

Credible Auctions via MPC Gadgets: Bounding Information Leakage Under Abort

Matheus V. X. Ferreira   

University of Virginia, Charlottesville, VA, USA

Abstract

The design of credible auctions – mechanisms where a revenue-maximizing auctioneer has no incentive to deviate from the protocol – faces a fundamental cryptographic barrier when the auctioneer controls shill bidders. While a natural approach is to use Secure Multi-Party Computation (MPC) to remove the trusted auctioneer, Cleve’s (1986) impossibility of fair coin flipping implies that monolithic MPC protocols grant the auctioneer a “free option”: they can learn the auction’s outcome and unilaterally abort if the revenue is unsatisfactory. Previous attempts to mitigate this abort asymmetry using cryptographic commitments and ex-ante penalties fail for heavy-tailed distributions. We overcome this impossibility by introducing the *MPC Decomposition Principle*. Rather than encrypting the entire mechanism, we use MPC strictly as an information-restriction tool. We isolate the *winner determination problem* into a minimal MPC gadget that computes and reveals the winner’s identity but no payment information. This qualitative restriction mathematically bounds the information leaked upon an abort. By combining this gadget with sequential revelation and finite economic penalties, we design the *Sequential Revelation Auction* (SRA). We prove that bounding the information leakage strictly bounds the value of the free option, showing that a penalty of $\kappa \geq \sum_{i=1}^n \text{REV}(F_i)$ is sufficient for credibility, and tight in the worst case: for the equal-revenue distribution – where this threshold coincides exactly with the optimal revenue – every smaller penalty admits a profitable deviation. Leveraging low-round-complexity MPC, the SRA resolves an open question of [1] and [11] by providing the first constant-round, incentive-compatible, revenue-optimal credible auction for all product distributions.

2012 ACM Subject Classification Theory of computation → Algorithmic mechanism design; Security and privacy → Cryptography

Keywords and phrases Multi-Party Computation, Credible Auctions, Revenue-optimal Auctions, Cryptographic Auctions

Digital Object Identifier 10.4230/LIPIcs.AFT.2026.18

Related Version *Full Version*: <https://arxiv.org/abs/2609.27402>

Acknowledgements I thank Justin Taylor and Vamsi Policharla for helpful discussions on constant-round MPC with identifiable abort via homomorphic encryption.

1 Introduction

The study of mechanism design traditionally assumes a trusted intermediary who faithfully implements the mechanism’s rules. In practice – ranging from large-scale digital advertising to decentralized finance [4] – the auctioneer is a rational agent who may deviate if doing so is both undetectable and increases their revenue. The literature on *credible auctions*, initiated by [1], formalizes this problem. In the full information setting (where the auctioneer can observe all the actions from bidders), they proved a fundamental impossibility: the ascending price auction is the unique revenue-optimal, incentive-compatible auction without undetectable profitable deviations, a desideratum we call credibility. Unfortunately, this auction requires an unbounded amount of communication in the worst case.

To circumvent this impossibility, [11] proposed transforming the auction into a game of incomplete information using cryptography. Their *Deferred Revelation Auction (DRA)* requires bidders to cryptographically commit to their valuations, revealing them only after



© Matheus V. X. Ferreira;

licensed under Creative Commons License CC-BY 4.0

8th Conference on Advances in Financial Technologies (AFT 2026).

Editors: Aggelos Kiayias and Maria Kyrpoulou; Article No. 18; pp. 18:1–18:22

Leibniz International Proceedings in Informatics



LIPICs Schloss Dagstuhl – Leibniz-Zentrum für Informatik, Dagstuhl Publishing, Germany

the auctioneer commits to the execution path. Moreover, participants deposit collaterals to incentivize opening of commitments. However, this approach hits a fundamental barrier when applied to bidder valuations drawn from heavy-tailed distributions (e.g., the equal-revenue distribution), where [11] proved that no finite ex-ante penalty can deter a malicious auctioneer. In the DRA, this failure occurs because the auctioneer’s shill bidder can wait for honest bidders to reveal their commitments, observe their valuations, and then refuse to open its own commitment. This forced abort triggers a restart where the auctioneer now knows the bids. For heavy-tailed distributions, the expected revenue gained from exploiting this acquired knowledge is infinite, meaning no finite collateral can deter the abort.

A natural alternative is to abandon commit-reveal schemes and use Secure Multi-Party Computation (MPC) to encrypt the entire auction. However, this monolithic MPC approach hits a fundamental cryptographic barrier first identified by [8] in the context of fair coin flipping: the *abort problem*. Cleve established that in protocols without an honest majority (which includes an auctioneer commanding unbounded shill bids), an adversary can always learn their output and unilaterally abort the protocol before honest parties receive theirs.

In economic terms, Cleve’s abort asymmetry grants the auctioneer a “free option”. If the auctioneer encrypts the entire optimal mechanism (allocation and payments) within an MPC protocol, they can run the protocol, observe the provisional revenue, and if it is unsatisfactory, abort the execution and restart. Since the abort reveals that all active bidders have valuations below a certain threshold, the restart is not an independent trial but a biased one. This allows the auctioneer to safely explore the bidders’ valuations using shills, effectively extracting full surplus and violating the credibility of the mechanism.

In this work, we introduce the *Sequential Revelation Auction (SRA)*, which resolves this impossibility and achieves credibility for *any* product distribution over bidder valuations. We overcome the free option problem by introducing the *MPC Decomposition Principle*. Rather than following the traditional paradigm of “encrypting everything,” we use MPC strictly as an *information-restriction tool*. Instead of computing the entire mechanism inside the MPC, we isolate the smallest possible piece of information necessary: the identity of the winner. We introduce a *Winner Determination Gadget (WDG)*, a minimal MPC functionality that takes committed bids and solves the Winner Determination Problem (WDP) by outputting *only* the identity of the winner, revealing *no payments* or *bids*.

This decomposition shifts the information asymmetry. If the auctioneer uses a shill to probe the price and subsequently aborts the gadget, they learn only *who* was winning, but not *how much* bidders were willing to pay. We formally quantify the value of this leaked information. Our main technical contribution, the Information Leakage Lemma (lemma 32) bounds the extra revenue the auctioneer can extract from the knowledge acquired during an abort. Therefore, by setting an ex-ante abort penalty κ to at least this threshold, we ensure that the cost of aborting always exceeds the value of the free option.

1.1 Technical Overview

We assume the existence of an MPC protocol with identifiable abort that implements any ideal functionality f . This protocol allows a set of parties to compute any ideal functionality $f(v_1, \dots, v_n) = (y_1, \dots, y_n)$ where v_i is the input that player i contributes to evaluating f and y_i is the output only he observes. A key feature of this model is that any player i can abort the execution at any time learning y_i and preventing others from learning y_{-i} ; however, if they do so, all other parties learn the identity of the deviating player.

In the SRA, we instantiate f as the WDP: given a set of bids $\mathbf{v}$, output some $i^* \in \arg \max_i v_i$. Then, we ask the winner to reveal v_{i^*} . It is not hard to see that the only undetectable deviation available to the auctioneer is to modify $\mathbf{v}$ by adding a set shill bids

$r_1 < r_2 < \dots < r_k$ into the input and if chosen as the winner, to refuse to reveal v_{i^*} in which case the protocol recomputes the WDP without i^* . The core logic of our credibility proof relies on an “unraveling” argument to argue that this abort does not leak sufficient information to violate credibility. By bounding the information leaked during such an abort, we show that setting the penalty $\kappa \geq \sum_{i=1}^n \text{REV}(F_i)$ ensures that aborting is never profitable where $(F_1, \dots, F_n)$ is the value distribution profile of real bidders and $\text{REV}(F_i)$ is the optimal revenue (lemma 7) an auctioneer can obtain in a single-bidder auction with a bidder with valuation drawn from F_i .

1.2 Our Contribution

We provide the first constant-round, incentive-compatible, revenue-optimal auction that is credible for all product distributions over bidder valuations, resolving the open problem left by [11]. Our contributions are threefold:

1. **The MPC Decomposition Principle:** We introduce a new design methodology for credible auctions based on *decomposing* the auction computation. Rather than implementing the entire mechanism (allocation and payments) inside MPC – which, as we show in Section 3, inevitably leaks enough information to violate credibility – we isolate the minimal subroutine needed: a Winner Determination Gadget that computes *only* the Winner Determination Problem. Payments are then determined *outside* the MPC through sequential revelation. This decomposition principle is the key conceptual contribution: it shows that for credibility, minimizing the computed function is more important than encrypting everything.
2. **General Credibility:** We prove that for any product distribution $\mathbf{F}$, the SRA is revenue-optimal, incentive-compatible, and credible provided the penalty κ satisfies $\kappa \geq \sum_{i=1}^n \text{REV}(F_i)$ (Theorem 39). This overcomes the fundamental barrier identified by [11]: for heavy-tailed distributions (e.g., the equal-revenue distribution) where they exhibit a single non-MHR bidder for which *no* finite penalty makes the Deferred Revelation Auction credible. Our protocol circumvents this impossibility by using MPC to change the information structure of the game.
3. **Tightness and Round Complexity:** We show that the penalty bound $\kappa \geq \sum_{i=1}^n \text{REV}(F_i)$ is tight: for the equal-revenue distribution, any penalty $\kappa < \sum_{i=1}^n \text{REV}(F_i)$ admits profitable deviations (Theorem 40). Additionally, the SRA runs in constant rounds on the equilibrium path (assuming standard constant-round MPC constructs). This complements the ascending protocol of [10], which achieves general credibility for n i.i.d. bidders in constant rounds *in expectation* but has unbounded worst-case round complexity. Our guarantee differs along two axes: the round bound on the equilibrium path is deterministic rather than in expectation, and our credibility guarantee holds for arbitrary *independent, non-identical* distributions $F_1, \dots, F_n$ rather than the i.i.d. case.

