Exact False-Name and One-User Coalition Incentives with Linear Miner Revenue

September 1, 2026

Abstract

Let $h \geq 1$ and let D be a distribution on nonnegative values whose public finite description permits exact sampling in expected polynomial time and deterministic polynomial-time computation of a positive rational median m_D and the exact rational probabilities $\Pr[V > m_D]$ and $\Pr[V = m_D]$. In the infinite-capacity ideal-MPC model, we construct a weakly anonymous randomized transaction fee mechanism with exact nonnegative rational payments and miner revenue. Pointwise, an unconfirmed bid pays zero, a confirmed bid pays at most its report, and miner revenue is at most total payments. On finitely encoded bids comparable with m_D in polynomial time, the mechanism runs in expected time polynomial in the public-description length, $\log h$, and the total bid-encoding length. It is ex post false-name user incentive compatible for every fixed outsider bid vector and every finite randomized deviation. For every miner coalition share $\rho \in [0, 1]$, it is exact Bayesian miner incentive compatible with at least h independent truthful draws from D and exact Bayesian one-user side-contract-proof with at least h such honest outsiders and one arbitrary-valued colluding user, against every finite randomized deviation selected before the honest draws and the mechanism's private coins. For a truthful bid vector $V \sim D^\ell$ and every $\ell \geq h$,

$$\mathbb{E}[\mu] = \begin{cases} m_D/2, & 1 \leq h < 64, \\ hm_D/128, & h \geq 64. \end{cases}$$

The same piecewise expression is therefore the infimum of expected revenue over all $\ell \geq h$, and it is always at least $hm_D/128$.

Note. This paper was fully AI generated through a harness created by Richard Peng that discovers open problems and runs them through an automated research pipeline.

Contents

1 Introduction

A transaction fee mechanism (TFM) decides which submitted transactions are confirmed, what each confirmed user pays, and how much revenue reaches the miner. These choices connect congestion and delay in fee markets with the long-run security role of transaction fees as block subsidies decline [HLM21, CKWN16]. The miner, however, is not a trusted auctioneer: it may submit fake transactions or coordinate with a user. Modern transaction fee mechanism design therefore treats user truthfulness, miner incentives, and collusive deviations as simultaneous design constraints [Rou24].

We study these constraints with infinite capacity and ideal multiparty computation (MPC). Infinite capacity removes competition for block space but not identity manipulation. A strategic user may submit any finite number of identities, although at most one identity represents its value-bearing transaction; payments made by the remaining identities are still real costs. We require false-name user incentive compatibility (UIC) ex post, for every fixed collection of outsider bids. Exact Bayesian miner incentive compatibility (MIC) and one-user side-contract-proofness (1-SCP) require that a finite false-name deviation selected before the honest values and the mechanism's private coins are observed offer no expected gain. Here exact means zero additive incentive slack. We also ask for pointwise payment and budget feasibility and a positive truthful-revenue guarantee uniform over every population size $\ell \geq h$, when all ℓ users draw values independently from a public distribution D and bid truthfully.

This combination of requirements is sensitive to the implementation and participation assumptions. In the plain one-shot model, Chung and Shi show that UIC together with 1-SCP forces zero miner revenue, even with infinite block space [CS23]. Shi, Chung, and Wu introduce the MPC-assisted TFM model and show that, in their universal-participation, standard-utility formulation, exact simultaneous user, miner, and one-user coalition incentives still force zero expected revenue when no positive lower bound on honest participation is known [SCW23]. A known lower bound h, together with a public value distribution, changes this premise.

Known-h MPC-assisted mechanisms. Wu, Shi, and Chung study the same broad known-h, infinite-block, MPC-assisted setting with a public D and fake bids [WSC24]. For $h = 1$, their parity posted-price mechanism has exact incentives against unrestricted finite strategic padding and earns $m_D/2$, where m_D is a median of D. For general h, their threshold construction permits arbitrary strategic bids and earns $\Theta(hm_D)$, with exponentially small rather than zero Bayesian incentive slack. Their exact linear-programming (LP) construction gives expected revenue $hm_D/4$ when the total number d of coalition bids satisfies

$$d \leq \frac{1}{8}\sqrt{\frac{h}{2\log h}}.$$

We address the exact general-h regime with no a priori numerical bound on a finite false-name vector. This comes with a coalition tradeoff: our theorem protects against one arbitrary-valued colluding user, whereas the approximate and bounded-bid results above accommodate larger user coalitions in their respective regimes. Our small-h branch uses the parity construction; the distinct ingredient is the large-h mechanism.

Our contribution. For every $h \geq 1$ and every explicitly described distribution D satisfying the exact-computation assumptions below, we construct a weakly anonymous infinite-capacity ideal-MPC mechanism family $\Pi_{h,D}$. It is implemented by ParityMechanism for $1 \leq h < 64$ and by SybilPotentialMechanism for $h \geq 64$. On finitely encoded bids comparable with m_D, these

mechanisms output confirmation bits, exact nonnegative rational payments, and exact nonnegative rational miner revenue in expected time polynomial in the public-description length, $\log h$, and the total bid-encoding length. They satisfy the pointwise payment and budget constraints and ex post false-name UIC. For every miner coalition share $\rho \in [0,1]$, they satisfy exact Bayesian MIC in the presence of at least h independent truthful users from D and exact Bayesian 1-SCP in the presence of at least h such outsiders and one arbitrary-valued colluding user. These guarantees cover every randomized finite-identity deviation chosen before the honest draws and private coins. Moreover, for every $\ell \geq h$, on a truthful bid vector of ℓ independent draws from D, expected miner revenue is exactly $m_D/2$ when $h < 64$ and exactly $hm_D/128$ when $h \geq 64$. The full statement uses the explicit-distribution and deviation conventions formalized in Definitions 2.1 and 2.2.

