volume 10958 of *Lecture Notes in Computer Science*, pages 3–18. Springer, 2019.

[24] Eunchan Park, Taeung Yoon, Hocheol Nam, Deepak Maram, and Min Suk Kang. On frontrunning risks in batch-order fair systems for blockchains. In *Proceedings of the 2025 ACM SIGSAC Conference on Computer and Communications Security*, pages 918–932, 2025.

[25] Kaihua Qin, Liyi Zhou, and Arthur Gervais. Quantifying blockchain extractable value: How dark is the forest? In *2022 IEEE Symposium on Security and Privacy (SP)*, pages 198–214. IEEE, 2022.

[26] Tom Robinson. The largest theft in history: North korea behind the $6.4 billion Binance hack. https://www.elliptic.co/insights/binance-largest-hack-in-history/, October 2026. Elliptic. Updated: November 1, 2026, 08:00 UTC.

[27] Jerome H. Saltzer and Michael D. Schroeder. The protection of information in computer systems. *Proceedings of the IEEE*, 63(9):1278–1308, September 1975.

[28] Stergios Skaperdas. Contest success functions. *Economic Theory*, 7:283–290, 1996.

[29] Thomas Thiery, Francesco D'Amato, Julian Ma, Barnabé Monnot, Terence Tsao, Jacob Kaufmann, and Jihoon Song. Eip-7805: Fork-choice enforced inclusion lists. https://eips.ethereum.org/EIPS/eip-7805, 2024. Ethereum Improvement Proposal.

[30] Christof Ferreira Torres, Ramiro Camino, et al. Frontrunner jones and the raiders of the dark forest: An empirical study of frontrunning on the ethereum blockchain. In *30th USENIX Security Symposium (USENIX Security 21)*, pages 1343–1359, 2021.

[31] TRM Labs. North korea steals USD 6.36 billion from Binance in the largest hack in history. https://www.trmlabs.com/resources/blog/north-korea-steals-usd-6-36-billion-from-binance-in-the-largest-hack-in-history, October 2026. October 29, 2026.

[32] Uniswap Labs. UniswapX overview. https://developers.uniswap.org/docs/liquidity/uniswapx/overview. Accessed: 2026-11-16.

[33] U.S. Department of the Treasury, Office of Foreign Assets Control. FAQ 646: How do I block digital currency? https://ofac.treasury.gov/faqs/646, 2021. October 15, 2021. Accessed: 2026-11-16.

[34] U.S. Department of the Treasury, Office of Foreign Assets Control. Specially designated nationals and blocked persons list (SDN). https://ofac.treasury.gov/specially-designated-nationals-and-blocked-persons-list-sdn-human-readable-lists, 2026. Accessed: 2026-11-10.

[35] Fabian Vogelsteller and Vitalik Buterin. Eip-20: Token standard. https://eips.ethereum.org/EIPS/eip-20, 2015. Ethereum Improvement Proposal.

[36] Anton Wahrstätter, Jens Ernstberger, Aviv Yaish, Liyi Zhou, Kaihua Qin, Taro Tsuchiya, Sebastian Steinhorst, Davor Svetinovic, Nicolas Christin, Mikolaj Barczentewicz, et al. Blockchain censorship. In *Proceedings of the ACM Web Conference 2024*, pages 1632–1643, 2024.

[37] Fredrik Winzer, Benjamin Herd, and Sebastian Faust. Temporary censorship attacks in the presence of rational miners. In *2019 IEEE European Symposium on Security and Privacy Workshops (EuroS&PW)*, pages 357–366. IEEE, 2019.

[38] Sen Yang, Kartik Nayak, and Fan Zhang. Decentralization of ethereum's builder market. In *2025 IEEE Symposium on Security and Privacy (SP)*, pages 1512–1530. IEEE, 2025.

[39] Sen Yang, Fan Zhang, Ken Huang, Xi Chen, Youwei Yang, and Feng Zhu. Sok: Mev countermeasures: Theory and practice. *arXiv preprint arXiv:2212.05111*, 2022.

