Foundations of Platform-Assisted Auctions

Hao Chung², Ke Wu³, and Elaine Shi^{*1}

¹CMU, Oblivious Labs Inc., and 0xPARC ²CMU ³CMU and UMich

Abstract

Today, many auctions are carried out with the help of intermediary platforms like Google and eBay. These platforms serve as a rendezvous point for the buyers and sellers, and charge a fee for its service. We refer to such auctions as *platform-assisted auctions*. Traditionally, the auction theory literature mainly focuses on designing auctions that incentivize the buyers to bid truthfully, assuming that the platform always faithfully implements the auction. In practice, however, the platforms have been found to manipulate the auctions to earn more profit, resulting in high-profile anti-trust lawsuits.

We propose a new model for studying platform-assisted auctions in the permissionless setting, where anyone can register and participate in the auction. We explore whether it is possible to design a dream auction in this new model, such that honest behavior is the utility-maximizing strategy for each individual buyer, the platform, the seller, as well as platform-seller or platformbuyer coalitions. Through a collection of feasibility and infeasibility results, we carefully characterize the mathematical landscape of platform-assisted auctions.

Interestingly, our work reveals exciting connections between cryptography and mechanism design. We show how cryptography can lend to the design of an efficient platform-assisted auction with dream properties. Although a line of works have also used multi-party computation (MPC) or the blockchain to remove the reliance on a trusted auctioneer, our work is distinct in nature in several dimensions. First, we initiate a systematic exploration of the game theoretic implications when the service providers (e.g., nodes that implement the MPC or blockchain protocol) are strategic and can collude with sellers or buyers. Second, we observe that the full simulation paradigm used in the standard MPC literature is too stringent and leads to high asymptotical costs. Specifically, because every player has a different private outcome in an auction protocol, to the best of our knowledge, running any generic MPC protocol among the players would incur at least n^2 total cost where n is the number of buyers. We propose a new notion of simulation called *utility-dominated emulation* that is sufficient for guaranteeing the game-theoretic properties needed in an auction. Under this new notion of simulation, we show how to design efficient auction protocols with quasilinear efficiency, which gives an n-fold improvement over any generic approach.

^{*}Author ordering is randomized

Contents

1	Introduction	3
	1.1 New Model for Platform-Assisted Auctions	4
	1.2 Inefficient Information-Theoretic Feasibility	5
	1.3 Impossibility Results	6
	1.4 Cryptography Meets Platform-Assisted Auctions	7
	1.5 Additional Definitional Contributions	9
	1.6 Additional Related Work	9
2	Model	11
	2.1 Platform-Assisted Auction	11
	2.2 Strategy Space and Utility	12
3	Preliminaries	14
4	Incentive Compatibility Definitions	15
	4.1 Computational Incentive Compatibility	15
	4.2 Strong Computational Incentive Compatibility	16
	4.3 Information-Theoretic Incentive Compatibility	17
	4.4 Preliminary: Myerson's Lemma	18
5	Inefficient Information-Theoretic Feasibility	19
6	Impossibility Results	20
-	$6.1 \text{bIC} + 1\text{-pbIC} \Longrightarrow \text{Almost Zero Platform Revenue} \qquad \dots $	20
	6.2 bIC + 1-pbIC \implies Public Outcome Necessary	24
	6.3 bIC + 2-pbIC \implies Impossibility	27
	6.4 bIC + psIC + 1-Message \implies Small Revenue $\ldots \ldots \ldots$	30
	6.4.1 The Ex Post Setting \ldots	31
	6.4.2 The Bayesian Setting	33
	6.5 Information Theoretic bIC + psIC \implies Small Revenue	34
7	Utility-Dominated Emulation	37
	7.1 Model for the Ideal Auction	37
	7.2 Definition: Utility-Dominated Emulation	38
8	Compiling an Ideal Auction to a Real-World Protocol	40
	8.1 Compiler	40
	8.2 Proofs of Robustness	45
	8.3 Proofs of Utility-Dominated Emulation: When the Platform is Strategic	45
	8.4 Proofs of Utility Dominated Emulation: When the Platform is Honest	50
9	Ideal Auction: 2nd-Price with Reserve and Fixed Platform Fees	51
	9.1 Ideal Auction	51
	9.2 Applying the Ideal-to-Real Compiler	52
	9.3 Proof of Theorem 9.1	52

\mathbf{A}	Preliminaries: Cryptographic Building Blocks	60
	A.1 Vector Commitment	60
	A.2 Argument of Knowledge	60
	A.3 Publicly Verifiable Non-Malleable Timed Commitments	61
в	Deferred Proofs for the Ascending Auction	63

1 Introduction

Traditionally, auctions are conducted in physical spaces where an auctioneer serves as a trusted intermediary who oversees the process. Today, however, many auctions are run by online platforms such as Google Ad Exchange or eBay, which not only serve as an intermediary and a rendezvous point between sellers and buyers, but also provide value-added services to both sellers and buyers. Typically, all players agree on some *fee structure* such that the platform gets remunerated for its service.

The new status quo of online *platform-assisted auctions* raises new challenges for the auction designer. Traditionally, the auction theory literature [Mye81,NRTV07] focuses on designing auctions that incentivize the buyers to bid honestly (e.g., the good old second-price auction [Vic61]); and we take it for granted that the auctioneer will always implement the auction's rules honestly. With platform-assisted auctions, however, this "fully trusted auctioneer" assumption is no longer always true. Specifically, the platform can manipulate the auction to earn more profit. For example, if we are running a second-price auction, the platform may inflate the second price to a buyer to make it pay more; and deflate the second price to the seller to withhold some profits from the seller. Such platform deviations have been observed in the real world. In 2023, the US Department of Justice filed an antitrust suit against Google for manipulating its Ad Exchange auctions which harms users and monopolizes the market [oJ23, Mon23].

Further, platform deviations are exacerbated by the fact that modern platform-assisted auctions are often run in a *permissionless* setting. Specifically, the identities of the buyers are unknown a-priori, and anyone (including the platform or the seller) may enter the auction with fake identities and inject fake bids (also known as shill bidding [CK04]). By contrast, the vast majority of works in the traditional auction theory literature focus on *permissioned* auctions where the identities of the buyers are known a-priori. The permissionless nature of platform-assisted auctions broadens the strategy space and is a new challenge we need to tackle in designing modern platform-assisted auctions.

Can we have a dream platform-assisted auction? In this paper, we want to rethink the design of platform-assisted auctions in light of these new requirements. We consider a scenario where a seller wants to sell a finite number of k identical items to buyers, assisted by a platform. We ask what is a *dream* platform-assisted auction, and whether we can design a platform-assisted auction with such dream properties. Our work systematically explores what kind of fundamental tradeoffs we face when designing platform-assisted auctions, and illustrates how platform-assisted auctions can benefit from cryptographic tools.

We argue that a dream platform-assisted auction should satisfy the following properties:

- Incentive compatibility for $X \in \{buyer, platform, seller\}$. We want that following the prescribed protocol honestly maximizes the utility for any individual buyer (or the platform, seller respectively). Henceforth, we use the shorthands *bIC*, *pIC*, and *sIC* to refer to buyer, platform, and seller incentive compatibility, respectively.
- Strategy proof when the platform colludes with $X \in \{buyer, seller\}$. To make sure that the platform cannot become a monopoly, we not only need platform incentive compatibility (pIC), we also want to make sure that strategic deviations do not help the platform even when it may collude with some of the buyers or the seller. We say that a platform-assisted auction satisfies platform-seller incentive compatibility (psIC) or c-platform-buyer incentive compatibility (c-pbIC) iff honest behavior is profit-maximizing for the platform-seller coalition, or any platform-buyer coalition that consists of at least one and at most c buyers, respectively.

Although there are some other collusion-resilience properties one can formulate, e.g., resilience against buyer-seller coalitions, our formulation prioritizes properties that mathematically capture notions of *anti-trust* against the *platform*. Therefore, we believe that defending against coalitions centered around the platform is our top concern.

Our Results and Contributions

We reveal the mathematical structure of platform-assisted auctions through a combination of novel feasibility and infeasibility results. Our results also reveal new connections between mechanism design and cryptography. On one hand, we show how cryptography can lend to the design of platform-assisted auctions; and on the other hand, we show that since mechanism design asks only for game-theoretic properties, we can avoid using generic cryptographic primitives such as multiparty computation, and instead opt for weaker notions of "simulation" that allow us to achieve asymptotical performance improvement over known generic approaches.

We now give an overview of our results and contributions.

1.1 New Model for Platform-Assisted Auctions

We formulate a new model and lay the groundwork for formally reasoning about platform-assisted auctions. We assume that each buyer and the seller communicates with the platform through a *pairwise private channel. There is no buyer-buyer or buyer-seller communication*. Besides exchanging pairwise messages between the platform and the buyers/seller, all players are additionally allowed to post messages to a *broadcast channel*¹, and everyone can read messages posted to the broadcast channel. In our feasibility results, we would like to minimize the usage of broadcast since it is an expensive operation (especially in a permissionless environment). Depending on which consensus protocol is used to realize the broadcast, every broadcast message incurs at least quadratic and possibly higher communication. Moreover, if the broadcast is instantiated with a blockchain, posting messages to the blockchain typically incurs transaction fees. On the other hand, any impossibility result that holds even in the presence of a broadcast channel should also hold when allowing direct communication among the buyers and the seller.

We stress that in certain auctions, some information about the auction (e.g., its outcomes) naturally becomes publicly visible — in such cases, we may assume that a broadcast channel exists for free for posting selected information about the auction. For example, consider an ad auction for the Super Bowl — in this case, the winners of the auction are publicly announced since anyone can view the ads on natural television.