Theorem 1.1 (Exact false-name and coalition incentives)**.** *Let $h \geq 1$ be an integer. Let D be a publicly described distribution on nonnegative values such that its finite description permits exact sampling in expected polynomial time and deterministic polynomial-time computation of a positive rational median $m = m_D$ and the exact rational numbers*

$$q_> = \Pr_{V\sim D}[V > m], \qquad q_= = \Pr_{V\sim D}[V = m],$$

all with polynomially bounded bit length. There is an infinite-capacity ideal-MPC mechanism family $\Pi_{h,D}$, implemented by ParityMechanism *when $1 \leq h < 64$ and by* SybilPotentialMechanism *when $h \geq 64$, with the following properties.*

1. *It is weakly anonymous and, on finitely encoded bids that can be compared with m in polynomial time, runs in expected time polynomial in the public-description length, $\log h$, and the total bid-encoding length.*

2. *On every finite bid vector it outputs confirmation bits $x_i \in \{0,1\}$, nonnegative rational payments p_i, and nonnegative rational miner revenue μ satisfying, pointwise,*

$$x_i = 0 \Longrightarrow p_i = 0, \qquad x_i = 1 \Longrightarrow p_i \leq b_i, \qquad \mu \leq \sum_i p_i.$$

3. *It is ex post false-name user incentive compatible for every fixed outsider bid vector and every randomized deviation with finitely many fake identities. For each miner coalition share $\rho \in [0,1]$, it is exact Bayesian miner incentive compatible in the presence of at least h independent truthful users from D, and exact Bayesian 1-side-contract-proof in the presence of at least h independent truthful outsiders from D and one colluding user of arbitrary value. Bayesian deviations may be randomized and use finitely many fake identities, but are chosen without observing honest draws and before private mechanism coins.*

4. *For every integer $\ell \geq h$,*

$$\mathbb{E}_{V\sim D^\ell,\,\Pi_{h,D}}[\mu] = \begin{cases} m_D/2, & 1 \leq h < 64, \\ hm_D/128, & h \geq 64. \end{cases}$$

Consequently, the infimum is determined exactly:

$$\inf_{\ell \geq h} \mathbb{E}_{V\sim D^\ell,\,\Pi_{h,D}}[\mu] = \begin{cases} m_D/2, & 1 \leq h < 64, \\ hm_D/128, & h \geq 64, \end{cases}.$$

In particular,

$$\inf_{\ell \geq h} \mathbb{E}_{V\sim D^\ell,\,\Pi_{h,D}}[\mu] \geq \frac{hm_D}{128}.$$

Ideas of the construction. We first randomize at a possible atom of D at its median so that each honest bid produces an independent fair eligibility bit. In the small-h regime, the parity of these bits masks every fixed finite collection of strategic bits. For large h, miner revenue is supported on a central binomial window and normalized so that truthful expected revenue is exactly $hm_D/128$ for every honest count $\ell \geq h$. The difficulty is that a deviator can change the submitted length: eligible fake identities incur payments, but ineligible ones can pad the input for free. The large-h mechanism therefore gives an ineligible bid a small free-confirmation probability that never increases with the submitted length. The same length-dependent parameter controls the payment discount for an eligible bid. Thus the designated user's direct utility never increases under padding; once the honest-count premise applies, an additional identity either lowers the relevant free-confirmation chance or raises the eligible real bid's payment. This utility loss is a *Sybil potential* that covers the largest possible revenue gain from padding. A binomial comparison lemma and an exact rational-coin implementation make this argument uniform over every finite deviation without introducing incentive slack.

Restricted and reformulated positive regimes. Other positive-revenue results alter a different part of the model. Gafni and Yaish's infinite-capacity Good Toy Auction on the discrete type space $\{0\} \cup \mathbb{N}$ satisfies their dominant-strategy incentive compatibility (DSIC), myopic miner incentive compatibility (MMIC), SCP, and off-chain-agreement proofness (OCA-proofness) notions [GY24]. Its DSIC definition permits one report and is not the present all-finite false-name guarantee. Chen, Simchi-Levi, Zhao, and Zhou study bounded i.i.d. values, finitely many users and slots, Bayesian user incentives, one-user collusion, and asymptotic revenue approximation, with fake-transaction concerns handled through additional sequencing and commitment structure [CSLZZ25]. Tang and Yao instead obtain positive results in a proof-of-stake model with modified intertemporal utilities and further bid-grid and parameter restrictions [TY23]. These are restricted or reformulated versions of the positive-revenue problem, not weaker theorem statements under the present definitions.

Organization. Section 2 specifies the computational and incentive model and develops exact median balancing. Section 3 explains the mechanisms and the incentive accounting. Section 4 proves the parity guarantees for small h, then establishes the parameters, revenue, and exact incentive properties of the large-h construction.

2 Preliminaries

Notation and randomness. For an event $\mathcal{A}$, let $\mathbf{1}_{\mathcal{A}}$ denote its indicator. For $\alpha \in [0,1]$, Bernoulli(α) denotes a bit that is one with probability α. For $n \geq 0$, $\mathrm{Bin}(n, 1/2)$ denotes the law of a sum of n independent fair bits, and D^ℓ denotes the product law of ℓ independent draws from D. Unless stated otherwise, distinct primitive randomized decisions use fresh, mutually independent private coins.

2.1 Input and incentive model

We specify the computational input and the deviations quantified by the incentive properties.

Definition 2.1 (Explicit distribution)**.** *The public description of a distribution D on $\mathbb{R}_{\geq 0}$ is a finite binary string of length L. It permits exact sampling from D in expected time polynomial in L and, deterministically in time polynomial in L, outputs a positive rational median $m = m_D$ together with the exact rational values $q_> = \Pr[V > m]$ and $q_= = \Pr[V = m]$. Their numerators and denominators have bit length polynomial in L. Here a median satisfies $\Pr[V \leq m] \geq 1/2$*

and $\Pr[V \geq m] \geq 1/2$. *The mathematical mechanism accepts arbitrary nonnegative real bids. Its computational guarantee applies to finite bid encodings for which comparison with* m *is polynomial-time; nonnegative binary rationals are a canonical case. Input size includes* L*, the binary encoding of* h*, and the entire submitted bid vector.*

Definition 2.2 (Ideal-MPC false-name model)**.** *On* $b = (b_1, \ldots, b_N)$*, a mechanism outputs confirmations* $x_i \in \{0, 1\}$*, payments* $p_i \geq 0$*, and total miner revenue* $\mu \geq 0$*. Every monetary output is an exact rational represented by a binary numerator and denominator. Pointwise feasibility means* $p_i = 0$ *if* $x_i = 0$*,* $p_i \leq b_i$ *if* $x_i = 1$*, and* $\mu \leq \sum_i p_i$*. Infinite capacity means that the mechanism is defined on every finite bid vector and imposes no additional bound on the number of confirmations.*