Practical Collateral Requirements

The threshold $\kappa \geq \sum_{i=1}^n \text{REV}(F_i)$ is practically reasonable, not merely finite. On the canonical stress test – the equal-revenue distribution $F(v) = 1 - 1/v$ for $v \geq 1$, where [11] showed the Deferred Revelation Auction needs *infinite* collateral even against a single bidder – the SRA is credible at $\kappa = 1$ per bidder, and since $v \geq 1$ on the support, no bidder ever posts more than their own valuation. Feasibility gives $\text{OPT}(\mathbf{F}) = \mathbb{E} [\max_i \bar{\phi}_i(v_i)^+] \leq \sum_{i=1}^n \mathbb{E} [\bar{\phi}_i(v_i)^+] = \sum_{i=1}^n \text{REV}(F_i)$, so the requirement is $\text{OPT}(\mathbf{F}) \leq \kappa \leq n \cdot \text{OPT}(\mathbf{F})$: between one and n times the revenue the auction generates. It is also a *liquidity* requirement rather than an expected cost, since every bidder who follows the protocol recovers κ in full.

1.3 Related Work

Secure Multi-Party Computation

Secure Multi-Party Computation (MPC) was introduced by [19] with the “Millionaires’ Problem” and [3] with “Coin Flipping by Telephone.” The first large-scale commercial application of MPC was the Danish Sugar Beet Auction [4], where farmers traded production contracts. In this system, robustness against bidder dropout was handled by re-running the computation or assuming the servers were honest. This approach to *safety* was acceptable because the threat model did not involve a malicious auctioneer. In contrast, our work focuses on *credibility*, where the auctioneer is the primary adversary. We show that the restart mechanism, essential for safety, becomes a vulnerability for credibility. For our auction to have small round complexity, we rely on works showing that any functionality can be realized in a constant number of rounds. Two rounds suffice for security *with abort* [13, 16], relying on powerful cryptographic primitives like *Multi-Key Fully-Homomorphic Encryption* (MK-FHE) and *indistinguishability obfuscation* (iO). Our protocol requires the stronger guarantee of *identifiable abort*, since the penalty must be charged to a named party, and identifiability provably costs additional rounds: [7] show that four rounds are necessary and sufficient in the plain model. Constant-round protocols with identifiable abort are known [15, 2, 9], which is all our result requires. While these primitives are currently computationally expensive, they establish the theoretical feasibility of avoiding the high round complexity of previous credible mechanisms [10].

Credible Auctions

The credibility framework was established by [1], who showed that unconditional non-trivial credible auctions are impossible. [11] introduced cryptographic assumptions to the problem, constructing two-round revenue optimal credible auctions for MHR distributions. However, their negative result for heavy-tailed distributions left open the question of general feasibility. Our work resolves this open problem. [10] also achieved general credibility, for n i.i.d. bidders with a finite monopoly price, in constant rounds in expectation but with unbounded worst-case round complexity. In contrast, our SRA runs in constant rounds (assuming standard constant-round MPC constructs) on the equilibrium path, and handles independent but non-identical distributions. Like theirs, our worst-case round complexity is unbounded: the number of rounds grows linearly with the number of bidders that abort, but aborting is a dominated strategy for all bidders (real or shills).

[6] extended DRA to settings where bidders can broadcast messages, which is relevant to our MPC-based approach since MPC with identifiable abort requires broadcast channels. [12] extended the DRA to multi-item settings with matroid constraints but also require similar distributional assumptions to [11] and [10].

There is a rich literature on implementing auction mechanisms using MPC [18, 5]. Most of this work focuses on *privacy* (keeping bids secret) or *fairness* (preventing early bid peeking). Our focus on *credibility* is distinct: we view the auctioneer as the adversary and specifically target the strategic implications of the abort capability. Our result highlights that for credibility, “encrypting everything” is not the solution; rather, careful information management via minimizing the computed function is key.

1.4 Paper Organization

Section 2 defines the credibility framework. Section 3 illustrates the failure of black-box MPC. Section 4 details the SRA protocol. Section 5 provides a warm up by showing that the SRA is credible for the equal-revenue distribution. This addresses the open problem of [11]. Section 6 bounds the information an abort leaks, and Section 7 uses that bound to prove credibility for every product distribution, which addresses the open problem of [11]. Section 8 provides the matching lower bound. Proofs omitted for space, together with a worked equal-revenue warmup, appear in the full version.

2 Preliminaries

We consider a single-item auction setting with a set of honest, risk-neutral bidders $N = \{1, \dots, n\}$. Additionally, there exists a countable set of potential shill bidders $S = \{n+1, n+2, \dots\}$. Honest bidders $i \in N$ have private valuations $v_i \in V_i \subseteq \mathbb{R}_{\geq 0}$ drawn independently from cumulative distribution functions F_i with probability density functions f_i . Shill bidders $s \in S$ have a fixed valuation $v_s = 0$. Crucially, shill bidders share the auctioneer's utility function (revenue maximization) and do not behave as self-interested agents maximizing their own allocation utility. We denote the valuation profile of honest agents by $\mathbf{v} = (v_1, \dots, v_n) \in \mathbf{V} = \times_{i \in N} V_i$. For value profile $\mathbf{v} = (v_i)_{i \in N \cup S}$ and subset $A \subseteq N \cup S$, we write $\mathbf{v}_A = (v_i)_{i \in A}$ for the restriction of $\mathbf{v}$ to bidders in A . For any function f , we write $f^+ := \max\{f, 0\}$ for the positive part.

We begin by formally defining the core concepts of optimal auctions.

► **Definition 1 (Mechanism).** A direct revelation mechanism $M = (x, p)$ consists of an allocation rule $x : \mathbf{V} \rightarrow [0, 1]^n$, where $x_i(\mathbf{v})$ is the probability that bidder i receives the item and feasibility requires $\sum_{i \in N} x_i(\mathbf{v}) \leq 1$ for all $\mathbf{v} \in \mathbf{V}$, together with a payment rule $p : \mathbf{V} \rightarrow \mathbb{R}^n$, where $p_i(\mathbf{v})$ is the expected payment from bidder i .

We assume agents have quasi-linear utility functions.

► **Definition 2 (Utility and DSIC).** The utility of bidder i with valuation v_i who reports v'_i while the others report $\mathbf{v}_{-i}$ is $u_i(v'_i, \mathbf{v}_{-i}; v_i) = v_i \cdot x_i(v'_i, \mathbf{v}_{-i}) - p_i(v'_i, \mathbf{v}_{-i})$. A mechanism is Dominant Strategy Incentive Compatible (DSIC) if truthful reporting maximizes utility for every honest player $i \in N$, valuation v_i , and profile $\mathbf{v}_{-i}$: $u_i(v_i, \mathbf{v}_{-i}; v_i) \geq u_i(v'_i, \mathbf{v}_{-i}; v_i)$ for all $v'_i \in V_i$.

► **Theorem 3 ([17]).** A direct revelation mechanism (x, p) is DSIC if and only if:

1. **Allocation rule is monotone.** For every bidder i and declarations $\mathbf{v}_{-i}$, $x_i(v_i, \mathbf{v}_{-i})$ is non-decreasing in v_i .
2. **Payment identity.** The payment rule is determined by the allocation rule: for every bidder i , v_i , and $\mathbf{v}_{-i}$, $p_i(v_i, \mathbf{v}_{-i}) = v_i \cdot x_i(v_i, \mathbf{v}_{-i}) - \int_0^{v_i} x_i(z, \mathbf{v}_{-i}) dz$.

► **Definition 4 (Revenue).** The ex-post revenue of a mechanism $M = (x, p)$ for valuation profile $\mathbf{v}$ is the sum of payments, $R^M(\mathbf{v}) = \sum_{i \in N} p_i(\mathbf{v})$; its expected revenue is $\text{REV}^M(\mathbf{F}) = \mathbb{E}_{\mathbf{v} \sim \mathbf{F}} [R^M(\mathbf{v})]$, and we write $R(\cdot)$ for short when M is clear from context. We write $\text{OPT}(\mathbf{F}) = \max_{M \text{ DSIC}} \text{REV}^M(\mathbf{F})$ for the optimal expected revenue over DSIC mechanisms, and, for a single distribution F_i , $\text{REV}(F_i) = \sup_{p \geq 0} p \cdot (1 - F_i(p))$ for the optimal monopoly revenue obtainable from bidder i alone by posting a price. Throughout, $\text{REV}(\cdot)$ applied to a single distribution always denotes this monopoly revenue, while revenue of a mechanism always carries the mechanism in the superscript.

Throughout we assume each F_i has a vanishing revenue tail, $\lim_{v \rightarrow \infty} v(1 - F_i(v)) = 0$; this is the hypothesis of Lemma 7 below, and it implies $\text{REV}(F_i) < \infty$, so the penalty $\sum_i \text{REV}(F_i)$ is finite.

► **Definition 5** (Regular Distribution). *A differentiable cumulative distribution function F with density f is said to be regular if the virtual valuation function $\phi(v) = v - \frac{1-F(v)}{f(v)}$ is monotone non-decreasing.*

► **Definition 6** (Ironed Virtual Valuation). *For any distribution F , the ironed virtual valuation function $\bar{\phi}(v)$ is constructed by applying Myerson's ironing procedure to $\phi(v)$ to ensure monotonicity. Specifically, let $H(q) = \int_0^q \phi(F^{-1}(t))dt$ for $q \in [0, 1]$. We define G as the convex hull of H , i.e., the largest convex function such that $G(q) \leq H(q)$ for all q . Then, $\bar{\phi}(v) = G'(F(v))$. If F is regular, $\bar{\phi}(v) = \phi(v)$. We define the inverse $\bar{\phi}^{-1}(y) := \inf\{v : \bar{\phi}(v) \geq y\}$, which is well-defined even when $\bar{\phi}$ has flat regions due to ironing.*

► **Lemma 7** (Virtual surplus identity; [17]). *For any DSIC mechanism M allocating a single item to agents with independent private values drawn from distributions F_i satisfying $\lim_{v \rightarrow \infty} v(1 - F_i(v)) = 0$, the expected revenue equals the expected virtual surplus:*

$$\text{REV}^M(\mathbf{F}) = \mathbb{E}_{\mathbf{v} \sim \mathbf{F}} \left[\sum_i \bar{\phi}_i(v_i) x_i(\mathbf{v}) \right]$$

where $\bar{\phi}_i(v_i)$ is the ironed virtual valuation function of bidder i .