[40] Mengqian Zhang, Yuhao Li, Xinyuan Sun, Elynn Chen, and Xi Chen. Maximal extractable value in batch auctions. In *Proceedings of the 26th ACM Conference on Economics and Computation*, pages 510–510, 2025.

[41] Francesco Zola, Jon Ander Medina, and Raúl Orduna. Assessing the impact of sanctions in the crypto ecosystem: Effective measures or ineffective deterrents? In *Computer Security. ESORICS 2024 International Workshops: DPM, CBT, and CyberICPS, Bydgoszcz, Poland, September 16–20, 2024, Revised Selected Papers, Part I*, page 292–308, Berlin, Heidelberg, 2024. Springer-Verlag.

# Appendix A. Technical Arguments for Section 4

## A.1. Additional Conventions

(1) **Restricted operations.** Before $p$ takes effect, the restricted party can always submit valid $x \in C(p)$ with nonce and funds independent of the authority's transactions. Its attached success-contingent payment is made only if $x$ succeeds and is not clawed back when $p$ later takes effect.
(2) **Protocol constraints.** Apart from authority priority, protocol constraints consist only of causal ordering (nonce or funding dependencies) and uniform inclusion obligations. A contract revert does not invalidate the outer transaction; capacity is nonbinding.
(3) **Payments.** Existing payment commitments have a finite total payable amount that does not automatically increase with new offers; additional funding and counteroffers are new messages. The restricted party can fully fund any finite payment using funds unavailable to existing commitments.
(4) **Rounds and network.** Cutoffs are public and exogenous; each finite interval has finitely many cutoffs, messages, and candidate producers. Submissions exactly at a cutoff enter the next round. By delay independence (Section 3.1), conditional on the causal past, the delays of finitely many subsequent messages, including adaptive ones, jointly fall in specified open intervals with positive probability.
(5) **Fee requirements.** These are fixed before the late offer or have a finite upper bound independent of it.

## A.2. Late Conditional Payments

**Lemma A.1** (Late conditional payment)**.** *Fix a round with cutoff $T$, $s \in (T-\Delta, T)$, $G > 0$, an event $E$ concerning its block, and a positive-probability event $H$ determined by the history before $s$. There exist a constant $K$ and a positive-probability event $H' \subseteq H$ such that, after the restricted party publicly broadcasts at $s$ a payment of $K+G$ conditional on $E$, every block satisfying $E$ pays at least $G$ more than any block not satisfying $E$ on $H'$.*

*Proof.* By Appendix A.1(3), all payments offered before $s$, plus the ordinary fees of the restricted party's transaction sent at $s$, have a finite sum. Thus some constant $K$ bounds it on a positive-probability subset of $H$. Since $T-s < \Delta$, Appendix A.1(4) allows us to further require that the restricted party's message sent at $s$ reach all candidate producers before $T$, while all other messages sent in $[s, T)$ arrive after $T$, retaining positive probability. On this event $H'$, blocks not satisfying $E$ earn at most $K$, whereas those satisfying $E$ earn at least $K+G$. □

## A.3. Proofs

*Proof of Theorem 4.1.* Fix an authority strategy, $0 \le \theta < \Delta$, and $G > 0$. In executions with the restricted party silent, consider $(T^- - \Delta, T - \theta)$ for each cutoff $T$ and its predecessor $T^-$. Since $\theta < \Delta$, these intervals cover the time axis. Countably many rounds and finite $t_c$ therefore allow a fixed round and a closed subinterval $[a, b]$ with $\Pr[t_c \in [a, b]] > 0$, giving $T^- < a + \Delta$ and $b + \theta < T$. Choose $s \in (\max\{b+\theta, T^-, T-\Delta\}, T)$. The restricted party remains silent before $s$, preserving the distribution of the preceding history. On $t_c \in [a, b]$, all authority messages submitted from $t_c$ onward are sent after $a$. By Appendix A.1(4), they all arrive after $T^-$ with positive conditional probability. Messages submitted before $t_c$ cannot suffice to enforce $p$ (Section 3.2), so $p$ has not taken effect in any round ending at or before $T^-$, nor at $s$. Choose a constant $c$ bounding the fee requirement (Appendix A.1(5)) on a positive-probability subevent $H$ of this event; $H$ is determined by the history before $s$.