A strategic user of value $v \geq 0$ *may submit any finite vector of identities and, before mechanism coins are drawn, designate at most one as its real transaction. Only the designated identity has intrinsic value. If* I *is the set of submitted identities and* $j \in I$ *is designated, realized user utility is*

$$vx_j - \sum_{i \in I} p_i;$$

if no identity is designated, the value term is absent. For a coalition controlling a fraction $\rho \in [0, 1]$ *of the miners, let* I_C *contain every identity submitted by any coalition member. If its user designates* $j \in I_C$*, realized coalition utility is*

$$vx_j - \sum_{i \in I_C} p_i + \rho\mu.$$

The value term is omitted when no identity is designated, including for a miner-only coalition. Honest user behavior is one truthful bid designated as the user's real transaction, honest miners submit no identities, and aborting has utility zero. The ideal functionality enforces the outcome rule above. Accordingly, miner deviations in this model consist of submitting a finite vector of identities or aborting; they do not include censoring honest inputs, observing those inputs before submission, or altering the realized outcome.

Ex post false-name user incentive compatibility compares truthful single-identity participation with every finite false-name strategy for every fixed vector of outsider bids, taking expectation only over private mechanism coins. For Bayesian miner incentive compatibility, at least h *honest users have independent values from* D*. For Bayesian* 1*-side-contract-proofness, at least* h *honest outsiders have independent values from* D*, and one colluding user has an arbitrary value. A Bayesian deviation may depend on the colluding user's value but is selected without observing the honest draws and before private mechanism coins are generated. Randomized deviations are allowed. Bayesian expectations additionally average over the independent honest draws. "Exact" means that the corresponding incentive inequality has zero additive slack. Weak anonymity means that permuting the bid vector induces the same permutation of the joint law of confirmations and payments and leaves the law of* μ *unchanged.*

2.2 Exact coins and balanced eligibility

For a rational $\alpha = a/s \in [0, 1]$, with integers $0 \leq a \leq s$ and $s \geq 1$, the desired object is a bit whose success probability is exactly α. The following rejection sampler implements this object using fair bits.

RationalCoin: Exactly sample a rational Bernoulli variable.
Input: A rational number $\alpha = a/s \in [0,1]$, represented by binary integers $0 \le a \le s$ and $s \ge 1$.
Output: A bit Z.
Guarantee: $\Pr[Z = 1] = \alpha$ exactly.

1. Set $k = \lceil \log_2 s \rceil$.
2. Draw k fair bits as a uniform integer $R \in \{0, \dots, 2^k - 1\}$. If $R \ge s$, repeat this step.
3. Output $Z = \mathbf{1}_{\{R<a\}}$.

Lemma 2.3 (Exact rational-coin sampling)**.** RationalCoin *outputs one with probability exactly* α *and uses expected* $O(1 + \log s)$ *bit operations and fair random bits.*

Proof. Conditioned on acceptance, R is uniform on $\{0, \dots, s-1\}$, exactly a of whose s elements are below a. Thus the output probability is a/s. If $s > 1$, then $2^{k-1} < s \le 2^k$, so each iteration is accepted with probability greater than $1/2$; for $s = 1$ it is accepted immediately. The expected number of iterations is therefore less than two, and an iteration uses $O(1 + \log s)$ random bits and bit operations. □

The sampler permits exact randomization at an atom of the median. This is the only place where the exact probabilities in the public description of D are needed.

Define

$$\tau = \begin{cases} \dfrac{1/2 - q_>}{q_=}, & q_= > 0, \\ 0, & q_= = 0. \end{cases} \tag{1}$$

Exact rational arithmetic on the public values $q_>$ and $q_=$ computes τ in polynomial time and gives it polynomial bit length. For each submitted bid b, its eligibility bit $E(b)$ is sampled independently according to

$$E(b) = \begin{cases} 0, & b < m, \\ \text{Bernoulli}(\tau), & b = m, \\ 1, & b > m. \end{cases} \tag{2}$$

The middle case is implemented by the one-argument call RationalCoin(τ).

Lemma 2.4 (Exact median balancing)**.** *The number* τ *lies in* $[0,1]$*. If* $V \sim D$*, then* $E(V)$ *is a fair bit. Eligibility bits produced from independent truthful values are mutually independent.*

Proof. The median property gives $q_> \le 1/2 \le q_> + q_=$. If $q_= > 0$, these inequalities imply $0 \le 1/2 - q_> \le q_=$, so $\tau \in [0,1]$ and

$$\Pr[E(V) = 1] = q_> + \tau q_= = \frac{1}{2}.$$

If $q_= = 0$, the median inequalities force $q_> = 1/2$, and the same identity holds. Independence follows from independence of the values and the fresh coin used for each identity. □

3 Overview

The construction has a common fair-bit reduction followed by two mechanisms. The reduction converts every honest value into an independent fair eligibility bit without assuming that the distribution is continuous. For $h < 64$, a parity rule then gives the required revenue and, in the presence of an honest fair bit, makes strategic reports unable to affect its expectation. For $h \geq 64$, parity no longer gives revenue linear in h; instead we normalize revenue on a central binomial window. The main difficulty in this second branch is that an ineligible false-name bid pays nothing but changes the submitted length, and can therefore change the revenue distribution. A decreasing free-confirmation probability acts as a *Sybil potential*: the possible revenue gain from such padding is matched by a loss in the colluding user's direct utility.

The common fair-bit reduction. Let $m = m_D$ and write $q_> = \Pr[V > m]$ and $q_= = \Pr[V = m]$. At an atom of the median, the deterministic threshold test need not be balanced, so we split that atom using

$$\tau = \frac{1/2 - q_>}{q_=} \quad \text{when } q_= > 0, \qquad \tau = 0 \quad \text{when } q_= = 0.$$

The eligibility rule is zero below m, one above m, and a Bernoulli(τ) bit at m. The median inequalities ensure $\tau \in [0, 1]$, and for an honest draw

$$\Pr[E(V) = 1] = q_> + \tau q_= = \frac{1}{2}.$$

Thus independent truthful bids yield independent fair bits. The exact public median data are used here essentially: they let RationalCoin implement the split with no approximation error; a sampler-only interface does not in general provide the exact median masses used in this step. The same threshold structure also makes the sign of a user's gain from eligibility agree with the sign of $v - m$.