Henceforth, we often refer to public information posted on the blockchain (or revealed to the public through a social channel) the *public outcome* of the auction.

Comparison with earlier models. Our model strictly generalizes earlier models such as the credible auction model [AL20] and blockchain-backed auctions [BK18, TAF⁺23, CFK23, MHAG22, Ope23]. Not only so, we avoid some important limitations of earlier models as mentioned below. The elegant credible auction framework [AL20] also considers a setting where buyers communicate with the platform through pairwise private channels. However, their model differs from ours in two important ways. First, their model does not allow posting messages to a broadcast channel. This makes some of their impossibility results overly pessimistic. In particular, as mentioned above, in some auctions, a social broadcast channel naturally presents itself, and makes (part of) the

¹Note that a broadcast channel with availability guarantees cannot be directly realized from a star topology where all players talk only to the platform.

auction's outcome publicly visible. Their model and their impossibility results do not necessarily extend to such auctions. Second, they assume that the platform and the seller are the same entity. In other words, an auction that is considered incentive compatible in their framework may still permit strategies where the platform cheats the seller to make additional profit — such deviations have been documented in real-life anti-trust lawsuits [Mon23].

Various works [BK18, TAF⁺23, CFK23, MHAG22, Ope23] have considered blockchain-backed auctions where all protocol messages are posted to a blockchain — such auctions can be used for selling both digital goods [MHAG22, Ope23] as well as physical items (e.g., Taylor Swift concert tickets [Pan22]). Blockchain-backed auctions suffer from high cost both in terms of bandwidth consumption (when the blockchain is actually realized with a consensus protocol) and transaction fees.

By separating the seller and the platform's roles, our modeling choices allows us to explicitly capture platform strategies that involve cheating the seller, and we can ask natural questions such as what kind of fee structure is possible between the platform and the seller. By basing our model on pairwise channels but additionally permitting broadcast messages, we can ask natural questions such as how to minimize the use of broadcast in the auction design.

1.2 Inefficient Information-Theoretic Feasibility

We show that the ascending auction with reserve and fixed platform fees (detailed in Section 5) satisfies almost all desired properties, and meanwhile achieves approximate revenue optimality. Informally speaking, the auction starts at a fixed reserve, and gradually raises the price until the number of remaining bidders is upper bounded by the number of items. The seller gets all the fees and the final outcome is sent on a broadcast channel. Since this auction does not employ any cryptography we call it an information-theoretic feasibility as stated below:

Theorem 1.1 (Informal: information-theoretic feasibility). The ascending auction with an appropriate reserve and fixed platform fees satisfies bIC, pIC, and 1-pbIC. Moreover, if buyers' true values are drawn from a natural distribution, the mechanism additionally satisfies Bayesian sIC and Bayesian psIC, and is almost revenue-optimal.

Despite these desirable properties, the ascending auction of Section 5 suffers from a few important limitations.

- 1. The platform can only get a fixed fee that is independent of the auction's revenue, and it would be nice to support other fee structures if possible.
- 2. The protocol requires that the platform post messages to a broadcast channel, which we ideally would like to avoid if possible.
- 3. The protocol satisfies only 1-pbIC, i.e., it provides incentive compatibility when the platform colludes with at most one buyer, but not more.
- 4. The protocol satisfies psIC and sIC only in the Bayesian sense, whereas we would ideally like to achieve the stronger notion of ex post incentive compatibility if possible. Roughly speaking, *Bayesian* incentive compatibility assumes that the strategic players have some a-priori belief about honest buyers' true values, but they are allowed to adaptively adjust their actions as additional information potentially gets revealed during the protocol. By contrast, *ex post* incentive compatibility requires that the strategic players are still incentivized to behave honestly even with full knowledge of honest buyers' true values.

5. Finally, the protocol's round complexity is as large as the number of possible values in the (discretized) value domain. In practice, to have sufficient precision in encoding buyers' values, we often choose the value domain to be exponentially large, thus leading to an exponentially large round complexity.

This raises the natural question, *can we avoid the above limitations*? We show that the first four limitations are in fact inherent in some sense, whereas the last limitation can be avoided through the use of cryptography. Below we elaborate on these results.

1.3 Impossibility Results

We prove a collection of impossibility results that together uncover the necessary mathematical structure of platform-assisted auctions.

Necessity of fixed platform fees. We prove that any auction that simultaneously satisfies bIC and 1-pbIC must impose a certain fee structure, i.e., the platform must get fixed fees that are independent of the auction's revenue:

Theorem 1.2 (Informal: bIC + 1-pbIC \implies fixed fee structure). Any (possibly multi-round) auction that is bIC and 1-pbIC (in the ex post setting) must pay a fixed fee to the platform that is independent of the auction's revenue (even assuming the existence of a broadcast channel). The same fee structure restriction also holds for any mechanism that simultaneously satisfies bIC, pIC, and 1-pbIC in the Bayesian setting. Further, for computationally bounded players the same fee structure restriction holds ignoring negligibly small differences.

Theorem 1.2 also explains why in the ascending auction of Theorem 1.1, the platform must get a fixed fee.

Necessity of broadcast. We prove an impossibility result which shows the significance of the broadcast channel:

Theorem 1.3 (Informal: bIC + 1-pbIC \implies broadcast necessary). Suppose we do not allow posting to a broadcast channel; and further, either there is no public-key infrastructure or the seller is offline and does not send any messages during the auction. Then, no (possibly multi-round) auction can simultaneously satisfy bIC and 1-pbIC in the ex post setting. Moreover, no auction can simultaneously satisfy Bayesian notions of bIC, pIC, and 1-pbIC. The impossibility holds even when the players are computationally bounded.

This impossibility also implies that posting to the broadcast channel is necessary in the ascending auction in Theorem 1.1.

Impossibility of 2-pbIC. We prove that it is not possible to simultaneously satisfy bIC and 2-pbIC, which explains why the ascending auction of Theorem 1.1 satisfies only 1-pbIC:

Theorem 1.4 (Informal: bIC and 2-pbIC \implies impossible). No auction can simultaneously satisfy bIC and 2-pbIC. Further, this impossibility holds even for multi-round auctions, even for Bayesian notions of bIC and 2-pbIC, even when allowing a broadcast channel, and even for computationally bounded players.

Impossibility of ex post psIC. We prove the following theorem which shows that simultaneously asking for (Bayesian) bIC and ex post psIC will severely constrain the design space such that the resulting auction must have small revenue. This partly explains why the ascending auction of Theorem 1.1 achieves only Bayesian psIC and Bayesian sIC.

Theorem 1.5 (Informal: (Bayesian) bIC + ex post psIC \implies impossible). Any (possibly multiround) auction that simultaneously satisfies information-theoretic (Bayesian) bIC and ex post psIC must be revenue dominated by a posted price auction. Moreover, this restriction holds even when allowing broadcast.

In the above, "revenue dominated by posted price" means that the auction's revenue is dominated by some posted price auction for every value vector.

Theorem 1.5 holds in the information theoretic setting, for ex post psIC, and for even multiround protocols. The proof of this theorem can be modified to show that if we restrict ourselves to one-round protocols, then the same impossibility holds even for Bayesian psIC and even for computationally bounded players.

Corollary 1.6 (Informal: extension of Theorem 1.5). Any 1-round auction that simultaneously satisfies bIC and psIC must be revenue dominated by a posted price auction; moreover, this restriction holds even for Bayesian notions of bIC and psIC, even for computationally bounded players, and even when allowing broadcast.

Interestingly, Corollary 1.6 can be viewed as a strengthening of an impossibility shown by Akbarpour and Li [AL20]. Translating their result using our terminology, they effectively proved that in the information theoretic setting, any 1-round auction that simultaneously satisfies even Bayesian notions of bIC and psIC² cannot achieve revenue optimality. We strengthen their impossibility in multiple dimensions:

- We explicitly characterize how small the revenue is by comparing the revenue with some posted price auction, whereas they only prove that any one-round, revenue-maximizing (Bayesian) bIC auction cannot additionally satisfy Bayesian psIC.
- Our impossibility holds even for computationally bounded players whereas their modeling and proofs are restricted to the information theoretic setting.

1.4 Cryptography Meets Platform-Assisted Auctions

We ask whether we can have an *efficient* auction that achieves the same incentive compatibility properties of the ascending auction of Theorem 1.1. We show that the answer is affirmative if we can employ the help of cryptography.

As mentioned, one challenge that comes with platform-assisted auctions is that there is no party entrusted to honestly implement the auction's rules. We can mitigate this challenge and restrict the strategy space by using a multi-party computation (MPC) protocol to realize a trusted party. A common misconception is that simply employing an MPC protocol would trivialize the problem of designing a platform-assisted auction, and bring us back to the classical land of auction design. This is not true partly because modern platform-assisted auctions are permissionless by nature where strategic players can inject fake bids; and partly because we ask for additional collusion resilience properties — both of these challenges are typically not considered in the classical auction theory literature. Because of these new challenges, all the impossibility results of Section 1.3 would still apply even allowing the use of MPC.

 $^{^{2}}$ Since they consider the platform and seller as a single entity, our psIC notion is aligned with their "credible" notion.

Efficiency limitations of generic MPC. As mentioned, we want to use MPC to help restrict the strategy space. However, using generic MPC techniques would incur a high cost. In particular, in an auction, all the buyers would obtain a different private outcome. To the best of our knowledge, for this setting, it is not known how to achieve generic MPC with subquadratic communication and computation cost. For example, even with Threshold Fully Homomorphic Encryption (TFHE) [AJLA⁺12,DGLS22,BGG⁺18], we would need all players to participate in a joint decryption protocol to help each buyer decrypt its outcome, resulting in at least quadratic total communication. Other state-of-the-art approaches [GLM⁺24] incur O(n|C|) total communication where n is the number of players and |C| denotes the circuit size — but for any n-ary function that must read every input, the circuit size is at least linear in n, making the total communication at least quadratic again. If we allow the use of indistinguishability obfuscation (iO) [GGH⁺13,JLS21], it is possible to adapt Hubacek and Wichs' techniques [HW15] to get a protocol with quasilinear communication; but the computation would still be quadratic.