Myerson's result identifies the revenue-optimal auction as the mechanism that allocates the item to the bidder with the highest ironed virtual value, provided it is non-negative. Ironing makes ties a positive-measure event rather than a null one – $\bar{\phi}_i$ is flat on each ironed interval, so every profile landing inside one produces a tie – and the Winner Determination Gadget of Section 4 publishes which of the tied parties is provisionally winning. We therefore fix the comparison and the tie-break together, as a single ranking, rather than treating the latter as an afterthought.

► **Definition 8** (Ranking). *For participants $i \neq j$ submitting bids b_i, b_j , we write $\bar{\phi}_i(b_i) < \bar{\phi}_j(b_j)$ to mean that the mechanism ranks j above i : either $\bar{\phi}_i(b_i)$ is numerically smaller than $\bar{\phi}_j(b_j)$; or the two are numerically equal and $b_i < b_j$; or both of these are equalities and $j < i$. We write $\leq$ for the reflexive version. A higher ironed virtual value therefore wins, ties go to the higher raw bid, and any remaining tie goes to the smaller index; the last clause separates any two distinct participants, so this is a strict total order and every finite set of participants has a unique maximum.*

Two conventions keep this unambiguous. First, the overload applies only between two participants' virtual values; a comparison against a constant, such as $\bar{\phi}_i(v_i) \geq 0$, is always the ordinary numeric one. Second, the relation is determined by the pairs (i, b_i) and (j, b_j) rather than by the two numbers alone, so it is antisymmetric even where the numbers coincide: $\bar{\phi}_i(b_i) \leq \bar{\phi}_j(b_j)$ and $\bar{\phi}_j(b_j) \leq \bar{\phi}_i(b_i)$ hold together only when $i = j$. The bids and identities are always clear from context.

The ranking is monotone in a participant's own bid: for $z < z'$ we have $\bar{\phi}_i(z) \leq \bar{\phi}_i(z')$ numerically because $\bar{\phi}_i$ is non-decreasing, and either that inequality is strict, in which case the first clause ranks z' above z , or the two agree numerically and the second clause does.

► **Definition 9** (Myerson mechanism). *The mechanism allocates to the highest-ranked participant whose ironed virtual value is non-negative, $x_i(\mathbf{b}) = \mathbb{1}_{\bar{\phi}_i(b_i) \geq 0} \cdot \mathbb{1}_{\bar{\phi}_j(b_j) < \bar{\phi}_i(b_i) \forall j \neq i}$, and charges the winner the threshold payment induced by that rule, $p_i(\mathbf{b}) = \inf\{z \geq 0 : x_i(z, \mathbf{b}_{-i}) = 1\}$, losers paying nothing.*

Winning under Definition 9 requires i to outrank the fixed best competitor and to satisfy $\bar{\phi}_i(b_i) \geq 0$; both conditions are upward closed in i 's own bid by the monotonicity of Definition 8, so the allocation rule is monotone and, with its threshold payments, the mechanism is DSIC by Theorem 3.

We now turn from the economic setting to the protocol design problem. The key tension is between *safety* (the protocol should handle dropouts gracefully) and *credibility* (the auctioneer should not be able to exploit the protocol). Our definitions make this tension precise.

2.1 Credibility

In this section, we formalize the notion of credibility we consider. Intuitively, a mechanism is credible if the auctioneer cannot profit by deviating from the protocol in a way that is indistinguishable from honest behavior.

► **Definition 10** (Auction Protocol). *An auction protocol is a pair (M, Π) , where $M = (x, p)$ is a mechanism and Π is a cryptographic protocol that purportedly realizes M .*

► **Definition 11** (Safe Deviation [1]). *A deviation by the auctioneer in protocol Π is safe if, for every honest bidder i , there exists a strategy profile of the other bidders $N \cup S \setminus \{i\}$ such that i 's view under the deviation is computationally indistinguishable from i 's view under the honest execution of Π with that strategy profile.*

To illustrate, consider a second-price auction. The deviation where the auctioneer shill bids is a *safe* deviation since from any real bidder's perspective, the shill is indistinguishable from a real bidder. In contrast, an *unsafe* deviation would be one where the auctioneer refuses to allocate the item to the winner. This aims to capture an auctioneer that is willing to deviate as long as deviations cannot be detected by honest bidders.

► **Definition 12** (Credibility). *An auction protocol (M, Π) is credible if there exists no safe deviation for the auctioneer within Π that yields a strictly higher expected revenue than the honest execution of M without shills.*

We can now state our primary design goal:

► **Definition 13** (Optimal Credible Auction Design Problem). *The goal is to find an auction protocol (M, Π) maximizing $\text{REV}(M)$ subject to M being DSIC for the honest bidders N and (M, Π) being credible.*

One further requirement is what makes that problem hard, and it is the reason the strawman of Section 3 fails. A protocol must remain usable when a participant simply stops responding.

► **Definition 14** (Safety against Honest Abort). *A protocol Π satisfies Safety against Honest Abort if for any set of bidders P , if a non-auctioneer player $i \in P$ aborts, the protocol can continue and the expected revenue – that is, the expected sum of payments from the remaining bidders – equals the expected revenue of the mechanism executed on the set $P \setminus \{i\}$.*

To satisfy Definition 14, a protocol must ensure that a single bidder's refusal to participate does not destroy the auction's value, which in practice requires it to restart or to switch to a mode that ignores the aborting party's input. Paradoxically, this creates the following vulnerability: it hands the auctioneer a "free option" to abort (perhaps through a shill) whenever the outcome is unfavorable, knowing the protocol *must* offer a recourse that preserves revenue – so the auctioneer retries the auction having learned the outcome of the first attempt. Section 3 turns this observation into a concrete example.

2.2 Cryptography

This section introduces basic cryptographic definitions used in the paper.

► **Definition 15** (Cryptographic Commitment). *A commitment scheme lets a committer publish $c = \text{Commit}(v, r)$, a commitment to a message v under randomness r . We require it to be computationally hiding – for distinct v, v' the distributions $\{\text{Commit}(v, r)\}_r$ and $\{\text{Commit}(v', r)\}_r$ are computationally indistinguishable – perfectly binding – no two pairs $(v, r) \neq (v', r')$ with $v \neq v'$ satisfy $\text{Commit}(v, r) = \text{Commit}(v', r')$ – and non-malleable, so that from c no adversary can produce a commitment whose value bears a non-trivial relation to v . See [14] for the formal definitions.*

A multi-party computation (MPC) with identifiable abort is a protocol that realizes an ideal functionality $\mathcal{F}$ that takes inputs $\mathbf{z} = (z_1, \dots, z_n) \in (\mathcal{Z} \cup \{\perp\})^n$ from n parties and computes a function $f(\mathbf{z}) = (y_1, \dots, y_n)$, where $y_i \in \mathcal{Y} \cup \{\perp\}$. An input of $\perp$ means that the party contributed with no input, and an output of $\perp$ means that the party gets no output.

While an ideal MPC would guarantee output delivery to all parties, this is known to be impossible in the presence of a dishonest majority [8]. To build intuition, consider the simple “Coin Flipping” problem where two parties want to generate a uniform random bit $r = r_1 \oplus r_2$. In the ideal world, a trusted party receives r_1, r_2 and sends r to both. In any real protocol, one party must send their value (or open their commitment) last. If this last party sees the result and doesn’t like it (e.g., they wanted Heads but the result is Tails), they can simply abort. This allows the last party to bias the output (forcing an abort is equivalent to rejecting the coin flip). This impossibility result implies that in any auction protocol with a potentially malicious auctioneer (dishonest majority), we must allow for the possibility of *abort*.

This is particularly relevant in the context of auctions, where the auctioneer is allowed to control any number of shill bidders. We therefore rely on the standard relaxation of security with *identifiable abort*, which guarantees that if the adversary prevents honest parties from learning the output, their identity is revealed.

► **Definition 16** (Ideal Functionality). *Let $\mathcal{F} : \mathcal{Z}^n \rightarrow \mathcal{Y}^n$ be an n -party function. Ideally, a trusted party computes $(y_1, \dots, y_n) = \mathcal{F}(z_1, \dots, z_n)$ and sends y_i to party P_i .*

We use the standard simulation-based notion of security with identifiable abort, in which the adversary may abort at any point during the execution; we state the ideal experiment explicitly, since the timing of the abort is what drives our incentive analysis.