The small-h branch. For $1 \leq h < 64$, ParityMechanism takes the public pair (h, D) and a finite bid vector, forms the eligibility bits, and outputs

$$x_i = E_i, \qquad p_i = mE_i, \qquad \mu = m\mathbf{1}_{\{\sum_i E_i \text{ is odd}\}}.$$

An eligible fake costs m, while an ineligible fake has no effect on the designated identity. In the Bayesian incentive comparisons, after fixing all strategic eligibility outcomes and all but one honest outcome, one independent fair bit remains; it makes the final parity fair. Consequently the conditional expected miner revenue is always $m/2$, regardless of any finite strategic padding. This observation proves both coalition properties in this regime, while the elementary payoff $a(v - m)$ for eligibility probability a proves ex post false-name user truthfulness. The same argument gives honest revenue $m/2 \geq hm/128$.

The large-h mechanism. Fix $h \geq 64$ and set $c = h/128$, $T = \lfloor h/4 \rfloor$, and $r = 9/10$. For a submitted length $N \geq h$, let $B_N \sim \mathrm{Bin}(N, 1/2)$ and introduce

$$\delta_N = 36cr^N, \qquad w_N = (1 - \delta_N)m, \qquad P_N = \Pr[T \leq B_N \leq N - T], \qquad A_N = \frac{c}{(1 - \delta_N)P_N}.$$

Below h we put $\delta_N = 1/4$ and set miner revenue to zero; this completes the mechanism on every finite input while preserving the monotonicity of δ_N. On input $(h, D, b_1, \ldots, b_N)$, SybilPotentialMechanism samples $E_i = E(b_i)$ and outputs

$$x_i = \begin{cases} 1, & E_i = 1, \\ \text{Bernoulli}(\delta_N), & E_i = 0, \end{cases} \qquad p_i = w_N E_i, \qquad \mu = w_N A_N \mathbf{1}_{\{T \leq K \leq N - T\}}, \quad K = \sum_i E_i,$$

with the displayed revenue rule used when $N \geq h$. Thus eligible bids pay a discounted threshold price, whereas ineligible bids receive only the small free-confirmation probability δ_N. If all N bids are honest, then $K \sim \mathrm{Bin}(N, 1/2)$, and the choice of A_N gives the exact normalization

$$\mathbb{E}[\mu] = w_N A_N P_N = cm = \frac{hm}{128},$$

independently of N. The central window simultaneously enforces budget feasibility: a binomial tail estimate yields $A_N \leq h/48 \leq T$, so positive revenue implies at least $T \geq A_N$ paying bids.

Padding control and the Sybil potential. Two quantitative ingredients make the normalization robust to false names. First, with $Q_N = 1 - P_N$ and

$$d_N = c\frac{Q_N}{P_N},$$

the tail estimate $Q_N \leq 2e^{-N/8} \leq 2r^N$ gives

$$\delta_n - \delta_N \geq d_N \quad (h \leq n < N), \qquad w_N A_N = cm + md_N \leq cm + w_N. \tag{3}$$

The first inequality says that increasing the submitted length causes enough potential drop to cover the largest possible revenue excess md_N; the second says that one eligible fake's payment already covers that excess.

Second, the binomial padding lemma isolates what happens when all added identities are ineligible. For $B_L \sim \mathrm{Bin}(L, 1/2)$, define

$$F_{L,N} = \Pr[T \leq B_L \leq N - T].$$

Binomial symmetry and unimodality show $F_{L,N} \leq P_N$. Moreover, $F_{N-1,N} = P_N$, and the same equality holds after adding one fixed eligible bit. Hence, in the presence of at least h honest fair bits, arbitrarily many fixed ineligible identities cannot raise expected revenue above cm. When $N - 1 \geq h$ of the bits are honest and fair, the eligibility status of one additional designated identity leaves expected revenue exactly cm. This latter equality supplies the honest coalition baseline needed for exact 1-side-contract-proofness.

Incentive accounting. At total length N, if a report b has eligibility probability $a(b)$, its expected direct utility is

$$u_N(v, b) = \delta_N v + a(b)(1 - \delta_N)(v - m). \tag{4}$$

The coefficient of $a(b)$ changes sign at m, so truthful reporting maximizes this expression. Its truthful optimum is

$$U_N(v) = \begin{cases} \delta_N v, & v \leq m, \\ v - m + \delta_N m, & v \geq m, \end{cases}$$

which is nonincreasing in N. Adding false names therefore cannot improve a user's direct payoff for any fixed outsider bid vector, and eligible fakes only add payments; this proves ex post false-name UIC.

For MIC, the strategic vector is chosen before the honest draws, so the proof may condition on all fake eligibility outcomes while retaining the honest fair bits. If no fake is eligible, the padding lemma bounds expected revenue by cm; if some fake is eligible, its payment w_N, together with (3), covers the entire possible revenue excess even for a miner share $\rho = 1$. The same estimates handle

most of the 1-SCP comparison. Its only additional case is an eligible designated identity padded solely by ineligible fakes. If n is the truthful length and $N > n$ is the padded length, then

$$U_n(v) - (v - w_N) \geq m(\delta_n - \delta_N) \geq md_N.$$

This loss of user utility covers the same md_N revenue excess, and hence also its ρ-fraction. These cases exhaust every finite deviation; averaging the conditional bounds also covers randomized strategies and yields the exact Bayesian incentive inequalities.

Finally, both mechanisms are symmetric in the submitted bids and hence weakly anonymous. Eligibility implies that every positive payment is at most the corresponding bid; odd parity in the small branch and the central-window threshold in the large branch bound revenue by total payments pointwise. The mechanisms use $O(N)$ calls to RationalCoin and exact rational arithmetic; in the large branch, P_N is computed from an $O(N)$-term binomial sum. This gives exact rational outputs and expected running time polynomial in the public-description length, $\log h$, and the full bid-encoding length. Together with the incentive arguments, this proves every guarantee of Theorem 1.1. In particular, for every honest length $\ell \geq h$,

$$\mathbb{E}_{V \sim D^\ell,\, \Pi_{h,D}}[\mu] = \begin{cases} m_D/2, & 1 \leq h < 64, \\ hm_D/128, & h \geq 64, \end{cases}$$

and hence the infimum over $\ell \geq h$ is at least $hm_D/128$.