At $s$, the restricted party publicly broadcasts $x$ with ordinary bid $c$ and an embedded payment conditional on $E = \{x \text{ succeeds in the round}\}$, of the amount given by Lemma A.1. The lemma yields a positive-probability event $H' \subseteq H$. On $H'$, $x$ is includable. By Appendix A.1(1)–(2), placing $x$ first, followed in dependency order by transactions required by uniform inclusion obligations, gives a block $B^*$ satisfying all protocol constraints except authority priority, in which $x$ succeeds before $p$. It pays at least $G$ more than any block in the round where $x$ does not succeed. By $I$, $x$ takes effect before $p$ in $B^*$, with $t_{\mathrm{sub}}(x) = s > t_c + \theta$. An ordering rule guaranteeing $P_\theta$ must almost surely select a block where $x$ does not succeed on $H'$, thereby forgoing $B^*$. □

*Proof of Theorem 4.3.* Number rounds from the first cutoff after $t_c$. The inclusion guarantee makes $p$ take effect within the block of round $d$ at the latest. Since round lengths are at most $\lambda$, cutoff $k$ is no later than $t_c + k\lambda$. Consider any restricted operation $x$ with $t_{\mathrm{sub}}(x) \ge t_c + \theta$. If $D < d$, then $\theta \ge (d - D)\lambda$, so $x$ enters no earlier than round $d - D + 1$: under immediate execution it executes after $p$; for $1 \le D < d$, it takes effect no earlier than the end of round $d$. If $D \ge d$, it takes effect no earlier than the end of round $D$.

In every case, $p$ takes effect first, rejecting $x$ or invalidating its request. Thus $P_\theta$ holds. □

*Proof of Proposition 1.* Exclusion rewards count at face value toward their round's revenue. Fix an authority strategy and $\theta \geq 0$. As in the proof of Theorem 4.1 with $\theta = 0$, choose the first attack round, $[a, b]$, and a time after $b$ within that round such that $t_c \in [a, b]$ and $p$ has not taken effect earlier with positive probability. Starting in that round, at a fixed time after the preceding cutoff and less than $\Delta$ before each cutoff, apply Lemma A.1 with $G = 1$ as follows. For cutoffs at or before $b + \theta$, take $E = \{p \text{ has not taken effect by round end}\}$. In the first round ending after $b+\theta$, submit an includable $x$ after $b+\theta$, bidding as in the proof of Theorem 4.1. Under immediate execution, take $E = \{x \text{ succeeds}\}$; under delayed effects, let $E$ require successful registration of $x$ and no enforcement by round end, then require no enforcement by round end in subsequent rounds until $x$ takes effect at the end of round $D$.

By the proposition's omission assumption, Appendix A.1(1)–(2), and $N$, every round admits a feasible block satisfying $E$, so every revenue-maximizing block satisfies $E$. Under immediate execution, $I$ makes $x$ take effect before $p$. Finitely many successive applications of Lemma A.1 give a joint event of positive probability. Since $t_{\text{sub}}(x) > b + \theta \geq t_c + \theta$ and $\theta$ is arbitrary, $w = \infty$. □

# Appendix B. Proofs for the Freeze Contest

**Notation.** By Proposition 4, $\sigma_f = (1 - \nu)\tau_f$. We use the following notation and expected payoffs:

$$\tau_f := \frac{\beta_f^r}{\beta_f^r + \beta_x^r},\ \tau_x := 1 - \tau_f,\ q' := (1 - \nu)q,\ \mu' := (1 - \nu)\mu,$$
$$U_f = \Psi(1 - q) + q'\Psi\tau_f - \beta_f, \quad U_x = (1 - \mu'\tau_f)(V - \beta_x),$$
$$z := \beta_f/\beta_x, \quad y := z^r, \quad k_1 := 2 - \mu' + \mu' r,$$
$$k_2 := 1 - \mu', \quad k_3 := (1 - \mu')(1 - r),$$
$$Q := (1 + y)^2 - \mu' y(1 + y - r), \quad \Phi(z) := \mu' z(1 + y)^2 Q^{-1},$$
$$Q = 1 + k_1 y + k_2 y^2 = (1 + y)(1 + k_2 y) + \mu' r y.$$