► **Definition 17** (Ideal model with identifiable abort). *Fix an n -party ideal functionality $\mathcal{F}$ and a set $I \subset [n]$ of corrupted parties. In the ideal execution, a simulator $\mathcal{S}$ controlling the parties in I interacts with the trusted party computing $\mathcal{F}$ as follows.*

1. *Each honest party $i \notin I$ sends its input z_i to the trusted party; $\mathcal{S}$ sends inputs $\{z_i\}_{i \in I}$ of its choice (possibly $\perp$).*
2. *The trusted party computes $(y_1, \dots, y_n) = f(\mathbf{z})$ and sends $\{y_i\}_{i \in I}$ to $\mathcal{S}$.*
3. *$\mathcal{S}$ replies with either *continue*, in which case each honest party i receives y_i ; or *(abort, i^*)* for some $i^* \in I$ of its choice, in which case every party outputs $(\perp, i^*)$.*

► **Definition 18** (Realization with Identifiable Abort). *A protocol Π securely realizes $\mathcal{F}$ with identifiable abort if for every probabilistic polynomial-time adversary $\mathcal{A}$ statically corrupting a set $I \subset [n]$ in the real execution of Π , there exists a probabilistic polynomial-time simulator $\mathcal{S}$ in the ideal model of Definition 17 such that the joint distribution of the honest parties’ outputs and the adversary’s view is computationally indistinguishable in the two executions. See [14] and [15] for the formal treatment.*

Note that Definition 17 already encodes the worst case for our analysis: the adversary learns its own outputs $\{y_i\}_{i \in I}$ *before* deciding whether honest parties receive theirs, and an abort names a corrupted party. Every deviation we analyze in Section 7 is a choice of continue versus (abort, i^*) at Step 3 of Definition 17, so no generality is lost by restricting attention to that decision.

► **Theorem 19** ([15, 2, 9]). *Under standard cryptographic assumptions and the availability of a broadcast channel, for any ideal functionality $\mathcal{F}$, there exists a protocol Π that securely realizes $\mathcal{F}$ with identifiable abort in the presence of a static adversary controlling any number of corrupted parties $t < n$ in a constant number of rounds of communication.*

3 Strawman

In this section, we analyze the credibility of a straightforward attempt to realize a revenue-optimal auction using generic Multi-Party Computation (MPC). We define the ideal functionality $\mathcal{F}_{OPT}$ for the single-item Myerson auction – which receives bids, computes the optimal allocation and payments, and distributes the results – and demonstrate that any protocol realizing $\mathcal{F}_{OPT}$ while satisfying *Safety against Honest Abort* (Definition 14) is vulnerable to a welfare-extraction attack by a malicious auctioneer.

■ **Algorithm 1** Optimal Auction Functionality $\mathcal{F}_{OPT}$.

Input: Bid profile $\mathbf{b} = (b_1, \dots, b_n)$ for bidders P .

Computation:

1. Determine allocation x :

$$x_i(\mathbf{b}) = \begin{cases} 1 & \text{if } i = \operatorname{argmax}_{j \in [n]} \{\bar{\phi}_j(b_j) \mid \bar{\phi}_j(b_j) \geq 0\} \\ 0 & \text{otherwise} \end{cases}$$

2. Compute payment p :

$$p_i(\mathbf{b}) = \begin{cases} \inf\{b_i : \bar{\phi}_i(b_i) \geq \max_{j \neq i} \{0, \bar{\phi}_j(b_j)\}\} & \text{if } x_i = 1 \\ 0 & \text{otherwise} \end{cases}$$

return (x, p) .

We now consider a protocol Π that realizes $\mathcal{F}_{OPT}$. To satisfy *Safety against Honest Abort* (Definition 14) – ensuring that a single bidder’s dropout allows the auctioneer to recover the revenue from the remaining bidders – practical implementations (like [4]) must be able to exclude the aborting party and complete the auction (e.g., via restart).

■ **Algorithm 2** Protocol that realizes $\mathcal{F}_{OPT}$ with Safety Restart.

Input: Bid profile $\mathbf{b} = (b_1, \dots, b_n)$ for bidders P .

Input: MPC protocol Π that realizes $\mathcal{F}_{OPT}$ with identifiable abort

Round 1:

Parties invoke Π on inputs $\mathbf{b}$.

if Π outputs $\perp$ due to abort by $k \in P$ **then**

 Identify aborter k .

 Update $P \leftarrow P \setminus \{k\}$.

Restart round 1 with remaining players P .

end

return (x, p) of Π to all participants.

► **Theorem 20.** *Let Π be any protocol that realizes $\mathcal{F}_{OPT}$. If Π satisfies Safety against Honest Abort, then there exists a safe deviation for the auctioneer that can extract the full social welfare.*

Proof. The core idea is that the auctioneer uses one shill to “probe” the payment – which in the Myerson mechanism reveals the bidder’s valuation – and then aborts to exploit this information in a restart.

Consider a setting with a single bidder with valuation v drawn from a distribution F with support $[\underline{v}, \bar{v}]$ (or $[\underline{v}, \infty)$). Suppose the auctioneer controls two shills S_1 and S_2 . In the first execution, S_1 bids $\underline{v}$ and S_2 bids a value B exceeding the upper bound of the support of F . Since S_2 has the highest bid, S_2 wins the auction. Upon learning the output, the auctioneer observes the payment charged to S_2 , which equals v (the threshold at which S_2 ’s virtual value exceeds the honest bidder’s). The auctioneer then instructs S_2 to abort. By safety, the protocol restarts with S_2 excluded but S_1 still active. Since Algorithm 2 restarts from scratch, S_1 may submit a fresh bid. The auctioneer now sets S_1 ’s bid to $v - \epsilon$ for arbitrarily small $\epsilon > 0$, extracting $v - \epsilon$ from the honest bidder. ◀

This impossibility result reveals the fundamental weakness of monolithic MPC for auction design: the protocol reveals the *entire outcome* (winner and price) to the adversary before honest parties can act. In particular, the adversary learns the highest honest valuation through the payment identity, and the safety restart mechanism provides a costless option to exploit this information.

Why committing bids does not help

One might attempt to fix this by requiring bids to be cryptographically committed, so that the auctioneer cannot modify shill bids after learning the outcome. However, this defense is easily circumvented: the auctioneer commits to k shill bidders with bids $1, 2, \dots, k$ and aborts all shills bidding above the highest honest bidder, effectively probing the honest valuation at no cost.

In the next section, we present the Sequential Revelation Auction, which avoids this vulnerability by decomposing the computation into minimal gadgets that reveal only the winner identity – never the price – and by imposing penalties that make aborting unprofitable.

4 The Protocol

The strawman analysis (Section 3) revealed that treating the MPC as a black box that realizes an auction functionality is fundamentally flawed for credibility: the adversary learns the full outcome upon abort, providing a “free option” to retry with extra information. Our protocol takes a fundamentally different approach. Rather than compute the entire auction mechanism – allocation *and* payments – inside MPC, we decompose the computation into a minimal *gadget* that reveals only the winner identity, and we compute payments *outside* the MPC via simultaneous revelation of committed bids. This design ensures that an abort reveals no payment information, and economic penalties make aborting unprofitable.

The key insight behind our design is that the auctioneer’s “free option” in the strawman protocol stems from learning payment information upon abort. If the MPC computed only the winner identity – but not the payment – an abort would reveal who is winning, but not how much they would pay. Since payments, not allocations, directly determine revenue, this limited information alone is not enough for the auctioneer to construct a profitable deviation, provided the restart is costly. We formalize this idea via the *Winner Determination Gadget*.

► **Definition 21** (Winner determination gadget). Let $\mathbf{c} = (c_i)_{i \in A}$ be a vector of commitments, broadcast by the bidders in the commit phase, before the gadget is invoked (Step 3 of Algorithm 3). The Winner Determination Gadget $\mathcal{F}_{\text{WDG}}^{\mathbf{c}}$ is the ideal functionality, parameterized by $\mathbf{c}$ and by the Myerson-optimal allocation rule x^* , that behaves as follows:

1. **Input:** Each party $i \in A$ submits a private input (v_i, ρ_i) .
2. **Computation:** Let $A' = \{i \in A : c_i = \text{Commit}(v_i, \rho_i)\}$ be the set of parties whose opening is consistent with their commitment. Parties in $A \setminus A'$ are treated as having contributed no input. The functionality computes the allocation $x^*(\mathbf{v}_{A'})$ and the identity i^* of the winner, if any.
3. **Output:** Every party $i \in A$ receives the pair $(i^*, A \setminus A')$, where $i^* = \perp$ if the item is unallocated.

Two features of Definition 21 deserve comment. First, the winner's identity is a *public* output, which the protocol needs – Algorithm 3 must detect when the provisional winner refuses to reveal, and the penalty can only be charged to a party the honest bidders can name – and which is harmless for credibility: wherever a real bidder wins, every shill is outbid and revealing weakly dominates aborting (Lemma 26), so the auctioneer never acts on the extra information and Lemma 32 is unaffected. Second, the functionality does not itself abort. Commitment verification is a computation on the parties' inputs, not a protocol event: a party with an inconsistent opening is treated as absent and its identity reported. Aborts are handled in the ideal model of Definition 17, where the adversary may, after learning its own output, send (abort, i^*) for a corrupted i^* , in which case all parties output $(\perp, \{i^*\})$. The gadget therefore reveals nothing about payments, or about others' valuations beyond what the allocation implies.

► **Example 22** (Walk-through of the SRA). Consider a single real bidder with valuation $v = 5$ drawn from the equal-revenue distribution $F(v) = 1 - 1/v$ on $[1, \infty)$, for which $\bar{\phi}(v) = v - (1 - F(v))/f(v) = v - v = 0$ for every v in the support. The optimal reserve is $r^* = 1$ and $\text{REV}(F) = 1$. The auctioneer introduces two shills with bids $r_1 = 2$ and $r_2 = 8$, committing them to the same transformation $\bar{\phi}$ as the real bidder. Since $\bar{\phi} \equiv 0$, all three bids carry the same ironed virtual value, so the ranking of Definition 8 is resolved entirely on the second coordinate: the item goes to the highest raw bid.

Round 1. The WDG is invoked with all three bids and identifies shill $r_2 = 8$ as the provisional winner. The auctioneer must decide between two options. *Reveal r_2 :* the shill wins, no real bidder is allocated the item, and the auctioneer collects revenue 0 – this is the cost of committing to a price floor of 8. *Hide r_2 :* abort, forfeit the penalty κ , and restart without r_2 .