4 Proof of the Main Theorem

4.1 The parity mechanism for small h

The first branch needs no quantitative tail estimates. Fix $1 \leq h < 64$ and a finite bid vector $b = (b_1, \ldots, b_N)$. Mathematically, sample independent $E_i = E(b_i)$, put $K = \sum_i E_i$, and define

$$x_i = E_i, \qquad p_i = mE_i, \qquad \mu = m\mathbf{1}_{\{K \text{ is odd}\}}. \tag{5}$$

The algorithmic presentation simply implements these already-defined random variables.

ParityMechanism: Implement the parity rule for the small-h regime.
Input: An integer $1 \leq h < 64$, the public description of D, and a finite bid vector $b = (b_1, \ldots, b_N)$.
Output: Confirmation bits $(x_i)_{i=1}^N$, payments $(p_i)_{i=1}^N$, and miner revenue μ.
Guarantee: The outputs obey (5); their incentive, feasibility, revenue, and resource guarantees are stated in Lemma 4.1.

1. Independently sample each E_i according to (2), using RationalCoin(τ) when $b_i = m$.
2. Set $x_i = E_i$ and $p_i = mE_i$ for every i.
3. Compute $K = \sum_{i=1}^N E_i$ and output $\mu = m\mathbf{1}_{\{K \text{ is odd}\}}$.

Lemma 4.1 (Small-regime parity guarantees)**.** *For $1 \leq h < 64$, ParityMechanism is weakly anonymous, pointwise feasible, and expected-polynomial-time. In the model of Definition 2.2, it is ex post false-name user incentive compatible, exact Bayesian miner incentive compatible, and exact Bayesian 1-side-contract-proof for every $\rho \in [0, 1]$ and every finite false-name strategy. If $V \sim D^\ell$ bids truthfully for any $\ell \geq h$, then $\mathbb{E}[\mu] = m/2$.*

Proof. An eligible bid satisfies $b_i \geq m$, is confirmed, and pays m; every other bid is unconfirmed and pays zero. If $\mu = m$, oddness implies $K \geq 1$, and hence

$$\mu = m \leq Km = \sum_i p_i.$$

Thus all pointwise constraints hold. The rule is permutation-equivariant. It makes $O(N)$ eligibility decisions, so Lemma 2.3 gives expected polynomial running time with exact rational outputs.

Fix outsider bids and a user of value v. If a report has eligibility probability a, its expected direct utility is $a(v - m)$. This is maximized by eligibility probability zero when $v < m$, by probability one when $v > m$, and by every probability when $v = m$; truthful reporting attains the maximum. Every eligible fake costs m and every ineligible fake has no effect on the designated identity. A strategy with no designated identity has utility at most zero, while truthful utility is nonnegative. This proves ex post false-name user incentive compatibility.

For either Bayesian coalition property, condition on the strategic bid vector, all of its eligibility outcomes, and all but one honest eligibility bit. At least one honest fair bit remains, so Lemma 2.4 makes the parity of K conditionally fair. Conditional expected miner revenue is therefore always $m/2$, regardless of finite padding. In a miner-only deviation, the miner share remains $\rho m/2$ and fake payments are nonnegative. In a one-user coalition, the designated identity's conditional direct payoff is either zero or $v - m$, and hence is at most the truthful payoff $\max\{0, v - m\}$; fake payments again can only decrease utility. The same conclusions hold after averaging over randomized deviations. Aborting gives zero and cannot improve either nonnegative honest baseline. Finally, when values are drawn independently from D and bid truthfully, there is an honest fair eligibility bit, so the same parity argument yields expected revenue $m/2$. □

4.2 Parameters and binomial padding for large h

We now turn to $h \geq 64$. The parameters below normalize honest revenue while making the free-confirmation probability decrease with the submitted length. Set

$$c = \frac{h}{128}, \qquad T = \left\lfloor \frac{h}{4} \right\rfloor, \qquad r = \frac{9}{10}. \tag{6}$$

For every integer $N \geq 0$, define

$$\delta_N = \begin{cases} 1/4, & N < h, \\ 36cr^N, & N \geq h, \end{cases} \qquad w_N = (1 - \delta_N)m. \tag{7}$$

For $N \geq h$, let $B_N \sim \mathrm{Bin}(N, 1/2)$ and put

$$P_N = \Pr[T \leq B_N \leq N - T], \qquad Q_N = 1 - P_N, \tag{8}$$

$$A_N = \frac{c}{(1 - \delta_N)P_N}, \qquad d_N = c\frac{Q_N}{P_N}. \tag{9}$$

The central interval is nonempty and has positive binomial mass because $T \leq N/4$.

Lemma 4.2 (Sybil-potential parameter bounds)**.** *For $h \geq 64$, the following properties hold.*

1. *The sequence $(\delta_N)_{N \geq 0}$ is nonincreasing and $0 < \delta_N \leq 1/4$.*

2. *For every $N \geq h$,*

$$Q_N \leq 2e^{-N/8} \leq 2r^N, \qquad P_N \geq \frac{1}{2}.$$

3. *For every $N \geq h$,*

$$A_N \leq \frac{h}{48} \leq T, \qquad d_N \leq 4cr^N = \frac{\delta_N}{9} < 1 - \delta_N.$$

4. *If $h \leq n < N$, then*

$$\delta_n - \delta_N \geq d_N. \tag{10}$$

Moreover,

$$w_N A_N = \frac{cm}{P_N} = cm + md_N \leq cm + w_N. \tag{11}$$

Proof. Write $B_N - N/2 = \sum_{i=1}^{N} Y_i$, where the independent variables Y_i are uniform on $\{-1/2, 1/2\}$. For $\lambda \geq 0$,

$$\mathbb{E}[e^{-\lambda Y_i}] = \cosh(\lambda/2) \leq e^{\lambda^2/8}.$$

Indeed, $\log\cosh z \leq z^2/2$ for $z \geq 0$ because its derivative satisfies $\tanh z \leq z$ and both sides vanish at zero. Markov's inequality with $\lambda = 1$ gives

$$\Pr[B_N - N/2 \leq -N/4] \leq e^{-N/4} \prod_{i=1}^{N} \mathbb{E}[e^{-Y_i}] \leq e^{-N/8}.$$

The upper tail has the same bound by symmetry. Since $T \leq N/4$, the event outside $[T, N-T]$ is contained in these two tails, proving $Q_N \leq 2e^{-N/8}$. Also $e^{1/8} > 10/9$, so $e^{-N/8} \leq r^N$. Finally, $r^4 < 2/3$ implies $r^{16} < 1/4$ and hence $r^{64} < 1/256$. Thus $Q_N \leq 2r^N \leq 2r^{64} < 1/2$, which gives $P_N \geq 1/2$.