Our approach: utility-dominated emulation. We observe that in mechanism design, since we are asking for only game-theoretic properties, we do not need the underlying MPC protocol to offer full simulation-based security. Instead, we propose a weaker notion of simulation called *utility-dominated-emulation* which suffices for mechanism design. Informally, let C denote some strategic player or coalition; the utility-dominated emulation notion asks that 1) for any real-world strategy S, there exists an ideal-world strategy S' such that C's utility in the ideal world dominates its utility in the real world; and 2) under an honest execution, all players' utilities are identically distributed in the real and ideal worlds. In this way, if honest behavior maximizes the coalition C's utility in the ideal world, honest behavior should also maximize its utility in the real world.

The utility-dominated emulation notion allows mechanism designers to focus on the game theoretic aspects of auction design without having to worry about the concrete cryptographic instantiation³. Specifically, the mechanism designer will be designing an auction in an ideal world where there is a trusted party that enforces the correct execution of the auction's rules. Because of the existence of this trusted party, the ideal auction need not employ any cryptography. We provide a light-weight cryptographic compiler such that given an ideal-world auction that satisfies the desired game theoretic properties, we can compile it to a real-world protocol that replaces the trusted party with actual cryptography, while retaining the same game-theoretic guarantees.

Instead of using generic MPC to instantiate the trusted party, the utility-dominated emulation paradigm allows us to devise a new compiler that achieves *quasi-linear* total communication and computation cost, which represents an *n*-fold efficiency improvement relative to known generic MPC protocols (or any protocol where all players post messages to the broadcast channel). Our result can be summarized with the following theorem:

Theorem 1.7 (Cryptographic auction with O(1) rounds and quasilinear efficiency). Assume that the strong repeated squaring assumption, the Decisional Diffie Hellman (DDH) and Decisional Composite Residuosity (DCR) assumptions (in suitable groups) all hold against quasi-polynomialtime adversaries. Then, assuming polynomially bounded players, there exists a platform-assisted auction that satisfies bIC, pIC, and 1-pbIC; moreover, assuming natural value distributions, the protocol additionally satisfies Bayesian psIC and Bayesian sIC. The protocol enjoys the following efficiency guarantees:

• the round complexity is O(1);

³The utility-dominated emulation paradigm applies broadly as long as the game-theoretic notion says we want the honest strategy to be the profit-maximizing for some individual or coalition, either in a Bayesian or ex post sense. It does not insist on the ideal auction being any specific auction.

- the number of bits posted to the blockchain is $\widetilde{O}_{\lambda}(1)$;
- every buyer and seller's computation and communication are $\widetilde{O}_{\lambda}(1)$; and
- the platform's computation and communication is $\widetilde{O}_{\lambda}(n)$ where n is the number of players.

In the above, we use the notation $O_{\lambda}(\cdot)$ to hide factors that depend on the security parameter λ as well as polylogarithmic factors.

In our cryptographic protocol, the strong repeated squaring is needed for ensuring an additional desirable property called robustness, i.e., as long as the platform is honest, all players will accept with all but negligible probability. If robustness is not required, we can get the same result from standard assumptions including the existence of Non-Interactive Zero-Knowledge (NIZK) and collision-resistant hashing secure against quasi-polynomially sized adversaries.

1.5 Additional Definitional Contributions

The vast majority of the mechanism design literature does not consider computationally bounded agents, and their modeling choices are often incompatible with computationally bounded players. We are faced with various definitional subtleties when trying to translate classical game theoretical concepts to a computationally bounded setting. We explicitly choose not to adopt classical modeling techniques and terminology that are incompatible with computationally bounded players, such as extensive-form game, information sets, and static protocols. Instead, we model an auction as a cryptographic protocol between players modeled as interactive Turing Machines, and we consider the round complexity of the protocol rather than the number of information sets.

One notable subtlety that arises in our modeling is how to define a suitable notion of credible [AL20] (or psIC in our terminology) that is compatible with computationally bounded players as explained below.

A computationally sound notion of "credible". The notion of "credible" (or psIC in our terminology) is proposed in the elegant work of Akbarpour and Li [AL20]. Translated to our terminology, an auction is said to be credible if the platform-seller coalition cannot benefit from any *safe* deviation. Specifically, a strategy is considered safe iff the following holds with probability 1: for every buyer, its view in the protocol has a plausible explanation there exist some inputs and random coins of all other players that are consistent with the buyer's view. However, the explanations for different buyers are allowed to be different.

Unfortunately, Akbarpour and Li's formulation of credible is incompatible with the cryptographic setting where players are computationally bounded. As a definitional contribution, we generalize the notion of credibility (or psIC) to the computationally bounded setting, by having every player explicitly output accept or reject in the protocol's syntax definition. An execution trace is considered safe if all honest players output accept. In Section 4.3, we discuss how our psIC notion compares with the "credible" notion in more detail (see paragraph entitled "Comparison of our psIC notion and credible").

1.6 Additional Related Work

Cryptography meets auction design. Several works have suggested to use multi-party computation or related cryptographic techniques to remove the reliance on a trusted auctioneer [NPS99, AN20, BCD⁺09, TAF⁺23]. Notably, MPC was deployed in real life in a Danish sugar beets auction [BCD⁺09]. The excellent survey of Alvarez and Nojoumin [AN20] also provides a more comprehensive review of this body of work. Besides cryptography, some works have also suggested the use of secure processors to remove the reliance on a trusted auctioneer $[KMS^{+}16]$.

We stress that even if the auction itself is based on a fully decentralized MPC protocol (or secure processors), it does not obviate the need for a middleman platform that provides valueadded services such as product discovery, recommender system, and serves as a rendezvous point between the sellers and the buyers. To the best of our knowledge, the aforementioned line of work focuses more on cryptographic protocol design, and does not explore the game theoretic implications when the platform can behave strategically to increase its revenue, such as injecting fake bids, or forming coalitions with either the seller or buyers.

In our work, we consider a single untrusted platform. In comparison, some earlier works consider a scenario where the platform is realized with two or three non-colluding parties (henceforth called *service providers*) [NPS99,BCD⁺09]. Even in this case, some of our impossibility results, including Theorem 1.2 and Theorem 1.4 still hold (assuming that the fees paid to the platform are divided among the parties that jointly realize the platform). Here, *c*-pbIC means that any service provider colluding with at most *c* buyers should be incentivized to behave honestly. Note that not relying on two or more non-colluding service providers also makes our approach much easier to deploy in practice.

Credible auction with cryptography or blockchain. Recent works by Ferreira et al. [FW20] and Chitra et al. [CFK23] showed that by employing either cryptographic commitments or a blockchain, we can circumvent the trilemma shown by Akbarpour and Li [AL20], and get a constant-round auction that is bIC and Bayesian psIC. Just like the original credible auction framework, these works treat the platform and the seller as the same entity, and thus their definitions do not protect against strategies where the platform cheats the seller. Further, both works suffer from *at least quadratic* total communication. Both Ferreira et al. [FW20] and Chitra et al. [CFK23] rely on collateral to prevent the auctioneer from injecting fake committed bids and refusing to open later. It might be possible to replace their commitments with timed commitments and get rid of the collateral in their work, but formalizing the resulting scheme is outside the scope of our paper.

Transaction fee mechanism design. Besides platform-assisted auctions, transaction fee mechanisms (TFMs) [Rou20,Rou21,CS23] are another example of decentralized mechanism design where miners (or consensus nodes), who partly implement the role of the auctioneer, can behave strategically and possibly form coalitions with users. TFMs in the plain model [CS23] (without cryptography) can be viewed as a restricted form of a platform-assisted auction, where 1) the protocol must be a direct revelation mechanism; 2) there is no seller and the revenue to the seller is effectively burned on the blockchain — as a result, the notions of sIC and psIC are not relevant in the TFM context. Further, our pIC notion corresponds to miner incentive compatibility, our bIC notion corresponds to user incentive compatibility, and our *c*-pbIC notion corresponds to *c*-side-contractproofness (SCP) from the TFM literature. Because our model can be seen as a generalization of both TFMs and the credible auction framework, an extra contribution we make is to *unify and elucidate the connections between these two previously separate lines of work*.

Our generalized model allows us to adopt a natural "computationally credible" relaxation (which is not applicable in TFM's stringent model), therefore, we can overcome the strong bIC + 1-pbIC impossibility in the TFM literature assuming finite number of items. We also borrow some proof techniques from the TFM literature to prove some of our impossibility results, e.g., Theorem 1.2 and Theorem 1.4.

Subsequent works [SCW23, WSC24] on TFM also showed how MPC can help circumvent some

of the impossibilities in the plain model. In comparison, these works [SCW23, WSC24] simply run a generic MPC protocol among the miners or MPC infrastructure providers, whereas our work introduces a new notion called utility-dominated emulation that allows us to achieve asymptotical efficiency improvements. While the existing works [SCW23, WSC24] rely on the full notion of simulation to abstract away the cryptography such that the mechanism designer can work in an idealized model without cryptography, our utility-dominated emulation provides the same benefits. Moreover, the existing works [SCW23, WSC24] need multiple MPC service providers among whom a threshold number must be honest, whereas our approach needs only a single untrusted platform to act as the service provider.