Suppose the auctioneer hides r_2 . Penalty κ is burned, and the auction restarts with $\{\text{bidder}, r_1\}$.

Round 2. The WDG runs on $\{\text{bidder}, r_1\}$. Since $v = 5 > r_1 = 2$, the real bidder wins. The bidder reveals $v = 5$ and deposits 5. Then r_1 is revealed as 2. The payment is the threshold of Definition 9 – the least bid at which the bidder still wins – and since all bids tie at $\bar{\phi} \equiv 0$ that threshold is settled on the raw bid, giving $r_1 = 2$ rather than the $\bar{\phi}^{-1}(\bar{\phi}(r_1)) = 1$ a numeric inverse would suggest. The auctioneer refunds $5 - 2 = 3$, keeping revenue 2.

On this particular realization the deviation nets $2 - \kappa$, which for $\kappa = 1$ is a gain of 1 over the honest revenue – so *ex post* the abort looks attractive. The decision to hide r_2 , however, must be made *before* v is known, and at that point the auctioneer is conditioning only on the

■ **Algorithm 3** Sequential revelation auction (SRA).

Parameters: Penalty κ per abort, MPC protocol Π_{WDG} that realizes $\mathcal{F}_{\text{WDG}}$,
Myerson mechanism (x, p) .

Input: Valuations v_i from bidders $i \in A$.

Setup: Bidders $i \in A$ sample ρ_i and send $c_i = \text{Commit}(v_i, \rho_i)$ to Π_{WDG} .

Output: An ex-post allocation $x(v_A)$ and payment $p(v_A)$.

- 1 **Commit Phase:**
- 2 Each bidder $i \in A$ deposits a collateral κ in escrow
- 3 Each bidder $i \in A$ samples randomness ρ_i and broadcast $c_i = \text{Commit}(v_i, \rho_i)$
- 4 **Reveal Phase:**
- 5 Invoke Π_{WDG} with input $\{(v_i, \rho_i, c_i)\}_{i \in A}$
- 6 **if** Π_{WDG} *aborts* **then**
- 7 Let k be the faulty bidder identified by Π_{WDG}
- 8 Exclude k : $A \leftarrow A \setminus \{k\}$
- 9 **Restart from Reveal Phase** with surviving bidders A
- 10 **end**
- 11 **else**
- 12 Π_{WDG} publicly outputs the identity i^* of the *provisional winner*
- 13 **if** $i^* = \perp$ **then**
- 14 The item is unsold; the auctioneer refunds all collaterals and the protocol halts
- 15 **end**
- 16 The provisional winner reveals their commitment by broadcasting (v_{i^*}, ρ_{i^*}) , then transfers v_{i^*} to the auctioneer
- 17 **if** the *provisional winner refuses to participate* **then**
- 18 Exclude the provisional winner: $A \leftarrow A \setminus \{i^*\}$
- 19 **Restart from Reveal Phase** with remaining bidders A
- 20 **end**
- 21 **else**
- 22 The provisional winner becomes the *final winner*. All bidders $j \in A \setminus \{i^*\}$ reveal their commitment by broadcasting (v_j, ρ_j)
- 23 **for** *Each* $j \in A \setminus \{i^*\}$ **do**
- 24 **if** j *fails to reveal* **then**
- 25 Exclude j : $A \leftarrow A \setminus \{j\}$
- 26 **end**
- 27 **end**
- 28 Auctioneer refunds $v_{i^*} - p_{i^*}(v_A)$ to the winner (where A is the set of currently active bidders)
- 29 The auctioneer returns the collateral to each bidder that did not abort (i.e., $i \in A$), and **burns** the collateral for all others
- 30 **end**
- 31 **end**

event that all real bids fall below 8. In expectation over v , the continuation revenue after the abort is at most $\text{REV}(F) = 1$, so the net expected payoff from hiding is at most $1 - \kappa \leq 0$. Thus the auctioneer has no incentive to hide r_2 in the first place. Section 6 is precisely the general form of this calculation: it bounds the value of the information the auctioneer buys by aborting.

4.1 Safety

The requirement that the winner i^* deposits v_{i^*} *before* other bidders reveal their commitments is structural, not incidental. It resolves a potential conflict between safety and credibility. Specifically, if i^* could observe the other bids before committing funds, they could abort upon discovering an unfavorable price (e.g., $p_{i^*} \approx v_{i^*}$), forcing a restart. By acting first, i^* is effectively bound: an abort after this point results in forfeiture of v_{i^*} , which by individual rationality exceeds any possible payment $p_{i^*}(v)$. Thus, the auctioneer need not restart; the forfeited deposit v_{i^*} serves as sufficient revenue.

► **Theorem 23.** *The SRA satisfies safety against honest aborts.*

Proof. Recalling Definition 14, we must demonstrate that upon any abort by a non-auctioneer participant, the protocol either (i) restarts with the remaining participants $P' \subset P$, preserving the original mechanism's outcome on P' , or (ii) terminates with revenue at least that of the mechanism on P' . We classify the possible abort conditions into three disjoint cases based on the protocol's progress:

1. **Abort during Computation (WDG):** If the MPC aborts, the Identifiable Abort property ensures the deviation is attributed to a specific party k . The protocol acts as if k never participated. The protocol restarts with $P \setminus \{k\}$, satisfying the definition.
2. **Abort during Revelation (Ex-Ante):** Suppose a bidder j aborts *before* the winner's deposit is secured.
 - If j is the winner i^* , they are excluded, and the protocol restarts. The revenue equals the revenue with $P \setminus \{i^*\}$.
 - If j is a loser, they are similarly excluded. The allocation rule implies the winner remains i^* (since $v_{i^*} > v_j$), and the payment remains determined by the highest remaining losing bid. Thus, the outcome is identical to the outcome with $P \setminus \{j\}$.
3. **Abort during Revelation (Ex-Post):** Suppose the winner i^* aborts *after* depositing v_{i^*} . The protocol guarantees revenue v_{i^*} . Since $v_{i^*} \geq p_{i^*}(v)$ (by individual rationality) this ensures i^* abort can only increase the auctioneer's revenue.

In all cases, the protocol recovers at least the revenue of the Myerson mechanism on the surviving set of bidders. ◀

4.2 Characterizing Deviations

We characterize the auctioneer's safe deviations in the SRA. Unlike real bidders, whose valuations are drawn from known distributions F_i , shill bidders are entirely controlled by the auctioneer. Crucially, the auctioneer must commit not only to each shill's bid b_j , but also to a monotone non-decreasing *ironed virtual value transformation* $\bar{\phi}_j : \mathbb{R}_+ \rightarrow \mathbb{R}$ that determines how shill j 's bid is compared against real bidders in the winner determination gadget. This transformation is analogous to the ironed virtual value function for real bidders, but is chosen strategically rather than derived from a distribution.

The auctioneer's strategy thus consists of:

1. Selecting a set of shill bidders S , each $j \in S$ with bid b_j and committed virtual value transformation $\bar{\phi}_j(\cdot)$.
2. When shill $j \in S$ is identified as the winner, deciding whether to *reveal* b_j or *hide* (abort and restart without j).
3. When shill $j \in S$ is not identified as the winner, deciding whether to *reveal* b_j or *hide* (abort and restart without j).

For credibility, we require that the honest mechanism weakly dominates *all* such deviations – regardless of the choice of shill bids $\{b_j\}_{j \in S}$ and transformations $\{\bar{\phi}_j\}_{j \in S}$.

► **Definition 24** (Threshold virtual value). *For bidding profile $\mathbf{b}$ and bidder $i \in N \cup S$, the threshold virtual value is $\beta_i(\mathbf{b}) := \max_{j \neq i} \bar{\phi}_j(b_j)$, the highest ironed virtual value among all other bidders.*

We analyze the auctioneer’s decision by considering whether a shill is outbid (does not win) or wins the allocation.

► **Observation 25.** *Shill $j \in S$ is identified as the winner if and only if $\beta_j(\mathbf{b}) < \bar{\phi}_j(b_j)$ and $\bar{\phi}_j(b_j) \geq 0$, where $\mathbf{b}$ is the bidding profile of all non-aborting bidders.*

Proof. The winner determination gadget implements the Myerson optimal allocation of Definition 9, which allocates to the highest-ranked participant with non-negative ironed virtual value. Shill j is therefore the winner exactly when it outranks every other non-aborting bidder, real or shill. ◀

Refinements

This section, we show that it is without loss of generality to impose assumptions on the deviations of the auctioneer. First, we assume a shill bidder never aborts during the WDG computation phase since the protocol restarts without him. This is strategically equivalent to the shill bidder aborting immediately after the WDG output is revealed (identifying them as the winner or loser). Thus, we can focus our analysis on deviations that occur during the revelation phase.

The next Lemma argues that if a shill is outbid by a real bidder, then revealing dominates hiding.

► **Lemma 26** (Outbid shills). *For shill $j \in S$, if $\bar{\phi}_j(b_j) < \beta_j(\mathbf{b})$, then revealing b_j weakly dominates aborting.*

Proof. By Observation 25, if $\bar{\phi}_j(b_j) < \beta_j(\mathbf{b})$ then shill j is not the winner, so some other bidder $i^* \neq j$ wins and reveals first. At this point revealing b_j acts as a *price floor*: the winner pays the least bid beating $\beta_{i^*}(\mathbf{b})$, and including b_j among the competitors can only raise that maximum, hence can only raise the payment. Aborting instead removes b_j from the competitor set, weakly decreasing revenue while incurring penalty κ . ◀

► **Definition 27** (Refined safe deviation). *A refined safe deviation is a safe deviation in which, when choosing whether to reveal or abort $j \in S$, the deviation reveals b_j if any bid (shill or real) has been revealed so far, and otherwise chooses freely.*

► **Lemma 28** (Refinement). *For any safe deviation there exists a refined safe deviation with weakly greater expected revenue. Moreover, the shills of any safe deviation may be relabelled $r_1, \dots, r_k$ so that $\bar{\phi}(r_1) < \bar{\phi}(r_2) < \dots < \bar{\phi}(r_k)$.*

Proof. The ordering is immediate: the ranking is a strict total order (Definition 8), so it linearly orders any finite set of shills whatever transformations $\{\bar{\phi}_j\}_{j \in S}$ they were committed to.