For every integer $s \geq 64$,

$$\frac{(s+1)r^{s+1}}{sr^s} = r\left(1 + \frac{1}{s}\right) < 1,$$

so sr^s is decreasing. Therefore, for $N \geq h$,

$$\delta_N = 36\frac{h}{128}r^N \leq 36\frac{h}{128}r^h \leq 18r^{64} < \frac{1}{4}.$$

Below h the sequence equals $1/4$, while above h it decreases by a factor r; this proves the first assertion, including the boundary at h.

The bounds $1 - \delta_N \geq 3/4$ and $P_N \geq 1/2$ imply

$$A_N \leq \frac{c}{(3/4)(1/2)} = \frac{h}{48}.$$

Since $T = \lfloor h/4 \rfloor \geq h/8$ for $h \geq 64$, this is at most T. Moreover,

$$d_N = c\frac{Q_N}{P_N} \leq 4cr^N = \frac{\delta_N}{9} < 1 - \delta_N.$$

For $N \geq h + 1$,

$$\delta_{N-1} - \delta_N = 36cr^{N-1}(1 - r) = 4cr^N.$$

If $h \leq n < N$, monotonicity now yields $\delta_n - \delta_N \geq 4cr^N \geq d_N$, proving (10). Finally,

$$w_N A_N = (1 - \delta_N)m\frac{c}{(1 - \delta_N)P_N} = \frac{cm}{P_N} = cm + mc\frac{Q_N}{P_N} = cm + md_N.$$

The inequality $md_N < w_N$ gives (11). □

The parameter estimates control the largest possible revenue value. To control its expectation after costless ineligible identities are appended, we next compare central-window probabilities with different numbers of fair bits.

For fixed $N \geq h$ and $0 \leq L \leq N$, let $B_L \sim \mathrm{Bin}(L, 1/2)$ and define

$$F_{L,N} = \Pr[T \leq B_L \leq N - T]. \tag{12}$$

Lemma 4.3 (Binomial padding comparison)**.** *For every $N \geq h$ and $0 \leq L \leq N$, $F_{L,N} \leq P_N$, with $F_{N-1,N} = P_N$. In addition,*

$$\Pr[T \leq B_{N-1} + 1 \leq N - T] = P_N.$$

Proof. Adding one independent fair bit gives

$$F_{L+1,N} - F_{L,N} = \frac{1}{2}\big(\Pr[B_L = T - 1] - \Pr[B_L = N - T]\big). \tag{13}$$

If $N - T > L$, the second probability vanishes. Otherwise put $a = T - 1$ and $b = N - T$. Then $a \leq b$ and $a + b = N - 1 \geq L$. If $a \leq L/2$, then $L - b \leq a \leq L/2$; symmetry and monotonicity of the binomial coefficients up to their center give $\binom{L}{a} \geq \binom{L}{L-b} = \binom{L}{b}$. If $a > L/2$, monotonicity after the center gives the same inequality directly. This monotonicity follows from

$$\frac{\binom{L}{k+1}}{\binom{L}{k}} = \frac{L - k}{k + 1}.$$

Thus (13) is nonnegative.

At $L = N - 1$, the indices $T - 1$ and $N - T$ sum to L, so the two probabilities in (13) are equal. Hence $F_{L,N}$ is nondecreasing in L and $F_{N-1,N} = F_{N,N} = P_N$. Finally, reflection $k \mapsto N - 1 - k$ maps the interval $[T - 1, N - T - 1]$ onto $[T, N - T]$. The symmetry of B_{N-1} therefore gives

$$\Pr[T \leq B_{N-1} + 1 \leq N - T] = \Pr[T - 1 \leq B_{N-1} \leq N - T - 1] = F_{N-1,N} = P_N.$$

□

4.3 The large-h mechanism and its revenue

The preceding objects now define the large-regime mechanism. On a bid vector $b = (b_1, \ldots, b_N)$, sample independent eligibility bits $E_i = E(b_i)$. Independently of them, sample $Z_i \sim \mathrm{Bernoulli}(\delta_N)$ for every i, ignoring Z_i when $E_i = 1$. Put

$$x_i = \begin{cases} 1, & E_i = 1, \\ Z_i, & E_i = 0, \end{cases} \qquad p_i = w_N E_i, \qquad K = \sum_{i=1}^{N} E_i, \tag{14}$$

and define

$$\mu = \begin{cases} 0, & N < h, \\ w_N A_N \mathbf{1}_{\{T \leq K \leq N - T\}}, & N \geq h. \end{cases} \tag{15}$$

Conditional on the bid vector, all primitive eligibility and free-confirmation coin draws are mutually independent; the displayed output variables are functions of those draws and the reports.

SybilPotentialMechanism: Implement the decreasing free-confirmation and central-window rules.
Input: An integer $h \geq 64$, the public description of D, and a finite bid vector $b = (b_1, \ldots, b_N)$.
Output: Confirmation bits $(x_i)_{i=1}^N$, payments $(p_i)_{i=1}^N$, and miner revenue μ.
Guarantee: The outputs obey (14)–(15); their feasibility, anonymity, and resource guarantees are stated in Lemma 4.4.

1. Compute δ_N and w_N. Independently sample each E_i according to (2), using $\textsc{RationalCoin}(\tau)$ when $b_i = m$.
2. For every i, if $E_i = 1$, set $(x_i, p_i) = (1, w_N)$. If $E_i = 0$, set $x_i \leftarrow \textsc{RationalCoin}(\delta_N)$ and $p_i = 0$.
3. Compute $K = \sum_i E_i$. If $N < h$, output $\mu = 0$; otherwise compute P_N and A_N from (8)–(9) and output $\mu = w_N A_N \mathbf{1}_{\{T \leq K \leq N-T\}}$.