2 Model

2.1 Platform-Assisted Auction

We consider a platform-assisted auction for selling k identical items. Throughout the paper, we assume that k is finite. We assume that every buyer i has a non-negative true value denoted $v_i \in \mathbb{R}_{\geq 0}$ for winning an item. Each buyer has unit demand, i.e., winning more than one item brings the same value as winning exactly one item.

A (possibly multi-round and randomized) platform-assisted auction parametrized by k is a *protocol* between the *platform*, a set of *buyers*, and a *seller*:

- 1. Inputs: An honest buyer i always uses its true value v_i as input to the protocol. The platform does not have any input.
- 2. Communication: During the protocol, each player \in {buyer, seller} communicates with the platform through a *pairwise private channel*. There is no buyer-buyer communication or buyer-seller communication. Additionally, the buyers, seller, and platform may also post messages to a *broadcast channel*, henceforth also called a *blockchain*. Messages posted to the blockchain are visible to every player.
- 3. **Private outcomes**: We may assume that at the end of the main protocol, the platform sends a single message to every buyer and seller to inform the player of its private outcome. Specifically, the platform sends either \perp indicating it wants to reject the execution, or it sends a private outcome of the following format:
 - Buyer i: its private outcome is of the form (x_i, p_i) where x_i denotes the number of items allocated to buyer i, and p_i denotes its payment.
 - Seller: its private outcome is of the form (t, μ_S) where $t \leq k$ denotes the total number of items sold, and μ_S denotes the seller's revenue.
- 4. Acceptance decision: Finally, the seller, the platform, and every buyer will decide whether to *accept* or *reject* the execution.

We require the following natural guarantees:

- 1. *Correctness:* If the platform is honest, all honest players agree on whether to accept or reject with probability 1. Further, if all players are honest, then every one accepts with probability 1.
- 2. Individual rationality: If any honest buyer *i* accepts, it must output some private outcome (x_i, p_i) where $x_i \in \{0, 1\}$, and $x_i \cdot v_i p_i \ge 0$, i.e., its utility is non-negative.

- 3. Consistency on number sold: As long as the platform is honest, it must be that $t = \sum_i x_i$, that is, the number of items sold as viewed by the seller is the same as the number of items acquired by all buyers. If the platform is strategic (or part of a strategic coalition), then it must be that $t \ge \sum_i x_i$, i.e., a strategic platform must still obtain enough items from the seller to distribute to the buyers.
- 4. Budget feasibility: If the platform is honest, it must be that $\mu_{\mathcal{S}} \leq \sum_{i} p_{i}$. The difference $\sum_{i} p_{i} \mu_{\mathcal{S}}$ represents the platform's revenue.

Definition 2.1 (View of buyer and seller). Fix an execution trace, the *view* of an honest buyer or seller includes its input, all random coins it has consumed and all protocol messages it has received (including messages posted to the blockchain).

Remark 2.2 (On the use of the blockchain). For our model to be general, we allow the buyers/platform to post messages real-time to the blockchain. However, in practice, it would be desirable to minimize posting to the blockchain to save cost. For all of our feasibility results that require the use of the blockchain, we only need the platform to post a small message to the blockchain at the very end of the protocol, i.e., we do not need the capability of posting to blockchain real-time. On the other hand, all of our infeasibility results hold even when the platform and buyers are allowed to post messages to the blockchain real-time during the protocol. This makes both of our feasibility and infeasibility results stronger.

Weak symmetry. In real life, the auction may look at additional auxiliary information (e.g., buyer's cryptographic identities, time-of-arrival of the bids) besides the amount of the bids in making decisions. Henceforth, all of this additional information is referred to as the *identity* of the buyer or bid. In this paper, we focus on auctions that satisfy a weak notion of symmetry, i.e., the auction makes use of such identity information only in tie-breaking. For example, if there are multiple buyers bidding at the same amount, the auction may give preference to the bids that arrive earlier, or whose identities are associated with higher reputation. The auction's outcome should not make use of identity information in any other way. More formally, we define *weak symmetry* as follows, same as earlier works [CRS24].

Definition 2.3 (Weak symmetry). An auction is called weakly symmetric if it can always be equivalently described in the following manner: given a value vector **b** where each bid may carry some extra identity information (e.g., cryptographic keys or timestamp), the honest (possibly multi-round) auction protocol is a realization of the following functionality. First, it sorts the vector **b** by the values. During the sorting step, if multiple values are the same, then arbitrary tie-breaking rules may be applied, and the tie-breaking can depend on the identity information and can be randomized. After this sorting step, the auction's algorithms depend only on the amount of the values and their relative position in the sorted vector.

2.2 Strategy Space and Utility

We consider a *permissionless* model where a player (buyer, seller, or platform) can create one or more fake identities to participate in the auction. For our computationally IC auctions, we assume that the identity space is super-polynomial in the security parameter λ , such that a probabilistic polynomial-time (PPT) adversary can make arbitrarily many fake identities as long as the total number is bounded by its running time. For the information theoretic setting, we may assume that the identity space is infinitely large. **Buyer strategy space.** A buyer can use an arbitrary input to the protocol rather than its true value; it can also deviate from the honest protocol arbitrarily. In particular, it can take future actions in an adaptive manner based on its view in the protocol so far. A buyer can also take on multiple identities and pretend to be multiple buyers.

Seller strategy space. A strategic seller can enter the auction pretending to be one or more buyers, and using arbitrary values as inputs. It can send arbitrary messages to the platform or the blockchain, and it can adapt its future actions based on its view in the protocol so far.

Platform and coalition strategy space. A strategic platform may deviate arbitrarily from the prescribed protocol, including creating and participating under fake identities; moreover, it can adapt its actions based on its view in the protocol so far. For a platform-seller or platform-buyer coalition, the strategy space is the union of the possible strategies of all players in the coalition.

Safe and unsafe execution traces. In a strategic execution, it is not guaranteed that all honest players will decide to accept. We say that an execution trace is *safe* if all honest players accept, otherwise it is said to be *unsafe*.

Utility. If the platform is honest, then we can define utilities as follows. If the platform rejects, then every one's utility is 0. If the platform accepts, then

- a buyer *i*'s utility is $v_i p_i$ if $x_i > 0$, otherwise its utility is 0;
- the seller's utility is its total revenue $\mu_{\mathcal{S}}$; and
- the platform's utility is $\sum_{i} p_i \mu_s$. In other words, the platform is remunerated the difference between the total payment and what the seller gets.

If the platform is strategic (or part of a strategic coalition), then the platform (or the coalition) has utility 0 on any unsafe execution trace. Otherwise, if the execution trace is safe, we can compute the strategic platform's (or coalition's) utility as follows:

- For a platform-seller coalition, its utility is the sum of all honest buyers' payments;
- For a strategic platform colluding with some buyers \mathcal{B} , we can compute its utility as follows. The number of items obtained by the coalition is $x^* = t - \sum_{i \in \mathcal{H}} x_i$ where \mathcal{H} is the set of honest buyers. Let \mathcal{B}^* be the top $\min(x^*, |\mathcal{B}|)$ buyers in \mathcal{B} with the highest true values. Then, the coalition's utility is $\sum_{i \in \mathcal{H}} p_i - \mu_{\mathcal{S}} + \sum_{j \in \mathcal{B}^*} v_j$. In other words, the coalition's utility is computed assuming it allocates its items to the colluding buyers with highest true values, and any leftover item has no value.
- For a strategic platform alone, we can compute its utility in the same way as above but assuming $\mathcal{B} = \emptyset$, i.e., its utility is simply $\sum_{i \in \mathcal{H}} p_i \mu_{\mathcal{S}}$.

Note that the motivation for forcing the platform or its coalition's utility to be zero upon an unsafe trace is aligned with the existing literature on credible auctions [AL20, FW20, CFK23]. Specifically, we consider a platform that values its reputation, and thus would not want to risk taking strategies that can impair its reputation. In fact, it may also make sense to set the platform or its coalition's utility to be $-\infty$ upon an unsafe trace — see the paragraph entitled "Comparison of our psIC notion and credible" in Section 4.3 for more detailed discussions.

3 Preliminaries

A function $f : \mathbb{N} \to [0, 1]$ is called *negligible*, if for every positive polynomial $p(\cdot)$ and all sufficiently large n, it holds that f(n) < 1/p(n).

A possibly randomized machine is said to be probabilistic polynomial time (PPT), iff it completes in total work upper-bounded by some polynomial function in its input length. Let $D(\cdot)$ be a polynomially bounded function. We say that a (non-uniform) PPT machine is *depth-D-bounded*, iff on any input *a* of length |a|, its depth is at most D(|a|).

Definition 3.1 (Computational indistinguishability against depth-bounded adversaries). Let $D(\lambda)$ be a polynomially bounded functions in λ . Consider probability ensembles $\{X_{\lambda}\}_{\lambda}$ and $\{Y_{\lambda}\}_{\lambda}$ indexed by the security parameter $\lambda \in \mathbb{N}$. We say that $\{X_{\lambda}\}_{\lambda}$ and $\{Y_{\lambda}\}_{\lambda}$ are *computationally indistinguishable against depth-D adversaries*, iff for every depth-*D* bounded non-uniform PPT machine, there exists a negligible function **negl** such that for any $\lambda \in \mathbb{N}$,

$$\left| \Pr_{x \stackrel{\$}{\leftarrow} X_{\lambda}} [\mathcal{A}(1^{\lambda}, x) = 1] - \Pr_{y \stackrel{\$}{\leftarrow} Y_{\lambda}} [\mathcal{A}(1^{\lambda}, y) = 1] \right| \leq \mathsf{negl}(\lambda).$$

As a special case, if there is no depth-D restriction on the adversary, we simply say that $\{X_{\lambda}\}_{\lambda}$ and $\{Y_{\lambda}\}_{\lambda}$ are *computationally indistinguishable*.