For the reveal-after-reveal property, suppose some bid has been revealed. The winner is then already determined, so every remaining shill j satisfies $\bar{\phi}_j(b_j) < \beta_j(\mathbf{b})$, and Lemma 26 says revealing b_j weakly dominates aborting. ◀

4.3 DSIC Property of the SRA

We now establish that when all skills are revealed, the SRA reduces to a DSIC mechanism.

► **Theorem 29** (DSIC under full revelation). *For a refined safe deviation of the Sequential Revelation Auction, if the auctioneer reveals all skills (i.e., never aborts), then this deviation is Dominant Strategy Incentive Compatible (DSIC) for the real bidders.*

Proof. Under full reveal no participant is excluded, so the allocation is exactly that of Definition 9 applied to $A = N \cup S$: the item goes to the highest-ranked participant with non-negative ironed virtual value, and the winner pays the least bid that would still beat $\beta_i(\mathbf{b})$. Since the allocation is monotone and its payments are the induced threshold payments, so Theorem 3 gives that the deviation is DSIC for the real bidders. ◀

We close this section by verifying that honest behavior is also optimal for the real bidders in the *extensive form*, not merely in the reported values. Two steps of Algorithm 3 warrant this check: the provisional winner must escrow the full v_{i^*} before learning the price, and the losers must open commitments that serve only to raise the winner's payment.

► **Proposition 30** (Honest play is weakly dominant for real bidders). *Fix $\kappa > 0$ and a real bidder $i \in N$ participating in the SRA. Conditional on reaching any information set of Algorithm 3, following the protocol weakly dominates aborting.*

Proof. A bidder who aborts forfeits the collateral κ and is excluded from the continuation, so it suffices to show that continuing yields a non-negative payoff. There are two cases.

If i is a loser, then i receives no allocation and makes no payment whether or not they open their commitment: their payoff from continuing is 0, while aborting costs $\kappa > 0$. (Real bidders are not revenue-maximizers, so the fact that opening raises the winner's payment does not enter i 's utility.)

If i is the provisional winner, continuing means escrowing v_i and later receiving the refund $v_i - p_i(\mathbf{v}_A)$, for a net payoff of $v_i - p_i(\mathbf{v}_A) \geq 0$ by individual rationality of the Myerson payment rule: $p_i(\mathbf{v}_A)$ is the least bid beating $\beta_i(\mathbf{v}_A)$, which is at most v_i whenever i wins. Aborting instead forfeits κ and yields no allocation, for a payoff of $-\kappa < 0$. ◀

5 Example: Equal Revenue Distribution

We now illustrate credibility for the canonical equal revenue distribution $F(v) = 1 - 1/v$ ($v \geq 1$). This distribution serves as a critical stress test for credibility due to two distinct properties:

1. **Heavy Tail:** The polynomial tail $\mathbb{P}[v > t] = 1/t$ implies that extremely high valuations occur with non-negligible probability. This allows shill-bidding strategies to chase huge payouts without the risk vanishing exponentially fast.
2. **Constant Revenue:** The “equal revenue” property ($r \cdot \mathbb{P}[v \geq r] = 1$) implies that every reserve price yields the identical expected revenue of 1. The distribution is “flat” with respect to optimization.

[11] demonstrated that for deferred revelation auctions (without MPC gadgets), *no finite penalty* suffices for this distribution – the heavy tail allows geometric shill strategies to extract unbounded profit. As a warmup, we show how a penalty of $\kappa \geq 1$ prevents unbounded profit for a natural class of shill strategies under the SRA. In Section 8, we prove this result is tight.

► **Example 31** (Warmup: Equal Revenue Distribution). Consider a single real bidder with valuation drawn from the equal revenue distribution $F(v) = 1 - 1/v$ for $v \geq 1$. Suppose the auctioneer employs a safe deviation using k shill bids $r_1 < r_2 < \dots < r_k$, where they hide $r_i, \dots, r_k$ if and only if $v \in [r_{i-1}, r_i)$. With a penalty $\kappa \geq \text{REV}(F) = 1$, we show this deviation yields expected revenue at most 1.

Without loss of generality, $r_j \geq 1$ for all j (shills below the support never win). Set $r_0 = 1$ (the lower bound of the support) and $r_{k+1} = \infty$. For the equal revenue distribution, the mechanism allocates to the highest bidder (any reserve $r \geq 1$ is optimal).

When $v \in [r_{i-1}, r_i)$ the auctioneer hides $r_i, \dots, r_k$ (each incurring penalty κ), while shills $r_1, \dots, r_{i-1}$ are outbid and revealed as price floors (by lemma 26). The bidder wins and pays r_{i-1} , giving ex-post revenue:

$$R(v) = r_{i-1} - (k - i + 1) \cdot \kappa.$$

We decompose the expected revenue into gross payments minus penalties. Using $\mathbb{P}[v \in [r_{i-1}, r_i)] = 1/r_{i-1} - 1/r_i$, the gross payment term telescopes:

$$\text{Gross} = \sum_{i=1}^{k+1} \left(\frac{1}{r_{i-1}} - \frac{1}{r_i} \right) r_{i-1} = \sum_{i=1}^{k+1} \left(1 - \frac{r_{i-1}}{r_i} \right).$$

For the penalties, each shill r_j is penalized whenever it *wins*, i.e., when $v < r_j$, which occurs with probability $1 - 1/r_j$. Thus the total expected penalty is $\kappa \sum_{j=1}^k (1 - 1/r_j)$. Setting $\kappa = 1$ and separating the last term of the gross sum ($1 - r_k/r_{k+1} = 1$):

$$\mathbb{E}[R(v)] = 1 + \sum_{j=1}^k \left[\underbrace{\left(1 - \frac{r_{j-1}}{r_j} \right)}_{\text{gross from interval } j} - \underbrace{\left(1 - \frac{1}{r_j} \right)}_{\text{penalty for shill } r_j} \right] = 1 + \sum_{j=1}^k \frac{1 - r_{j-1}}{r_j}.$$

Since $r_{j-1} \geq r_0 = 1$ for all $j \geq 1$, each summand is non-positive, giving:

$$\mathbb{E}[R(v)] \leq 1 = \text{REV}(F).$$

6 Bounding Information Leakage Under Abort

By Myerson's characterization, the expected revenue of any DSIC mechanism equals its expected virtual surplus. However, this classical equivalence holds only when the expectation is taken over the full, unconditioned prior distribution. In our credibility analysis, the auctioneer's decision to abort and restart the auction hinges on a conditional event – specifically, learning that the highest real bid is below a certain threshold. When we condition on such an upper-bound event – the event E_r of Lemma 32 below, on which no real bidder outranks the shill r – the boundary terms in Myerson's integration by parts no longer vanish, and expected revenue can exceed expected virtual surplus.

The *Information Leakage Lemma* quantifies precisely how much the conditional expected revenue can exceed the conditional expected virtual surplus. Crucially, we prove that this excess “leakage” is bounded by $\sum_{i=1}^n \text{REV}(F_i)$ (the sum of the optimal monopoly revenues for each bidder). This bounds the maximum option value the auctioneer can extract from restarting, dictating the minimum penalty κ required to deter deviations.

Intuition via a single bidder

Consider a single bidder with valuation $v \sim F$, reserve price r^* , allocation rule $x(v) = \mathbb{1}_{v \geq r^*}$, and payment rule $p(v) = r^* \mathbb{1}_{v \geq r^*}$. Under DSIC, the utility of a bidder of type v is given by the envelope formula $u(v) = \int_0^v x(t) dt$, which represents the bidder's accumulated information rent. Myerson's identity shows that the unconditional expected revenue is $\text{REV}(F) = r^*(1 - F(r^*)) = \mathbb{E}[\phi(v)^+]$, where $\phi(v)$ is the virtual value. Now, condition on the upper-bound event $E = \{v \leq c\}$ for some threshold $c > r^*$. Working with unnormalized expectations throughout, the truncated expected payment is $\mathbb{E}[p(v) \mathbb{1}_{v \leq c}] = \mathbb{E}[(vx(v) - \int_0^v x(t) dt) \mathbb{1}_{v \leq c}]$. When we perform integration by parts on the truncated information rent term $\int_0^c f(v) \int_0^v x(t) dt dv$, the boundary term at $v = c$ does not vanish, yielding:

$$\int_0^c f(v) \int_0^v x(t) dt dv = F(c) \int_0^c x(t) dt - \int_0^c F(v) x(v) dv.$$

Rearranging terms, the truncated expected revenue satisfies:

$$\mathbb{E}[p(v) \mathbb{1}_{v \leq c}] = \mathbb{E}[\phi(v)x(v) \mathbb{1}_{v \leq c}] + (1 - F(c)) \int_0^c x(t) dt.$$

The second term on the right-hand side is the non-vanishing boundary term representing the information leakage: the utility $u(c) = \int_0^c x(t) dt = c - r^*$ of the boundary type c weighted by the probability of exceeding c in the unconditioned world. Crucially, this leakage is bounded by

$$(1 - F(c))(c - r^*) \leq c(1 - F(c)) \leq \text{REV}(F),$$

where the last inequality holds because posting the price c cannot beat the optimal monopoly price. The Information Leakage Lemma generalizes this analysis to n bidders.