Lemma 4.4 (Large-regime feasibility and exact implementation)**.** *For $h \geq 64$, SybilPotentialMechanism implements the random variables in (14)–(15), is weakly anonymous, satisfies every pointwise feasibility constraint in Definition 2.2, produces exact rational monetary outputs, and runs in expected polynomial time.*

Proof. Steps 1–3 implement the eligibility, free-confirmation, payment, and revenue rules in (14)–(15) directly. An eligible bid is at least m, is confirmed, and pays $w_N = (1 - \delta_N)m \leq m$. An ineligible bid pays zero whether or not its free confirmation succeeds, and an unconfirmed bid always pays zero. If $\mu > 0$, then $N \geq h$, $K \geq T$, and $A_N \leq T$ by Lemma 4.2. Exactly the K eligible bids pay w_N, so

$$\mu = w_N A_N \leq w_N T \leq w_N K = \sum_i p_i.$$

If $\mu = 0$, budget feasibility is immediate. Every rule depends only on N, each bid's relation to m, independent identically distributed coins, and the symmetric statistic K, proving weak anonymity.

For exact implementation, compute

$$P_N = 2^{-N} \sum_{k=T}^{N-T} \binom{N}{k}$$

using the recurrence for consecutive binomial coefficients. This takes $O(N)$ arithmetic operations on integers with $O(N)$ bits. The rationals r^N, δ_N, P_N, A_N, and w_N have bit length polynomial in N, $\log h$, and the public-description length. Hence their exact arithmetic and comparisons take polynomial time. There are $O(N)$ exact rational-coin calls, each expected-polynomial-time by Lemma 2.3. This proves the resource claim and exactness of all outputs. □

Feasibility uses only the central-window threshold. The next lemma uses the padding comparison to identify the normalized expected revenue needed in all three incentive arguments.

Lemma 4.5 (Revenue under fair eligibility and ineligible padding)**.** *Assume $h \geq 64$.*

1. *If a length-N input has $L \geq h$ independent fair eligibility bits and every other eligibility bit fixed to zero, where $N \geq L$, then its expected miner revenue is at most cm.*

2. *If all $N \geq h$ eligibility bits are independent and fair, expected miner revenue is exactly $cm = hm/128$.*

3. *If a length-N input has $N-1 \geq h$ independent fair eligibility bits and one additional eligibility bit fixed to either zero or one, then, conditional on that fixed bit, expected miner revenue is exactly cm.*

Proof. In the first case, the eligibility count is B_L, and Lemma 4.3 gives

$$\mathbb{E}[\mu] = w_N A_N F_{L,N} \leq w_N A_N P_N = cm.$$

In the second case the central event has probability P_N, so equality holds. In the third case, its probabilities for fixed bit zero and fixed bit one are, respectively, $F_{N-1,N}$ and $\Pr[T \leq B_{N-1} + 1 \leq N - T]$. Both equal P_N by Lemma 4.3, and multiplication by $w_N A_N$ gives cm. □

4.4 Incentive properties for large h

We first isolate direct user utility. The decreasing sequence δ_N makes its truthful optimum nonincreasing in the total number of identities, which is precisely the monotonicity needed against false names.

Lemma 4.6 (False-name truthfulness from decreasing confirmation potential)**.** *For $h \geq 64$, SybilPotentialMechanism is ex post false-name user incentive compatible for every finite vector of outsider bids.*

Proof. At total submitted length N, let $a(b) = \Pr[E(b) = 1]$. The expected direct utility of a value-v designated identity reported as b is

$$\begin{aligned} u_N(v, b) &= a(b)(v - w_N) + (1 - a(b))\delta_N v \\ &= \delta_N v + a(b)(1 - \delta_N)(v - m). \end{aligned} \tag{16}$$

Thus truthful reporting maximizes direct utility: it has eligibility probability zero when $v < m$, probability one when $v > m$, and at $v = m$ the coefficient of $a(b)$ vanishes. Its optimum is

$$U_N(v) = \begin{cases} \delta_N v, & 0 \leq v \leq m, \\ v - m + \delta_N m, & v \geq m. \end{cases} \tag{17}$$

By Lemma 4.2, $U_N(v)$ is nonincreasing in N.

Fix s outsider bids. Truthful single-identity participation has length $n = s + 1$ and utility $U_n(v) \geq 0$. A deviation with no designated identity has no intrinsic value and makes only nonnegative payments, so its utility is at most zero. Otherwise its designated identity and finitely many fakes make the new length $N \geq n$. The designated identity contributes at most $U_N(v) \leq U_n(v)$, while all fake payments are nonnegative. This pointwise comparison for each randomized choice of reports remains valid after averaging. □

User truthfulness alone does not control the miner's revenue share. For a miner-only deviation, the padding lemma handles the case of no eligible fake, while one eligible fake already pays enough to cover the maximum revenue excess.

Lemma 4.7 (Miner deviations with finite false names)**.** *For $h \geq 64$,* SybilPotentialMechanism *is exact Bayesian miner incentive compatible for every $\rho \in [0,1]$ and every randomized finite fake-bid strategy chosen independently of at least h honest values.*

Proof. Let $L \geq h$ honest users bid truthfully. Their eligibility bits are independent and fair, so honest miner utility is ρcm by Lemma 4.5. Fix a submitted fake vector and condition on all of its eligibility outcomes; these are independent of the honest fair bits. Let N be the total length and let f be the number of eligible fakes.

If $f = 0$, all fake payments vanish, and the first part of Lemma 4.5 bounds conditional expected revenue by cm. If $f \geq 1$, each eligible fake pays w_N, while (11) bounds revenue pointwise by $cm + w_N$. Conditional coalition utility is therefore at most

$$\rho(cm + w_N) - f w_N = \rho cm - (f - \rho) w_N \leq \rho cm.$$

The inequality survives averaging over fake eligibility coins and a randomized choice of any finite fake vector. Aborting yields zero, also no more than the honest baseline. □

It remains to combine both effects for a coalition of miners and one user. The only new configuration has one eligible designated identity padded solely by ineligible fakes; the potential drop in (10) was chosen to pay for that configuration.

Lemma 4.8 (One-user coalition deviations)**.** *For $h \geq 64$,* SybilPotentialMechanism *is exact Bayesian 1-side-contract-proof for every $\rho \in [0,1]$ and every randomized finite joint deviation by the miners and one user, when at least h honest outsiders have independent values from D.*

Proof. Let $L \geq h$ be the number of honest outsiders, let the colluding user's value be $v \geq 0$, and set $n = L + 1$. Under truthful single-identity participation, the user's expected utility is $U_n(v)$. Conditional on either eligibility status of that identity, expected miner revenue is cm by the third part of Lemma 4.5. Thus truthful coalition utility is

$$U_n(v) + \rho cm. \tag{18}$$

Fix a finite deviating bid vector and condition on every eligibility outcome of its identities. Let N be the total length, and let f count eligible fake identities, excluding the designated identity when one exists.