Fact 3.2. Suppose that $\{X_{\lambda}\}_{\lambda}$ and $\{Y_{\lambda}\}_{\lambda}$ are probability ensembles taking value over [0, 1]. Suppose that $\{X_{\lambda}\}_{\lambda}$ and $\{Y_{\lambda}\}_{\lambda}$ are computationally indistinguishable for depth- $C \log(\cdot)$ -bounded adversaries where C is a suitably large universal constant. Then, there is a negligible function $negl(\cdot)$ such that for every $\lambda \in \mathbb{N}$, $|\mathbf{E}[X_{\lambda}] - \mathbf{E}[Y_{\lambda}]| \leq negl(\lambda)$.

Proof. We first prove the following useful claim.

Claim 3.3. Suppose $\{X_{\lambda}\}_{\lambda}$ and $\{Y_{\lambda}\}_{\lambda}$ are computationally indistinguishable for depth- $C \log(\cdot)$ adversaries. Then, it must be that there is a negligible function $\operatorname{negl}(\cdot)$ such that for every λ , for any $v \in [0, 1]$, $|\operatorname{CDF}_{X_{\lambda}}(v) - \operatorname{CDF}_{Y_{\lambda}}(v)| \leq \operatorname{negl}(\cdot)$ where $\operatorname{CDF}_{X}(\cdot)$ denotes the cumulative distribution function of some random variable X.

Proof. Suppose the claim is not true, i.e., there is some $\mathsf{poly}(\cdot)$ such that for infinitely many λ 's, there is some $v_{\lambda} \in [0, 1]$ such that $|\mathsf{CDF}_{X_{\lambda}}(v_{\lambda}) - \mathsf{CDF}_{Y_{\lambda}}(v_{\lambda})| > 1/\mathsf{poly}(\lambda)$. Then, we can construct a polynomially sized distinguisher \mathcal{B} of depth at most $C\log(\ell(\lambda))$ where $\ell(\lambda)$ is the bit-length of X_{λ} or Y_{λ} , that can effectively distinguish a random sample of X_{λ} or Y_{λ} with $1/\mathsf{poly}(\lambda)$ probability for infinitely many λ 's. Basically, $\mathcal{B}(1^{\lambda})$ receives a value that is sampled from either X_{λ} or Y_{λ} , if the value is at most v_{λ} , it outputs 1; else output 0. The depth of \mathcal{B} is upper bounded by the depth of a comparator, which is logarithmic in the input length. \Box

We now continue with the proof of Fact 3.2. It suffices to show that $\mathbf{E}[X_{\lambda}] \leq \mathbf{E}[Y_{\lambda}] + \mathsf{negl}(\lambda)$ since the other direction is symmetric. Consider the random variable X'_{λ} whose cumulative distribution function is defined as follows where $\mathsf{negl}(\cdot)$ is the negligible function in the above claim:

$$\mathsf{CDF}_{X'_{\lambda}}(v) = \max\left(\mathsf{CDF}_{X_{\lambda}}(v) + \mathsf{negl}(\lambda), 1\right)$$

Intuitively, X'_{λ} is obtained by forcing the largest negligible fraction of X_{λ} to 0. Because of the above claim, it holds that X'_{λ} is stochastically dominated by Y_{λ} . Therefore, we have that

$$\mathbf{E}[X_{\lambda}] \le \mathbf{E}[X_{\lambda}'] + \mathsf{negl}(\lambda) \le \mathbf{E}[Y_{\lambda}] + \mathsf{negl}(\lambda)$$

4 Incentive Compatibility Definitions

We now define our incentive compatibility (IC) notions. Throughout the paper, if the auction makes use of identity information to break ties (see also the definition of weak symmetry in Section 2.1), we require that our IC notions hold for any choice of identities.

4.1 Computational Incentive Compatibility

To define computational incentive compatibility, we assume that the execution is parameterized with a security parameter 1^{λ} and every player receives 1^{λ} as an additional input. The players are modeled as Turing Machines.

Normalized and discretized value domain. Because players are computationally bounded, we cannot have infinite precision to represent real-numbered values. We assume that every player's true value is drawn from some finite domain $\mathbb{U}_{\lambda} := \{0,1\}^{\ell(\lambda)}$ that is a subset of the non-negative reals, such that the true value can always be represented by $\ell(\lambda) \leq \operatorname{poly}(\lambda)$ bits. Since the value domain is finite, there is a maximum possible value. Therefore, by rescaling, we may assume without loss of generality that the value domain \mathbb{U}_{λ} encodes a finite subset of [0,1], written as $\mathbb{U}_{\lambda} \subset [0,1]$. Throughout the paper, we assume that $0 \in \mathbb{U}_{\lambda}$ for any λ .

We use \mathbb{U}_{λ}^{k} to denote a length-k vector where each coordinate is from \mathbb{U}_{λ} , and we use the notation \mathbb{U}_{λ}^{*} to denote an arbitrary-length vector where each coordinate is from \mathbb{U}_{λ} .

Definition 4.1 (Computational (ex post) incentive compatibility). Given an auction parametrized by λ over the value domain $\mathbb{U}_{\lambda} \subset [0, 1]$, we say that the auction satisfies (ex post) incentive compatibility w.r.t. a strategic player or coalition denoted \mathcal{C} , iff for any polynomial function $n(\cdot)$, for any probabilistic polynomial time (PPT)⁴ strategy S adopted by \mathcal{C} , there exists a negligible function $\operatorname{negl}(\cdot)$ such that for any $\lambda \in \mathbb{N}$, for any $n_H \leq n(\lambda)$ for any true value vector $\mathbf{v}_{-\mathcal{C}} \in \mathbb{U}_{\lambda}^{n_H}$ corresponding to honest buyers, and for any true value vector $\mathbf{v}_{\mathcal{C}} \in \mathbb{U}_{\lambda}^{n_C}$ of the buyers in \mathcal{C} where $n_{\mathcal{C}}$ is the number of buyers in \mathcal{C} ,

$$\mathbf{E}\left[\mathsf{util}_{\mathcal{C}}^{S}(1^{\lambda},\mathbf{v}_{-\mathcal{C}},\mathbf{v}_{\mathcal{C}})\right] \leq \mathbf{E}\left[\mathsf{util}_{\mathcal{C}}^{H}(1^{\lambda},\mathbf{v}_{-\mathcal{C}},\mathbf{v}_{\mathcal{C}})\right] + \mathsf{negl}(\lambda)$$

where H denotes the honest strategy, and moreover, the notation $\operatorname{util}_{\mathcal{C}}^{S}(1^{\lambda}, \mathbf{v}_{-\mathcal{C}}, \mathbf{v}_{\mathcal{C}})$ represents the random variable corresponding to coalition \mathcal{C} 's utility in the following randomized experiment:

- 1. execute the protocol with security parameter 1^{λ} , where C adopts the strategy S, and all remaining players act honestly with the honest buyers taking true values \mathbf{v}_{-C} as input;
- 2. output the utility of \mathcal{C} assuming that the buyers in \mathcal{C} have the true value vector⁵ \mathbf{v}_{C} .

Definition 4.2 (Computational Bayesian incentive compatibility). Given an auction parametrized by λ over the value domain $\mathbb{U}_{\lambda} \subset [0, 1]$, we say that the auction satisfies Bayesian incentive compatibility w.r.t. a strategic player or coalition denoted \mathcal{C} and the distribution \mathcal{D}_{λ} , iff for any polynomial $n(\cdot)$, for any PPT strategy S adopted by \mathcal{C} , there exists a negligible function $\mathsf{negl}(\cdot)$, such that for any λ , for any $n_H \leq n(\lambda)$, any true value vector $\mathbf{v}_{\mathcal{C}} \in \mathbb{U}_{\lambda}^{n_{\mathcal{C}}}$ corresponding to the buyers in \mathcal{C} where $n_{\mathcal{C}}$ denotes the number of buyers in \mathcal{C} ,

$$\mathbf{E}\left[\mathsf{util}_{\mathcal{C}}^{S}(1^{\lambda}, n_{H}, \mathbf{v}_{\mathcal{C}})\right] \leq \mathbf{E}\left[\mathsf{util}_{\mathcal{C}}^{H}(1^{\lambda}, n_{H}, \mathbf{v}_{\mathcal{C}})\right] + \mathsf{negl}(\lambda)$$

⁴Throughout this paper, we allow the strategic player or coalition's algorithm to be a non-uniform machine.

⁵Note that the true value vector \mathbf{v}_{C} of the buyers in C is only used to calculate C's utility, the coalition can adopt an arbitrary strategy which implies that it can submit an arbitrary bid vector.

where H denotes the honest strategy, and moreover, the notation $\operatorname{util}_{\mathcal{C}}^{S}(1^{\lambda}, n_{H}, \mathbf{v}_{\mathcal{C}})$ represents the random variable corresponding to coalition \mathcal{C} 's utility under the following execution:

- 1. sample a vector $\mathbf{v}_{-\mathcal{C}}$ of length n_H from $\mathcal{D}_{\lambda}^{n_H}$;
- 2. execute the protocol with security parameter 1^{λ} , where C adopts the strategy S, and all remaining players act honestly with the honest buyers taking true values \mathbf{v}_{-C} ;
- 3. output the utility of \mathcal{C} assuming that the buyers in \mathcal{C} have the true value vector \mathbf{v}_{C} .