We state the lemma for the class of mechanisms that arise as continuations of the SRA after some shills have been revealed: the Myerson optimal auction augmented with a price floor. Formally, for a shill bid $r \geq 0$ let $M^{(r)}$ denote the DSIC mechanism with allocation mrule

$$x_i(\mathbf{v}) = \mathbb{1}_{\bar{\phi}_j(v_j) < \bar{\phi}_i(v_i) \ \forall j \in N \setminus \{i\}} \cdot \mathbb{1}_{\bar{\phi}(r) < \bar{\phi}_i(v_i)} \cdot \mathbb{1}_{\bar{\phi}_i(v_i) \geq 0}$$

and the induced threshold payments, all comparisons being taken in the ranking of Definition 8. Placing r below every participant recovers the Myerson optimal auction, and $M^{(r_i)}$ is the continuation in which shills $r_1 < \dots < r_i$ have been revealed.

► **Lemma 32 (Information Leakage Lemma).** *Let there be n bidders with independent valuations drawn from distributions $F_1, \dots, F_n$ with vanishing revenue tails, $\lim_{v \rightarrow \infty} v(1 - F_i(v)) = 0$. For a shill bid r with $\bar{\phi}(r) \geq 0$, define the event*

$$E_r = \{\mathbf{v} \mid \forall i, \bar{\phi}_i(v_i) \leq \bar{\phi}(r)\}$$

on which no real bidder outranks r . Let $R(\mathbf{v}) = \sum_{i=1}^n p_i(\mathbf{v})$ and $VS(\mathbf{v}) = \sum_{i=1}^n x_i(\mathbf{v}) \bar{\phi}_i(v_i)$ denote the ex-post revenue and ex-post ironed virtual surplus, respectively, of $M^{(r')}$ for any floor $r' \leq r$. Then,

$$\mathbb{E}[R(\mathbf{v}) \mathbb{1}_{E_r}] \leq \mathbb{E}[VS(\mathbf{v}) \mathbb{1}_{E_r}] + M \cdot \mathbb{P}[E_r]$$

where $M = \sum_{i=1}^n \text{REV}(F_i)$ and $\text{REV}(F_i)$ is the expected revenue of the optimal monopoly posted price for distribution F_i .

Observe the event E_r is a product of half-lines, $E_r = \prod_{i \in N} \{v_i < c_i\}$, where c_i is the least bid at which i would overtake r . This product form is the only property of E_r the proof uses. Note that c_i need not be $\bar{\phi}_i^{-1}(\bar{\phi}(r))$: when $\bar{\phi}_i$ is flat at the level $\bar{\phi}(r)$ the tie is settled on the raw bid, and $c_i = r$. That case is exactly what the ranking is for, and the proof handles it directly.

► **Lemma 33** (Ironing Equivalence; [17]). *Let F_i be a distribution with virtual value function ϕ_i and ironed virtual value function $\bar{\phi}_i$. Let $x_i(\mathbf{v})$ be the allocation rule of the Myerson-optimal auction. Let c be any threshold that does not lie in the strict interior of an ironed interval of F_i . Then for any fixed $\mathbf{v}_{-i}$,*

$$\int_0^c x_i(v_i, \mathbf{v}_{-i}) \phi_i(v_i) f_i(v_i) dv_i = \int_0^c x_i(v_i, \mathbf{v}_{-i}) \bar{\phi}_i(v_i) f_i(v_i) dv_i.$$

This identity holds because the Myerson-optimal allocation rule x_i is constant on any ironed interval $I = [v, \bar{v}]$ of F_i (i.e., any interval on which $\bar{\phi}_i$ is flat). Consequently, on each complete ironed interval, the allocation rule $x_i(v_i, \mathbf{v}_{-i})$ can be factored out of the integral. Since $\bar{\phi}_i$ is defined as the conditional expectation of ϕ_i with respect to the distribution F_i on that interval, we have $\int_I \phi_i(v_i) f_i(v_i) dv_i = \bar{\phi}_i(v) \int_I f_i(v_i) dv_i$ for any $v \in I$. Because the threshold c never bisects an ironed interval, the domain $[0, c]$ can be perfectly partitioned into complete ironed intervals and regions where $\phi_i = \bar{\phi}_i$. Thus, the equivalence holds over the truncated domain. For the foundational proof, see [17, Section 4].

The proof of Lemma 32 is deferred to the full version.

7 Sequential Revelation Auction is Credible

We now present our main result. The warmup (Section 5) demonstrated credibility for a restricted class of strategies against the equal-revenue distribution. The Information Leakage Lemma (Section 6) established the key quantitative tool: the excess revenue an auctioneer can extract from conditional information is bounded by $\sum_{i=1}^n \text{REV}(F_i)$. We now combine these ingredients to prove credibility in full generality.

► **Theorem 34** (Credibility of the sequential revelation auction). *Consider a single-item auction with n bidders with independent valuations drawn from product distribution $\mathbf{F}$. Then the Sequential Revelation Auction (Algorithm 3) with penalty $\kappa \geq \sum_{i=1}^n \text{REV}(F_i)$ is credible.*

Proof roadmap

The proof reduces the auctioneer’s complex, adaptive strategy space to a finite number of deterministic cases, then bounds each one. The reduction has three layers:

1. **Refinement** (Lemma 28): Without loss of revenue, all shills share the common virtual value function and outbid shills are always revealed.
2. **Kuhn’s Theorem** (Lemma 38): Any adaptive (behavioral) strategy is equivalent to a mixture of pure strategies, each of which commits to a threshold: “abort all shills above r_i , reveal r_i and below.”
3. **Case analysis** (Lemma 37): For each threshold r_i , the expected revenue decomposes into virtual surplus on the “good event” (a real bidder wins) plus conditional revenue on the “bad event” (all real bidders are below r_i). The Information Leakage Lemma bounds the latter, and the penalty $\kappa \geq \sum_j \text{REV}(F_j)$ offsets it.

The central tension is a *commitment vs. free option* trade-off. Revealing a skill commits the auctioneer to a fixed price floor, whereas hiding it opens a “free option” to restart – but at cost κ . We show the option value is always bounded by $\sum_j \text{REV}(F_j)$, so the penalty neutralizes it.

We start by bounding the revenue on good events for the auctioneer – i.e., when a real bidder is guaranteed to be the winner. We show that conditioned on a real bidder being the winner, the expected revenue is equal to the expected maximum virtual surplus restricted to the event that a real bidder is the winner.

► **Lemma 35.** *Consider any safe deviation of SRA. Let r be the highest-ranked skill bid, and write $L = \bar{\phi}(r)$ for its ironed virtual value. Define the event $E = \{\mathbf{v} \in \mathbb{R}^N : \exists i \in N, \bar{\phi}_i(v_i) > L\}$ as the set of profiles where at least one real bidder outranks every skill, the comparison being taken in the ranking of Definition 8. Then the expected revenue of the safe deviation on event E equals the expected maximum virtual surplus restricted to the event E :*

$$\mathbb{E}[R(\mathbf{v}) \cdot \mathbb{1}_E] = \mathbb{E}\left[\max_{i \in N} \bar{\phi}_i(v_i) \cdot \mathbb{1}_E\right].$$

where $R(\mathbf{v}) = \sum_{i=1}^n p_i(\mathbf{v})$ denotes the ex-post revenue of the deviation.

The proof of Lemma 35 is deferred to the full version.

Bounding Refined Deviations

By Lemma 28, it suffices to consider refined safe deviations, where the auctioneer submits skill bids $r_1 < \dots < r_k$ – ordered by the ranking of Definition 8 – and processes them from highest to lowest. We define r_{k+1} to be a formal top element that every participant precedes.

Let $\mathcal{E}_j$ be the event that no real bidder outranks the j -th skill:

$$\mathcal{E}_j := \{\mathbf{v}_N \in \mathbb{R}_+^N : \bar{\phi}_i(v_i) \leq \bar{\phi}(r_j) \text{ for all } i \in N\}.$$

Because the ranking is a total order, $\mathcal{E}_j$ is the event that skill r_j is the provisional winner, by Observation 25 – not merely a superset of it. This is where the ranking earns its keep: had ties been left to a rule outside the comparison, a real bidder tied with r_j but carrying a higher raw bid would lie in $\mathcal{E}_j$ while r_j nonetheless loses, and the penalty accounting in Lemma 37 would charge for an abort that never happens. It is the exact complement of the event E of Lemma 35, and it coincides with the event E_{r_j} of the Information Leakage Lemma (Lemma 32), which is what lets the two lemmas compose in Lemma 37. Since $r_{i+1} < \dots < r_k$, the events are nested: $\mathcal{E}_{i+1} \subseteq \dots \subseteq \mathcal{E}_k$.

Let I_ℓ denote the information set where the auctioneer decides whether to reveal or abort r_ℓ . Reaching I_ℓ implies that all higher skills were aborted, and that $\mathbf{v}$ satisfies event $\mathcal{E}_\ell$. For any mechanism M , let $R(\mathbf{v}; M) = \sum_{i=1}^n p_i^M(\mathbf{v})$ denote the ex-post revenue. We write $R(\mathbf{v})$ when M is clear from context.

► **Lemma 36 (Revenue Decomposition).** *From information set I_ℓ , the ex-ante expected revenue is:*

$$\mathbb{E}[R(\mathbf{v})] = \mathbb{E}\left[VS(\mathbf{v}) \cdot \mathbb{1}_{\mathcal{E}_\ell^c}\right] + \mathbb{E}[R(\mathbf{v}) \cdot \mathbb{1}_{\mathcal{E}_\ell}]$$

The proof of Lemma 36 is deferred to the full version.

Next, we formalize the deterministic strategies. Any deterministic strategy corresponds to a case C_i for some $i \in \{0, 1, \dots, k\}$:

C_i : The auctioneer reveals skill r_i (if $i \geq 1$) and aborts all higher skills $r_{i+1}, \dots, r_k$. The case C_0 corresponds to aborting all skills.