Thus $\tau_f = y/(1 + y)$, $\tau_f\tau_x = y/(1 + y)^2$, $0 \leq k_3 \leq k_2 < 1 \leq k_1$, and $1 + y \leq Q \leq (1 + y)^2$.

## B.1. Contest Characterization

*Proof of Proposition 4.* Suppose both bids are positive. *Necessity.* By condition (ii), strictly increasing functions $G_f, G_x$, both zero at zero, satisfy $\tau_f/\tau_x = G_f(\beta_f)/G_x(\beta_x)$. By (iv) and (i), $G_f(u)/G_x(v) = G_x(u)/G_f(v)$ for all $u, v > 0$; setting $u = v$ gives $G_f = G_x =: G$. By (iii), $G(tu)/G(tv) = G(u)/G(v)$, so $\varphi(t) := G(tu)/G(u)$ is independent of $u$ and satisfies $\varphi(st) = \varphi(s)\varphi(t)$. Since $\varphi : (0, \infty) \to (0, \infty)$ is strictly increasing and multiplicative, $\varphi(t) = t^r$ for some $r > 0$ [5]. Thus $G(u) = G(1)u^r$ and $\tau_f/\tau_x = (\beta_f/\beta_x)^r$; condition (i) gives $\tau_f = \beta_f^r/(\beta_f^r + \beta_x^r)$. At $\beta_f/\beta_x = e$, the odds are $e^r$, uniquely determining $r$. *Sufficiency* follows by direct verification. Finally, the issuer's odds $\beta_f^r/\beta_x^r$ are concave in $\beta_f$ if and only if $r \leq 1$, as are the restricted party's odds in $\beta_x$. This is condition (v). □

## B.2. Equilibrium

*Proof of Lemma 7.1.* Since $\nu < 1$, $q', \mu' > 0$.

*Interiority.* The bid $\beta_x = V$ gives $U_x = 0$ and is not a best response. At $\beta_x = 0$, $\sigma_x = 0$, so the issuer's best bid is zero. At $\beta_f = 0$, however, every $\beta_x > 0$ gives $\tau_f = 0$; $U_x = V - \beta_x$ never attains its supremum $V$, while $\beta_x = 0$ yields only $(1 - \mu)V < V$. The restricted party has no best response. An equilibrium therefore requires $\beta_f > 0$ and $0 < \beta_x < V$.

*Sufficiency of first-order conditions.* For $r \leq 1$, $U_f$ is strictly concave in $\beta_f$. In $\beta_x$, $A := 1 - \mu'\tau_f = 1 - \mu' + \mu'\tau_x$ satisfies $A > 0$, $A' > 0$, and $A'' \leq 0$, so $U_x'' = A''(V - \beta_x) - 2A' < 0$. As $\beta_x \to 0^+$, $U_x \to (1 - \mu')V \geq (1 - \mu)V$; an interior stationary point therefore dominates $\beta_x = 0$ and is a global best response.

*Solution.* Differentiating Equation (2) using (1) gives the first-order conditions and equivalent bid expressions

$$\beta_f = q'\Psi r\tau_f\tau_x, \quad \mu' r\tau_f\tau_x(V - \beta_x) = (1 - \mu'\tau_f)\beta_x,$$
$$\beta_f = q'\Psi r y(1 + y)^{-2}, \quad \beta_x = \mu' r y V Q^{-1}.$$