Given Definitions 4.1 and 4.2, we can define the following notions:

- Computational (Bayesian) buyer incentive compatibility (bIC): Definition 4.1 (or Definition 4.2) must hold when C contains only an individual buyer assuming the platform follows the protocol honestly;
- Computational (Bayesian) seller incentive compatibility (sIC): Definition 4.1 (or Definition 4.2) must hold when C contains only the seller assuming the platform follows the protocol honestly;
- Computational (Bayesian) platform incentive compatibility (pIC): Definition 4.1 (or Definition 4.2) must hold when C contains only the platform;
- Computational (Bayesian) platform-seller incentive compatibility (psIC): Definition 4.1 (or Definition 4.2) must hold when C contains the platform and the seller;
- Computational (Bayesian) c-platform-buyer incentive compatibility (c-pbIC): Definition 4.1 (or Definition 4.2) must hold when C contains the platform and an arbitrary non-empty set of at most c buyers.

4.2 Strong Computational Incentive Compatibility

We also define a stronger variant of our computational IC properties. Intuitively, these stronger notions capture the idea that the only way to gain over the honest strategy is to break the cryptography. In other words, if the strategic player or coalition is restricted to strategies that do not involve breaking cryptography, then the $negl(\lambda)$ term in Definitions 4.1 and 4.2 would be forced to 0. More specifically, we require that any strategy that does nothing more than using non-truthful values, injecting fake bids, and dropping out should not do better than the honest strategy.

Input replacement strategies. We say that a coalition (or an individual buyer) C adopts an *(extended) input replacement* strategy if it uses as input an arbitrary bid vector $\mathbf{b}_{\mathcal{C}}$ which need not be the same as the true values $\mathbf{v}_{\mathcal{C}}$ of the buyers in C. Otherwise, C follows the honest protocol. With extended input replacement, the length of the strategic bid vector $\mathbf{b}_{\mathcal{C}}$ need not be the same as $\mathbf{v}_{\mathcal{C}}$, i.e., C may inject fake bids or have some colluding buyers drop out. By contrast, an input replacement strategy is more restrictive and requires that the two lengths be the same.

Definition 4.3 (Strong computational (Bayesian) incentive compatibility). We say that an auction satisfies strong computational (Bayesian) incentive compatibility w.r.t. the strategic player or coalition C, iff

- it satisfies Definition 4.1 (or Definition 4.2), and
- if C is restricted to PPT extended input replacement strategies, then Definition 4.1 (or Definition 4.2) is satisfied even when the $negl(\cdot)$ function is forced to 0.

Based on Definition 4.3, we can define strong computational (Bayesian) bIC, sIC, pIC, pbIC, and psIC, respectively.

Remark 4.4 (Computational IC by utility-dominated emulation implies strong computational IC.). Later in Section 7.2, we will define a notion of utility-dominated emulation which can be viewed as an alternative way to define computational IC (see Remark 7.3). This utility-dominated emulation framework also gives a new paradigm for designing and reasoning about the IC properties of cryptographic auctions. It is not hard to see that any auction that is a utility-dominated emulation of an incentive compatible ideal auction satisfies our strong computational IC notion.

4.3 Information-Theoretic Incentive Compatibility

Given an auction parametrized by λ over the family of value domains \mathbb{U}_{λ} , we say that the auction satisfies *information-theoretic* (Bayesian) incentive compatibility against some coalition C, iff Definition 4.1 (or Definition 4.2) holds but with the following modifications:

- it holds for not just PPT but even unbounded strategies;
- the $negl(\cdot)$ function in Definition 4.1 (or Definition 4.2) is forced to 0; and
- the restriction that the length of $\mathbf{v}_{-\mathcal{C}}$ and n_H are polynomially bounded is removed.

Equivalently, the definitions can also be rephrased as the following (expanding the ensemble notations for clarity):

Definition 4.5 (Information theoretic incentive compatibility). Given an auction parametrized by λ over the family of value domains \mathbb{U}_{λ} , we say that the auction satisfies *information-theoretic incentive compatibility* against some coalition C, iff for any λ , for any honest value vector $\mathbf{b}_{-C} \in \mathbb{U}^*_{\lambda}$, for any value vector $\mathbf{v}_{\mathcal{C}} \in \mathbb{U}^{|\mathcal{C}|}_{\lambda}$ of the coalition C, for any strategy S of the coalition C, $\mathbf{E}[\mathsf{util}^{\mathcal{S}}_{\mathcal{C}}(1^{\lambda}, \mathbf{b}_{-\mathcal{C}}, \mathbf{v}_{\mathcal{C}})] \leq \mathbf{E}[\mathsf{util}^{\mathcal{H}}_{\mathcal{C}}(1^{\lambda}, \mathbf{b}_{-\mathcal{C}}, \mathbf{v}_{\mathcal{C}})]$ where the notations $\mathsf{util}^{\mathcal{H}}(1^{\lambda}, \mathbf{b}_{-\mathcal{C}}, \mathbf{v}_{\mathcal{C}})$ and $\mathsf{util}^{\mathcal{S}}(1^{\lambda}, \mathbf{b}_{-\mathcal{C}}, \mathbf{v}_{\mathcal{C}})$ are defined in the same way as in Definition 4.1.

Definition 4.6 (Information theoretic Bayesian incentive compatibility). Given an auction parametrized by λ over the family of value domains \mathbb{U}_{λ} , we say that the auction satisfies *information-theoretic Bayesian incentive compatibility* against some coalition C, iff for any λ , for any n_H that corresponds to the number of honest buyers, for any value vector $\mathbf{v}_{\mathcal{C}} \in \mathbb{U}_{\lambda}^{|\mathcal{C}|}$ of the coalition C, for any strategy Sof the coalition \mathcal{C} , $\mathbf{E}[\mathsf{util}_{\mathcal{C}}^{S}(1^{\lambda}, n_H, \mathbf{v}_{\mathcal{C}})] \leq \mathbf{E}[\mathsf{util}_{\mathcal{C}}^{H}(1^{\lambda}, n_H, \mathbf{v}_{\mathcal{C}})]$ where the notations $\mathsf{util}^{H}(1^{\lambda}, n_H, \mathbf{v}_{\mathcal{C}})$ and $\mathsf{util}^{S}(1^{\lambda}, n_H, \mathbf{v}_{\mathcal{C}})$ are defined in the same way as in Definition 4.2.

Remark 4.7 (Note about non-cryptographic auctions over reals). In classical mechanism design (without cryptography), we typically consider a single auction whose value domain is over non-negative reals, and the auction is not parametrized by λ . This can be viewed as a special case of a family of auctions parametrized by λ , where $\mathbb{U}_{\lambda} = \mathbb{R}_{\geq 0}$ for any λ , and further, the auction protocol does not depend on λ — note that for the information theoretic setting, we can remove the constraint that \mathbb{U}_{λ} must be a finite domain where each value is encoded by polynomial in λ bits. In this sense, Definitions 4.5 and 4.6 can be interpreted for classical auctions over a real-valued domain as well.

Comparison of our psIC notion and credible. Our information theoretic psIC notion can be viewed as a desirable strengthening of the credible notion proposed by Akbarpour and Li [AL20]. Specifically, their notion requires that honest behavior be utility-maximizing among only strategies

that must be safe with probability 1. However, a coalition may consider adopting a strategy that risks generating an unsafe trace with extremely small probability, and the vast majority of times, the strategy results in a safe execution that benefits the coalition. In our formulation, we simply force the coalition's utility to be 0 on unsafe traces when computing the coalition's expected utility. In other words, Akbarpour and Li's formulation can also be equivalently viewed as forcing the coalition's utility to be $-\infty$ upon an unsafe trace in our framework.

Moreover, we define a safe trace as one where all honest players accept, whereas Akbarpour and Li requires a safe trace to be one where there exists a possibly different innocent explanation to every honest player. In the information theoretic setting, for any protocol that is credible by their notion, we can simply augment the protocol by having the platform providing an explanation to every honest player, and the player accepts if the explanation is valid. On the other hand, in a computationally bounded setting, searching for an innocent explanation may not be feasible within polynomial time. Therefore, our new modeling approach is necessary for the definitions to be compatible with computationally bounded agents.

4.4 Preliminary: Myerson's Lemma

We will rely on the famous Myerson's Lemma for our impossibility proofs. Below we state a version of Myerson's Lemma that holds regardless whether the value domain is continuous or discrete.

Below, we use the notation $\mathbf{x}(\mathbf{b}) \in [0,1]^{|\mathbf{b}|}$ to denote the probabilities that each buyer gets an item under the value vector \mathbf{b} (assuming that every one acts honestly). Similarly, we use the random variables $\mathbf{p}(\mathbf{b}) \in \mathbb{R}_{\geq 0}^{|\mathbf{b}|}$ and $\mu_{\mathcal{S}}(\mathbf{b}) \in \mathbb{R}_{\geq 0}$ to denote everyone's expected payment as well as the seller's revenue under \mathbf{b} . We use $x_i(\cdot)$ and $p_i(\cdot)$ to denote the *i*-th coordinate of \mathbf{x} and \mathbf{p} , that is, the *i*-th buyer's probability of getting an item and its expected payment. Note that \mathbf{x} , \mathbf{p} , and $\mu_{\mathcal{S}}$ are well-defined even for multi-round auctions.

Lemma 4.8 (Myerson's Lemma). Suppose an auction satisfies bIC under input replacement strategies, then the following must hold: for any buyer *i*, any value vector \mathbf{b}_{-i} corresponding to all other buyers,

- Monotone allocation rule: suppose $b'_i > b_i$, it must be that $x_i(\mathbf{b}_{-i}, b'_i) \ge x_i(\mathbf{b}_{-i}, b_i)$.
- Payment sandwich: suppose $b'_i > b_i$, then

$$b_i \cdot (x_i(\mathbf{b}_{-i}, b'_i) - x_i(\mathbf{b}_{-i}, b_i)) \le p_i(\mathbf{b}_{-i}, b'_i) - p_i(\mathbf{b}_{-i}, b_i) \le b'_i \cdot (x_i(\mathbf{b}_{-i}, b'_i) - x_i(\mathbf{b}_{-i}, b_i))$$

The above statement of the Myerson's Lemma (Lemma 4.8) holds even when the value domain is discrete — in this case, the variables \mathbf{b}_{-i} , b_i , and b'_i in the statement Lemma 4.8 are assumed to take values from the discrete value domain.