► **Lemma 37** (Bounding Deterministic Cases). *For any $i \in \{0, 1, \dots, k\}$, the expected deviation revenue under C_i is at most $\text{OPT}(\mathbf{F})$, provided $\kappa \geq \sum_{j=1}^n \text{REV}(F_j)$.*

The proof of Lemma 37 is deferred to the full version.

Finally, we restrict the auctioneer’s adaptive strategies to pure strategies.

► **Lemma 38** (Reduction to Pure Strategies). *The maximum expected revenue achievable by any behavioral strategy β in the refined deviation space is bounded by the maximum expected revenue of the deterministic cases $\{C_0, \dots, C_k\}$.*

The proof of Lemma 38 is deferred to the full version.

Due to space constraints, the proof of Theorem 34 is deferred to the full version.

► **Theorem 39.** *For any product distribution $\mathbf{F}$, the Sequential Revelation Auction (Algorithm 3) with penalty $\kappa \geq \sum_{i=1}^n \text{REV}(F_i)$ is a credible auction protocol that implements the revenue-optimal auction. Furthermore, on the equilibrium path without skills, the protocol terminates in a constant number of rounds.*

The proof of Theorem 39 is deferred to the full version.

8 Lower Bound

The following result shows that Theorem 34 is tight: the Sequential Revelation Auction (Algorithm 3) admits profitable safe deviations when the penalty κ is smaller than the optimal revenue.

► **Theorem 40** (Profitable Safe Deviation of the SRA). *Consider the Sequential Revelation Auction with $N \geq 1$ bidders whose valuations are drawn i.i.d. from the equal revenue distribution. If the penalty per hidden skill bid satisfies $\kappa < N$, then there exists a safe deviation with expected revenue strictly greater than N . In particular, since $\sum_{i=1}^N \text{REV}(F_i) = N = \text{OPT}(\mathbf{F})$ for this distribution, the sufficient penalty of Theorem 39 is exactly tight.*

We defer the full proof to the full version.

9 Conclusion

We introduced the Sequential Revelation Auction (SRA), a protocol that achieves credibility for any product distribution by combining a minimal MPC sub-computation – the Winner Determination Gadget – with sequential revelation and economic penalties. The gadget reveals only the winner identity, never payment information, which prevents the adversary from exploiting the “free option” inherent in monolithic MPC approaches. We proved that a penalty of $\kappa \geq \sum_{i=1}^n \text{REV}(F_i)$ is sufficient for credibility across all product distributions, and that it is tight in the worst case: for the equal-revenue distribution, where the threshold coincides exactly with the optimal revenue, every smaller penalty admits a profitable deviation.

Sharper penalties and richer adversaries

Our lower bound shows $\kappa \geq \sum_{i=1}^n \text{REV}(F_i)$ is necessary *for the SRA*; whether another design – richer gadgets, or a different revelation order – is credible at a smaller penalty is open, as is tightening the constant in $\text{OPT}(\mathbf{F}) \leq \kappa \leq n \cdot \text{OPT}(\mathbf{F})$. Our round bound holds on the equilibrium path only: the worst case grows linearly in the number of aborts, so relative to [10] the SRA buys a deterministic on-path bound and non-identical distributions

without improving the worst case. We also assume independent private values and a static adversary; correlation breaks both Myerson’s benchmark and the leakage bound, and the repeated-auction setting, where the auctioneer’s information advantage compounds across rounds, is outside the single-shot notion analyzed here.

Practical implementation

Constant-round MPC with identifiable abort from MK-FHE or iO is currently impractical, so the SRA is at present a feasibility result rather than a deployable protocol. The decomposition principle, however, is also a concrete efficiency argument: the gadget evaluates a single $\arg \max$ over n committed values, a circuit of size $O(n\ell)$ for ℓ -bit bids, rather than a full Myerson circuit with ironing and threshold-payment computation. This suggests efficient instantiations may be within reach of modern MPC constructions such as garbled circuits or secret sharing over small fields, which is the most promising route to a practical credible auction.

References

- 1 Mohammad Akbarpour and Shengwu Li. Credible auctions: A trilemma. *Econometrica*, 88(2):425–467, 2020. doi:10.3982/ECTA15925.
- 2 Carsten Baum, Emmanuela Orsini, Peter Scholl, and Eduardo Soria-Vazquez. Efficient constant-round mpc with identifiable abort and public verifiability. In *Advances in Cryptology – CRYPTO 2020*, pages 549–579, Cham, 2020. Springer. doi:10.1007/978-3-030-56880-1_20.
- 3 Manuel Blum. Coin flipping by telephone. In *COMPCON Spring*, pages 133–137, San Francisco, CA, 1981. IEEE.
- 4 Peter Bogetoft, Dan Lund Christensen, Ivan Damgård, Martin Geisler, Thomas Jakobsen, Mikkel Krøigaard, Janus Dam Nielsen, Jesper Buus Nielsen, Kurt Nielsen, Jakob Pagter, et al. Secure multiparty computation goes live. In *Financial Cryptography and Data Security (FC)*, pages 325–343, Berlin, Heidelberg, 2009. Springer. doi:10.1007/978-3-642-03549-4_20.
- 5 Felix Brandt. How to obtain full privacy in auctions. *International Journal of Information Security*, 5(4):201–216, 2006. doi:10.1007/s10207-006-0001-y.
- 6 Tarun Chitra, Matheus V. X. Ferreira, and Kshitij Kulkarni. Credible, optimal auctions via public broadcast. In *6th Conference on Advances in Financial Technologies (AFT 2024)*, volume 316 of *Leibniz International Proceedings in Informatics (LIPIcs)*, pages 19:1–19:16, Dagstuhl, Germany, 2024. Schloss Dagstuhl – Leibniz-Zentrum für Informatik. doi:10.4230/LIPIcs.AFT.2024.19.
- 7 Michele Ciampi, Divya Ravi, Luisa Siniscalchi, and Hendrik Waldner. Round-optimal multiparty computation with identifiable abort. In *Advances in Cryptology – EUROCRYPT 2022*, pages 335–364, Cham, 2022. Springer. doi:10.1007/978-3-031-06944-4_12.
- 8 Richard Cleve. Limits on the security of coin flips when half the processors are faulty. In *Proceedings of the Eighteenth Annual ACM Symposium on Theory of Computing (STOC)*, pages 364–369, New York, NY, USA, 1986. ACM. doi:10.1145/12130.12168.
- 9 Ran Cohen, Jack Doerner, Yashvanth Kondi, and Abhi Shelat. Secure multiparty computation with identifiable abort via vindicating release. In *Advances in Cryptology – CRYPTO 2024*, pages 36–73, Cham, 2024. Springer. doi:10.1007/978-3-031-68397-8_2.
- 10 Meryem Essaidi, Matheus V. X. Ferreira, and S. Matthew Weinberg. Credible, strategyproof, optimal, and bounded expected-round single-item auctions for all distributions. In *Proceedings of the 13th Innovations in Theoretical Computer Science Conference (ITCS)*, volume 215 of *Leibniz International Proceedings in Informatics (LIPIcs)*, pages 66:1–66:19, Dagstuhl, Germany, 2022. Schloss Dagstuhl – Leibniz-Zentrum für Informatik. doi:10.4230/LIPIcs.ITCS.2022.66.
- 11 Matheus V. X. Ferreira and S. Matthew Weinberg. Credible, truthful, and two-round (optimal) auctions via cryptographic commitments. In *Proceedings of the 21st ACM Conference on Economics and Computation (EC)*, pages 683–702, New York, NY, USA, 2020. ACM. doi:10.1145/3391403.3399495.

- 12 Aadityan Ganesh and Qianfan Zhang. Truthful, credible, and optimal auctions for matroids via blockchains and commitments. In *Proceedings of the 26th ACM Conference on Economics and Computation (EC)*, pages 923–943, New York, NY, USA, 2025. ACM. doi:10.1145/3736252.3742652.
- 13 Sanjam Garg, Craig Gentry, Shai Halevi, and Mariana Raykova. Two-round secure mpc from indistinguishability obfuscation. In *Theory of Cryptography Conference (TCC)*, pages 74–94, Berlin, Heidelberg, 2014. Springer. doi:10.1007/978-3-642-54242-8_4.
- 14 Oded Goldreich. *Foundations of Cryptography: Volume 2, Basic Applications*, volume 2. Cambridge University Press, Cambridge, 2004. doi:10.1017/CB09780511721656.
- 15 Yuval Ishai, Rafail Ostrovsky, and Vassilis Zikas. Secure multi-party computation with identifiable abort. In *Advances in Cryptology – CRYPTO 2014*, pages 369–386, Berlin, Heidelberg, 2014. Springer, Springer. doi:10.1007/978-3-662-44381-1_21.
- 16 Pratyay Mukherjee and Daniel Wichs. Two-round multiparty computation via multi-key fhe. In *Advances in Cryptology – EUROCRYPT 2016*, pages 735–763, Berlin, Heidelberg, 2016. Springer. doi:10.1007/978-3-662-49896-5_26.
- 17 Roger B. Myerson. Optimal auction design. *Mathematics of Operations Research*, 6(1):58–73, 1981. doi:10.1287/moor.6.1.58.
- 18 Moni Naor, Benny Pinkas, and Reuban Sumner. Privacy preserving auctions and mechanism design. In *Proceedings of the 1st ACM Conference on Electronic Commerce (EC)*, pages 129–139, Denver, CO, USA, 1999. ACM. doi:10.1145/336992.337028.
- 19 Andrew C Yao. Protocols for secure computations. In *23rd Annual Symposium on Foundations of Computer Science (FOCS)*, pages 160–164, Chicago, IL, 1982. IEEE. doi:10.1109/SFCS.1982.38.