Both bid expressions depend only on $z$. Their ratio makes the first-order conditions equivalent to $\Phi(z) = q'\Psi/V$, or

$$z(1 + y)^2 Q^{-1} = q\Psi(\mu V)^{-1}. \tag{3}$$

*Uniqueness.* Differentiation gives

$$E := \frac{d \ln \Phi}{d \ln z} = 1 + \frac{\mu' r y[(1 - r) + (1 + r)y]}{(1 + y)Q} \geq 1.$$

Thus $\Phi$ increases strictly from zero to infinity, and Equation (3) has a unique solution $z^*$. It yields a unique $(\beta_f^*, \beta_x^*)$, with $\beta_f^* > 0$ and $0 < \beta_x^* < V$, since $\mu' r y < Q$. □

Equation (3) also implies that $z^*$ increases strictly in $\Psi$ and $q$, decreases strictly in $V$, and satisfies $d \ln z^*/d \ln \Psi = d \ln z^*/d \ln q = 1/E \leq 1$.

## B.3. Participation

*Proof of Proposition 5.* Let $q_0$ be the probability that a participant co-occurs with $f$, and let $\ell$ be the participation rate among restricted parties with balance $V$. Then $q = \ell q_0$. A nonparticipant earns $(1-\mu)V$, whereas a participant earns $U_x^* - c$. By Equation (2), the expected gain before $c$ is

$$g(V; q) := (1 - \mu\sigma_f^*)(V - \beta_x^*) - (1 - \mu)V.$$

Equilibrium requires $g \leq c$ at $\ell = 0$, $g \geq c$ at $\ell = 1$, and $g = c$ for $0 < \ell < 1$.

(a) A freeze has been issued with probability $\mu$, and the transfer then takes effect first with probability $\sigma_x^*$, proving the first claim. Substitution gives

$$g/V = \mu\nu + \mu' D(y^*), \quad D(y) := (1 + k_3 y)/Q.$$

Since $k_3 < k_1$, $0 < D < 1$, so $\mu\nu V < g < \mu V$. *Dependence on $q$ and $\Psi$.* Since $k_3 < k_1$ and $k_2, k_3 \geq 0$, $D'(y) < 0$. Since $y^*$ increases strictly in $q$ and $\Psi$, $g$ decreases strictly in both. As $q \to 0$, $y^* \to 0$ and $g \to \mu V$; accordingly, define $g(V;0) := \mu V$.

*Dependence on $V$.* By Equation (3),

$$\mu' V D = q'\Psi(1 + k_3 y)[z(1+y)^2]^{-1}.$$

The final factor's logarithmic derivative with respect to $z$ is $rk_3 y/(1 + k_3 y) - 1 - 2ry/(1 + y) < r - 1 \leq 0$. Since $z^*$ decreases strictly in $V$, $g$ increases strictly in $V$. *Dependence on $\nu$.* Fix $q, \mu, V, \Psi, r$ and vary $\nu \in [0, 1)$ through $\mu' = (1 - \nu)\mu$, with $q' = (q/\mu)\mu'$. Regard $z^*$ as a function of $\mu'$. Logarithmically differentiating Equation (3), $\sigma_f^* = (\mu'/\mu)y/(1 + y)$, and $\beta_x^* = \mu' r V y/Q$ gives

$$\begin{aligned}
\frac{d\ln z^*}{d\ln \mu'} &= -\frac{\mu' y(1+y-r)}{QE} < 0,\\
\frac{d\ln \sigma_f^*}{d\ln \mu'} &= 1 - \frac{r\mu' y(1+y-r)}{(1+y)QE},\\
\frac{d\ln \beta_x^*}{d\ln \mu'} &= 1 + \frac{\mu' y(1+y-r)}{Q}\left[1 - \frac{r(1-k_2 y^2)}{QE}\right].
\end{aligned}$$

The derivative of $\ln \sigma_f^*$ is positive because $r\mu' y(1+y-r) \leq y(1+y) < (1+y)^2 \leq (1+y)QE$; that of $\ln \beta_x^*$ is at least one because $r(1 - k_2 y^2) \leq 1 \leq QE$. Thus $\sigma_f^*$ and $\beta_x^*$ increase strictly in $\mu'$, hence decrease strictly in $\nu$, while $\sigma_x^*$ increases strictly in $\nu$. Consequently, both $1 - \mu\sigma_f^*$ and $V - \beta_x^*$ increase strictly in $\nu$, as does $g$.