In the standard literature, Myerson's Lemma is sometimes stated using a unique payment rule in lieu of the payment sandwich of Lemma 4.8. Note that the unique payment rule does not hold in the case of a discrete value domain.

The Myerson's Lemma can be extended to the Bayesian setting stated as follows.

Lemma 4.9 (Myerson's Lemma in Bayesian setting). Suppose an auction satisfies Bayesian bIC under input replacement strategies w.r.t. some distribution \mathcal{D} , then for any $n \in \mathbb{N}$, and for any buyer *i*, the following must hold:

• Monotone allocation rule: suppose $b'_i > b_i$, it must be that $\underset{\mathbf{b}_{-i} \stackrel{\$}{\leftarrow} \mathcal{D}^n}{\mathbf{E}} [x_i(\mathbf{b}_{-i}, b'_i)] \ge \underset{\mathbf{b}_{-i} \stackrel{\$}{\leftarrow} \mathcal{D}^n}{\mathbf{E}} [x_i(\mathbf{b}_{-i}, b_i)].$

• Payment sandwich: suppose $b'_i > b_i$, then it must be

$$\mathbf{E}_{\mathbf{b}_{-i} \stackrel{\$}{\leftarrow} \mathcal{D}^{n}} \left[b_{i} \cdot \left(x_{i}(\mathbf{b}_{-i}, b_{i}') - x_{i}(\mathbf{b}_{-i}, b_{i}) \right) \right] \leq \mathbf{E}_{\mathbf{b}_{-i} \stackrel{\$}{\leftarrow} \mathcal{D}^{n}} \left[p_{i}(\mathbf{b}_{-i}, b_{i}') - p_{i}(\mathbf{b}_{-i}, b_{i}) \right] \\
\leq \mathbf{E}_{\mathbf{b}_{-i} \stackrel{\$}{\leftarrow} \mathcal{D}^{n}} \left[b_{i}' \cdot \left(x_{i}(\mathbf{b}_{-i}, b_{i}') - x_{i}(\mathbf{b}_{-i}, b_{i}) \right) \right].$$

5 Inefficient Information-Theoretic Feasibility

In this section, we show that the ascending auction where the platform gets a fixed fee satisfies bIC, pIC, 1-pbIC, as well as Bayesian sIC and Bayesian psIC. Like Akbarpour and Li [AL20], rather than introducing a model for continuous-time auctions, for simplicity, we shall assume a value domain \mathbb{U} that is discretized and consists of the values $0 = \theta_1 < \theta_2 < \ldots \theta_T$. Since we are in the information theoretic setting, we do not need to index the auction with the security parameter λ (see Remark 4.7).

We assume that honest buyers' true values are sampled independently from the distribution \mathcal{D} . We assume all buyers' have the same distribution \mathcal{D} for convenience, but our results readily extend to the setting where all buyers' values are sampled independently but from different distributions.

Given \mathcal{D} over \mathbb{U} with the cumulative distribution function F and probability density function f, we define the *virtual value* of θ_i as $\phi(\theta_i) := \theta_i - \frac{1 - F(\theta_i)}{f(\theta_i)} (\theta_{i+1} - \theta_i)$, and let $\phi(\theta_T) = \theta_T$. We say that the distribution \mathcal{D} is *regular*, iff $\phi(\cdot)$ is a strictly increasing function, i.e., $v > v' \Longrightarrow \phi(v) > \phi(v')$.

Ascending auction with reserve and fixed platform fee

Input: each buyer *i* has a true value v_i .

Auction protocol:

- Register: Let τ_0 be the smallest positive integer such that $\phi(\theta_{\tau_0}) \ge 0$. Every buyer *i* whose value is at least θ_{τ_0} sends register to the platform. Let *R* be the set of buyers that have registered.
- Auction: Each round $\tau = \tau_0, \tau_0 + 1, \dots, T$:
 - If a buyer *i* has not received (stop, _) from the platform, it checks whether $v_i > \theta_{\tau}$. If so, send ok to the platform; else send \perp to the platform, and stop sending messages in future rounds.
 - Let \mathcal{I} be the set of buyers that either sent \perp in round τ or failed to respond in round τ ; the platform lets $R \leftarrow R \setminus \mathcal{I}$.
 - If $|R| \leq k$, the platform posts τ to the blockchain and sends the following messages to the buyers and the seller:
 - (a) It sends (stop, 1) to all buyers in R.
 - (b) If $\tau \neq \tau_0$, the platform randomly chooses k |R| buyers from the set \mathcal{I} and sends (stop, 1) to them, and it (stop, k) to the seller. Else, it sends (stop, |R|) to the seller.
 - (c) To all remaining buyers, the platform sends (stop, 0).

At the end of round T, if |R| > k, the platform posts \perp to the blockchain.

• Acceptance: Any buyer or the seller accepts the auction iff one of the following holds:

- 1. It received $(\text{stop}, _{-})$ from the platform in some round τ , and the platform has posted the same τ to the blockchain.
- 2. It has not received (stop, _) from the platform, and the platform has posted \perp to the blockchain.
- *Outcomes:* The platform always accepts the auction and it gets no revenue (or a fixed fee that is independent of the auction). If a buyer or the seller accepts the auction, it can determine its private outcome as follows^a.
 - If the platform has posted \perp to the blockchain, then no item is sold.
 - Otherwise, let τ be the number posted to the blockchain. A buyer that has received (stop, 1) gets an item and pays θ_{τ} , otherwise it does not get an item. Let (stop, \tilde{k}) be the message the seller received from the platform earlier, the seller's private outcome is $(\tilde{k}, \tilde{k} \cdot \theta_{\tau})$.

 a Or equivalently, the platform sends the private outcomes to each player and each player rejects if the received private outcome does not agree with what it has computed on its own.

Theorem 5.1 (Ascending auction with reserve and fixed platform fee). The above ascending auction with reserve and fixed platform fee over a discretized value domain \mathbb{U} satisfies the following properties:

- it satisfies (information theoretic) bIC, pIC, and 1-pbIC regardless of the value distribution \mathcal{D} ;
- suppose buyers' true values are sampled independently from a regular distribution \mathcal{D} , then, the auction additionally satisfies (information theoretic) Bayesian sIC and Bayesian psIC;
- suppose the buyers' true values are sampled independently from a regular distribution \mathcal{D} , then, the auction is $k \cdot \text{tick}(\mathbb{U})$ -approximately revenue maximizing where $\text{tick}(\mathbb{U}) := \max_{2 \leq j \leq T} (\theta_j - \theta_{j-1})$.

In the above, ϵ -approximately revenue maximizing means that assuming that buyers' values are sampled independently from \mathcal{D} , the expected revenue is at most ϵ away from the bIC auction that maximizes the expected revenue.

The proof of Theorem 5.1 is deferred to Appendix B.

6 Impossibility Results

6.1 bIC + 1-pbIC \implies Almost Zero Platform Revenue

We first show that any auction that satisfies bIC and pbIC (even in the Bayesian sense) must suffer from almost no platform revenue. In fact, in the information theoretic setting, the platform revenue must be 0. For the computational setting, there is a small slack related to the discretization of the value domain. Typically, the value domain should have super-polynomially many ticks to have sufficient precision in the encoding scheme, in this case, the resulting slack is negligibly small.

Interpretation. In our formal model in Section 2.1, we require that the total revenue of the seller and the platform do not exceed the total payment. When this requirement is relaxed, the 0 platform revenue restriction should actually be interpreted as having *fixed platform revenue* that is independent of the auction's revenue. In practice, it is possible to charge a fixed entry fee to the seller and possibly each (winning) buyer on top of the revenue earned from the auction.

The proof techniques are inspired by the recent transaction fee mechanism literature [CS23, SCW23], and we extend the techniques to work for computationally IC and multi-round mechanisms.

Additional notations. For some countable domain $\mathbb{U} \subset \mathbb{R}$, we say that $a, b \in \mathbb{U}$ are adjacent, denoted $a \stackrel{\text{adj}}{\sim} b$ iff $\mathbb{U} \cap (a, b) = \emptyset$. We define

$$\mathsf{tick}(\mathbb{U}) := \max_{\substack{a,b \in \mathbb{U}, \ a \stackrel{\mathrm{adj}}{\sim} b}} |a - b|$$

Henceforth we will use the notation $\mathbf{x}(\mathbf{b}) \in [0, 1]^{|\mathbf{b}|}$ to denote the probabilities that each buyer gets an item under the value vector \mathbf{b} (assuming that every one acts honestly). Similarly, we use the random variables $\mathbf{p}(\mathbf{b}) \in \mathbb{R}_{\geq 0}^{|\mathbf{b}|}$ and $\mu_{\mathcal{S}}(\mathbf{b}) \in \mathbb{R}_{\geq 0}$ to denote everyone's expected payment as well as the seller's revenue under \mathbf{b} . We use $x_i(\cdot)$ and $p_i(\cdot)$ to denote the *i*-th coordinate of \mathbf{x} and \mathbf{p} , that is, the *i*-th buyer's probability of getting an item and its expected payment.