If $f \geq 1$, then $N \geq n$. Conditional on its eligibility, a designated identity's direct utility is at most $U_N(v)$: an ineligible one gives $\delta_N v$, while an eligible one gives $v - w_N$; comparison with (17) verifies the claim on both sides of m. With no designated identity, direct utility is zero. Thus direct utility is at most $U_N(v) \leq U_n(v)$. Using (11), conditional coalition utility is at most

$$U_n(v) - f w_N + \rho(cm + w_N) \leq U_n(v) + \rho cm.$$

Suppose $f = 0$ and no identity is designated. The coalition has no intrinsic value, all of its identities are ineligible, and its payments vanish. The first part of Lemma 4.5, applied to the L fair honest bits, bounds expected revenue by cm. This includes complete dropout, for which $N = L$.

Suppose $f = 0$ and a designated identity is ineligible. Then $N \geq n$, its direct utility is $\delta_N v \leq U_N(v) \leq U_n(v)$, all coalition payments vanish, and the same padding bound limits expected revenue to cm.

Finally suppose $f = 0$ and the designated identity is eligible. Here $N \geq n$. If $N = n$, there are no fakes; its direct utility $v - w_n$ is at most $U_n(v)$, and conditional expected revenue is exactly cm by Lemma 4.5. If $N > n$, all additional identities are ineligible and cost zero. Equation (17) gives

$$U_n(v) - (v - w_N) \geq m(\delta_n - \delta_N). \tag{19}$$

For $v \geq m$ this is equality; for $v \leq m$, the left side exceeds the right side by $(1 - \delta_n)(m - v)$. Because $n \geq h$ and $N > n$, (10) yields

$$U_n(v) - (v - w_N) \geq md_N.$$

At the same time, (11) bounds expected revenue by $cm + md_N$. Hence deviation utility is at most

$$U_n(v) - md_N + \rho(cm + md_N) = U_n(v) + \rho cm - (1 - \rho)md_N \leq U_n(v) + \rho cm.$$

These cases exhaust every conditioned finite deviation. Averaging over eligibility coins and randomized strategic choices preserves comparison with (18); aborting gives zero and cannot improve that nonnegative baseline. □

Proof of Theorem 1.1. Define $\Pi_{h,D}$ by ParityMechanism for $1 \leq h < 64$ and by SybilPotentialMechanism for $h \geq 64$. Exact balanced eligibility follows from Lemma 2.4. In the first regime, all incentive, pointwise feasibility, anonymity, complexity, and revenue claims follow from Lemma 4.1. In the second regime, feasibility, anonymity, and exact expected-polynomial implementation follow from Lemma 4.4; the three incentive properties follow from Lemmas 4.6 to 4.8. Each incentive proof allows an arbitrary finite strategic bid vector.

For every truthful input length $\ell \geq h$, Lemmas 4.1 and 4.5 gives

$$\mathbb{E}_{V \sim D^\ell, \Pi_{h,D}}[\mu] = \begin{cases} m_D/2, & 1 \leq h < 64, \\ hm_D/128, & h \geq 64. \end{cases}$$

When $h < 64$, $m_D/2 \geq hm_D/128$. Taking the infimum over $\ell \geq h$ proves the asserted lower bound. □

References

[CKWN16] Miles Carlsten, Harry A. Kalodner, S. Matthew Weinberg, and Arvind Narayanan. On the instability of Bitcoin without the block reward. In *Proceedings of the 2016 ACM SIGSAC Conference on Computer and Communications Security*, pages 154–167. Association for Computing Machinery, 2016.

[CS23] Hao Chung and Elaine Shi. Foundations of transaction fee mechanism design. In Nikhil Bansal and Viswanath Nagarajan, editors, *Proceedings of the 2023 Annual ACM-SIAM Symposium on Discrete Algorithms (SODA)*, pages 3856–3899. Society for Industrial and Applied Mathematics, 2023.

[CSLZZ25] Xi Chen, David Simchi-Levi, Zishuo Zhao, and Yuan Zhou. Bayesian mechanism design for blockchain transaction fee allocation. *Operations Research*, 73(4):1944–1964, 2025.

[GY24] Yotam Gafni and Aviv Yaish. Discrete and bayesian transaction fee mechanisms. In Stefanos Leonardos, Elise Alfieri, William J. Knottenbelt, and Panos Pardalos, editors, *Mathematical Research for Blockchain Economy*, Lecture Notes in Operations Research, pages 145–171, Cham, 2024. Springer.

[HLM21] Gur Huberman, Jacob D. Leshno, and Ciamac Moallemi. Monopoly without a monopolist: An economic analysis of the Bitcoin payment system. *The Review of Economic Studies*, 88(6):3011–3040, 2021.

[Rou24] Tim Roughgarden. Transaction fee mechanism design. *Journal of the ACM*, 71(4):30:1–30:25, 2024.

[SCW23] Elaine Shi, Hao Chung, and Ke Wu. What Can Cryptography Do for Decentralized Mechanism Design? In Yael Tauman Kalai, editor, *14th Innovations in Theoretical Computer Science Conference (ITCS 2023)*, volume 251 of *Leibniz International Proceedings in Informatics (LIPIcs)*, pages 97:1–97:22, Dagstuhl, Germany, 2023. Schloss Dagstuhl – Leibniz-Zentrum für Informatik.

[TY23] Wenpin Tang and David D. Yao. Transaction fee mechanism for proof-of-stake protocol, 2023.

[WSC24] Ke Wu, Elaine Shi, and Hao Chung. Maximizing Miner Revenue in Transaction Fee Mechanism Design. In Venkatesan Guruswami, editor, *15th Innovations in Theoretical Computer Science Conference (ITCS 2024)*, volume 287 of *Leibniz International Proceedings in Informatics (LIPIcs)*, pages 98:1–98:23, Dagstuhl, Germany, 2024. Schloss Dagstuhl – Leibniz-Zentrum für Informatik.