(b) The function $g$ is continuous and strictly increasing in $V$. As $V \to 0$, $g < \mu V \to 0$. As $V \to \infty$, $g \to \infty$: for $\nu > 0$, use the term $\mu\nu V$; for $\nu = 0$, use the preceding expression and $z^* \to 0$. Thus $g(\bar{V}; q_0) = c$ has a unique solution $\bar{V}$. The bounds $\mu\nu\bar{V} < c < \mu\bar{V}$ give $\bar{V} > c/\mu$ and, for $\nu > 0$, $\bar{V} < c/(\mu\nu)$. If $V \leq c/\mu$, then $g \leq \mu V \leq c$ for every $q$, strictly for $q > 0$, so $\ell = 0$. If $V \geq \bar{V}$, then for every $q < q_0$, $g(V; q) > g(V; q_0) \geq c$, so $\ell = 1$. If $c/\mu < V < \bar{V}$, then $g(V; 0) > c > g(V; q_0)$. Continuity and strict monotonicity in $q$ give a unique $\ell^* \in (0, 1)$ satisfying $g(V; \ell^* q_0) = c$, the only participation rate satisfying the equilibrium conditions. Since $g$ increases with $V$ and decreases with $q$, $\ell^*$ increases with $V$.

(c) Since $g$ increases strictly in $V$, $\bar{V}$ increases with $c$. Since $g$ decreases strictly in $\Psi$ and increases strictly in $\nu$, $\bar{V}$ increases with $\Psi$ and decreases with $\nu$. If an execution provider charges a fraction $\varphi \in (0, 1)$, the gain after delegation, $(1 - \varphi)g$, is positive for every $V > 0$. □

## B.4. Bounded Stakes

*Proof of Proposition 6.* (a) By $Q \leq (1 + y)^2$ and (3), $z^* \leq q\Psi/(\mu V)$, so

$$\tau_x^* = (1 + y^*)^{-1} \geq (1 + (q\Psi/(\mu V))^r)^{-1}.$$

Substitute into $\sigma_x^* = \nu + (1 - \nu)\tau_x^*$. (c) Since $\tau_f \tau_x \leq 1/4$ and $y < Q$, $\beta_f^* \leq q' r\Psi/4$ and $\beta_x^* = \mu' r y V/Q < \mu' r V$.

(b) By the comparative statics in $\nu$ established in Appendix B.3(a), $\sigma_x^*$ increases strictly in $\nu$. As $\nu \to 1$, $\sigma_f^* \leq 1 - \nu \to 0$; by part (c), both bids vanish, and the probability that the freeze takes effect first, $1 - q\sigma_x^*$, approaches $1 - q$. □

**Lemma B.1.** *If $\Psi \geq 2\mu' r V$, then $z^* \geq q' r$.*

*Proof.* Since $Q \geq 1 + y$, $\Phi(z) \leq \mu' z(1 + y)$. Setting $z_0 := q' r \leq 1$ gives $\Phi(z_0) \leq 2\mu' q' r \leq q'\Psi/V = \Phi(z^*)$; monotonicity of $\Phi$ yields $z^* \geq z_0$. □

## B.5. Value of Ordering Control

*Proof of Proposition 7.* Fix $A$. At issuer bid $b$, let $h_A(b)$ be the interception probability conditional on co-occurrence. Then $h_A(b) \leq \sigma_x(b, \beta_x^*) \leq 1$ and $h_A(\beta_f^*) = \rho\sigma_x^*$. Let $U_0, U_A$ denote payoffs without and with $A$, before maintenance costs. Bounded compensation gives

$$q h_A(b)(\Psi - \beta_x^*) \leq U_A(b) - U_0(b) \leq q h_A(b)\Psi.$$

Define $G_A := \sup_{b \geq 0} U_A(b) - U_0(\beta_f^*)$. Evaluating at $\beta_f^*$ gives the lower bound; baseline optimality and $h_A(b) \leq 1$ give the upper bound. The cost and compensation claims follow from $\Gamma > 0$ and $\beta_x^* \leq V$. No step uses the contest form or fixes the interception rate across bids. □

## B.6. Stronger Incentives for Ordering Control

*Proof of Proposition 8.* For $\Psi \geq \mu r V$, Proposition 6(c) gives $\Psi - \beta_x^* > \Psi - \mu' r V \geq 0$. Since $\sigma_x^* > 0$, $\Gamma > 0$. (a) By Appendix B.3(a), $\sigma_x^*$ increases strictly in $\nu$, and $\beta_x^*$ decreases strictly. Both factors in $\Gamma$ are positive, so $\Gamma$ increases strictly in $\nu$. As $\nu \to 1$, $\sigma_x^* \to 1$ and $\beta_x^* \to 0$, hence $\Gamma \to q\Psi$.

(b) Suppose $\Psi \geq 2\mu r V$, hence $\Psi \geq 2\mu' r V$. Write $\Gamma = q\nu(\Psi - \beta_x^*) + \Gamma_s$, where $\Gamma_s := q'\tau_x^*(\Psi - \beta_x^*)$. *First term.* Since $Q \geq 1 + y$, $d(y/Q)/d\ln z \leq r/4$. Also, $d\ln z^*/d\ln\Psi = d\ln z^*/d\ln q = 1/E \in (0, 1]$. Thus the derivatives of $\beta_x^* = \mu' r V y/Q$ with respect to $\ln\Psi$ and $\ln q$ are at most $\mu' r^2 V/4$. The first term's derivatives with respect to $\Psi$ and $q$ are therefore bounded below, respectively, by

$$q\nu\left[1 - \mu' r^2 V(4\Psi)^{-1}\right] \geq 0, \quad \nu\left[\Psi - \mu' r V - \mu' r^2 V/4\right] \geq 0.$$

*Second term: dependence on $\Psi$.* Fix $q$ and substitute (3) to obtain $\Gamma_s$ with $z = z^*$:

$$\Gamma_s = \mu' V\left\{z(1+y)Q^{-1} - q' r y[(1+y)Q]^{-1}\right\},$$

$d\Gamma_s/dz$ has the sign of

$$(1+y)Q - ry[k_2(1+y)^2 + \mu' r] - \frac{q' r^2 y}{z}\frac{1 - (k_1 + k_2)y^2 - 2k_2 y^3}{(1+y)^2}.$$

By Lemma B.1, $q' r^2 y/z \leq ry$, and the following factor is at most one. Using $(1 + y)Q = (1 + y)^2(1 + k_2 y) + \mu' r y(1 + y)$, the expression is at least

$$(1+y)^2(1+k_3 y) + \mu' r y(1+y-r) - ry > 0.$$

Since $z^*$ increases strictly in $\Psi$, so does $\Gamma_s$. *Second term: dependence on $q$.* Fix $\Psi$ and let $\omega := \mu' r V/\Psi \le 1/2$. Substituting Equation (3) gives $\Gamma_s$ below; using the preceding bounds on $Q$ and $\omega$, the logarithmic derivatives of its two $z$-dependent factors satisfy

$$\Gamma_s = \mu' V z(1+y) Q^{-1}\left(1-\omega y Q^{-1}\right),$$

$$L_1 = \frac{(1+y)^2(1+k_3 y)+\mu' r y(1+y-r)}{(1+y)Q} \ge \frac{1+y}{Q} \ge \frac{1}{1+y},$$

$$L_2 = -\frac{\omega r y(1-k_2 y^2)}{Q(Q-\omega y)} \ge -\frac{(r/2)y}{(1+y)(1+y/2)},$$

$$L_1 + L_2 \ge (1+y)^{-1}\left[1 - r y (1+y)^{-1}\right] > 0.$$

Since $z^*$ increases strictly in $q$, so does $\Gamma_s$. $\square$