Lemma 6.1. Let $\mathbb{U} \subset \mathbb{R}^{\geq 0}$ be a countable subset of non-negative real numbers. Let $x : \mathbb{U} \to [0,1]$ be a monotonically non-decreasing function. and let $p : \mathbb{U} \to \mathbb{R}^{\geq 0}$ and $\mu : \mathbb{U} \to \mathbb{R}^{\geq 0}$ be two other functions. Suppose that for any $b, b' \in \mathbb{U}$ where b < b', the following sandwich inequality holds: $b \cdot (x(b') - x(b)) \leq p(b') - p(b) \leq b' \cdot (x(b') - x(b))$. Further, suppose that for any $b', b \in \mathbb{U}$, it holds that $\mu(b') - \mu(b) \leq p(b') - p(b) - (b \cdot (x(b') - x(b)))$. Then, it holds that for any b, b', $\mu(b') - \mu(b) \leq \text{tick}(\mathbb{U}) \cdot |x(b') - x(b)|$.

Further, if the above conditions hold for a real-valued domain $\mathbb{U} = \mathbb{R}^{\geq 0}$, then, for any b, b', $\mu(b') = \mu(b)$.

Proof. We prove it for the case when b' > b since the proof for the other direction when b' < b is similar. Let $b_0 := b$ and $b_k := b'$, and suppose $b_0 < b_1 < b_2 < \ldots < b_{k-1} < b_k$, where for every $i \in [k-1], b_i \in \mathbb{U}$, and b_1, \ldots, b_{k-1} are all the values in \mathbb{U} between b_0 and b_k . For convenience, let $\mu_i = \mu(b_i), p_i = p(b_i)$, and $x_i = x(b_i)$ for $i \in \{0, 1, \ldots, k\}$.

$$\begin{split} u_k - \mu_0 &= \mu_k - \mu_{k-1} + \mu_{k-1} - \mu_{k-2} + \ldots + \mu_1 - \mu_0 \\ &\leq \sum_{i \in [k]} \left((p_i - p_{i-1}) - b_{i-1}(x_i - x_{i-1}) \right) \\ &\leq \sum_{i \in [k]} \left(b_i(x_i - x_{i-1}) - b_{i-1}(x_i - x_{i-1}) \right) \\ &= \sum_{i \in [k]} (b_i - b_{i-1})(x_i - x_{i-1}) \\ &\leq \sum_{i \in [k]} \operatorname{tick}(\mathbb{U}) \cdot (x_i - x_{i-1}) \leq \operatorname{tick}(\mathbb{U}) \cdot (x_k - x_0). \end{split}$$

1

The above concludes the proof when \mathbb{U} is a countable subset of the non-negative reals. For the continuous case when \mathbb{U} is the non-negative reals, the conclusion that $\mu(b') = \mu(b)$ follows directly by taking tick(\mathbb{U}) to be infinitesimally small.

Theorem 6.2 (bIC + pIC + 1-pbIC \implies Almost zero platform revenue). Let Π_{λ} be an auction for selling k items satisfies strong computational Bayesian bIC, strong computational Bayesian pIC, and strong computational Bayesian 1-pbIC w.r.t \mathcal{D}_{λ} , where \mathcal{D}_{λ} is a distribution over some finite

domain $\mathbb{U}_{\lambda} \subset [0,1]$. Then, for any λ , the platform's expected revenue under a random value vector drawn from $\mathcal{D}_{\lambda}^{n}$ is at most tick $(\mathbb{U}_{\lambda}) \cdot k \cdot (\ln n + O(1))$.

Further, if Π_{λ} satisfies strong computational bIC and strong computational 1-pbIC both in the ex post setting, then, for any λ , for any input vector $\mathbf{b} \in \mathbb{U}_{\lambda}^{n}$, the platform's expected revenue is at most tick $(\mathbb{U}_{\lambda}) \cdot k \cdot (\ln n + O(1))$.

As mentioned, by rescaling, we can without loss of generality assume that \mathbb{U}_{λ} is an encoding of a subset of [0, 1]. If we use λ bits to encode evenly spaced points in [0, 1], then tick(\mathbb{U}_{λ}) = $\Theta(1/2^{\lambda})$. Therefore, intuitively, the above theorem says that as long as the encoded subspace of [0, 1] is reasonably dense without leaving large gaps in between, then the platform's revenue should be negligibly small. Note also that for the Bayesian case, the above theorem additionally needs Bayesian pIC, whereas for the ex post version, the theorem holds without requiring pIC.

Proof of Theorem 6.2. Throughout the proof, we fix an arbitrary λ and henceforth write $\mathcal{D} := \mathcal{D}_{\lambda}$ and $\mathbb{U} := \mathbb{U}_{\lambda}$ for convenience.

Bayesian setting. We first prove the Bayesian case. Let $\mu(\mathbf{b})$ denote the platform's expected revenue under the input **b** in an honest execution. Given any buyer *i*, let

$$\bar{x}_i(\cdot) := \underset{\mathbf{b}_{-i} \overset{\$}{\leftarrow} \mathcal{D}^{n-1}}{\mathbf{E}} [x_i(\cdot, \mathbf{b}_{-i})], \qquad \bar{p}_i(\cdot) := \underset{\mathbf{b}_{-i} \overset{\$}{\leftarrow} \mathcal{D}^{n-1}}{\mathbf{E}} [p_i(\cdot, \mathbf{b}_{-i})], \qquad \bar{\mu}(\cdot) := \underset{\mathbf{b}_{-i} \overset{\$}{\leftarrow} \mathcal{D}^{n-1}}{\mathbf{E}} [\mu(\cdot, \mathbf{b}_{-i})]$$

Due to strong computational bIC and Myerson's Lemma (Lemma 4.9), for any buyer *i*, and for any b' > b, it must be $b \cdot (\bar{x}_i(b') - \bar{x}_i(b)) \le \bar{p}_i(b') - \bar{p}_i(b) \le b' \cdot (\bar{x}_i(b') - \bar{x}_i(b))$, and \bar{x}_i is a monotonically increasing.

Suppose the platform colludes with a buyer with true value b. Due to strong computational Bayesian 1-pbIC, the platform's increase in expected revenue when the buyer bids b' instead of b is upper bounded by the buyer's loss in utility. Otherwise, the buyer should strategically bid b' instead of its true value b to increase the coalition's joint utility. In other words,

$$\bar{\mu}(b') - \bar{\mu}(b) \le (b \cdot \bar{x}_i(b) - \bar{p}_i(b)) - (b \cdot \bar{x}_i(b') - \bar{p}_i(b')) = \bar{p}_i(b') - \bar{p}_i(b) - b \cdot (\bar{x}_i(b') - \bar{x}_i(b))$$

By Lemma 6.1, we have $\bar{\mu}(b') - \bar{\mu}(0) \leq \operatorname{tick}(\mathbb{U}) \cdot \bar{x}_i(b')$ for any b'.

Next, by strong computational Bayesian pIC, injecting a bid 0 should not increase the platform's revenue. Thus, we have

$$\mathbf{E}_{\mathbf{b}_{-i} \stackrel{\$}{\leftarrow} \mathcal{D}^{n-1}} [\mu(0, \mathbf{b}_{-i})] \leq \mathbf{E}_{\mathbf{b}_{-i} \stackrel{\$}{\leftarrow} \mathcal{D}^{n-1}} [\mu(\mathbf{b}_{-i})]$$

Consequently, for any n, we have

$$\begin{split} \mathbf{E}_{\mathbf{b} \leftarrow \mathcal{D}^{n}} [\mu(\mathbf{b})] &= \mathbf{E}_{b_{i} \leftarrow \mathcal{D}} \mathbf{E}_{\mathbf{b}_{-i} \leftarrow \mathcal{D}^{n-1}} [\mu(b_{i}, \mathbf{b}_{-i})] \\ &\leq \mathbf{E}_{b_{i} \leftarrow \mathcal{D}} \left(\mathbf{E}_{\mathbf{b}_{-i} \leftarrow \mathcal{D}^{n-1}} [\mu(0, \mathbf{b}_{-i})] + \mathsf{tick}(\mathbb{U}) \cdot \mathbf{E}_{\mathbf{b}_{-i} \leftarrow \mathcal{D}^{n-1}} [x_{i}(b_{i}, \mathbf{b}_{-i})] \right) \\ &= \mathbf{E}_{\mathbf{b}_{-i} \leftarrow \mathcal{D}^{n-1}} [\mu(0, \mathbf{b}_{-i})] + \mathsf{tick}(\mathbb{U}) \cdot \mathbf{E}_{\mathbf{b} \leftarrow \mathcal{D}^{n}} [x_{i}(\mathbf{b})] \quad (\star) \end{split}$$

The above holds for any buyer *i*. We will pick an *i* such that $\mathbf{E}_{\mathbf{b} \leftarrow \mathcal{D}^n} [x_i(\mathbf{b})]$ is minimized — it is not hard to see that in this case, $\mathbf{E}_{\mathbf{b} \leftarrow \mathcal{D}^n} [x_i(\mathbf{b})] \leq k/n$. Therefore, we have

$$(\star) \leq \underbrace{\mathbf{E}}_{\mathbf{b}_{-i} \stackrel{\$}{\leftarrow} \mathcal{D}^{n-1}} [\mu(0, \mathbf{b}_{-i})] + \mathsf{tick}(\mathbb{U}) \cdot k/n$$
$$\leq \underbrace{\mathbf{E}}_{\mathbf{b}_{-i} \stackrel{\$}{\leftarrow} \mathcal{D}^{n-1}} [\mu(\mathbf{b}_{-i})] + \mathsf{tick}(\mathbb{U}) \cdot k/n$$

Applying the above inductively on $\mathbf{E}_{\mathbf{b}_{-i} \stackrel{\$}{\leftarrow} \mathcal{D}^{n-1}}[\mu(\mathbf{b}_{-i})]$, and using the base condition that $\mu(0) = 0$,

we have that

$$\mathop{\mathbf{E}}_{\mathbf{b} \xleftarrow{\mathbb{S}} \mathcal{D}^n} [\mu(\mathbf{b})] \leq \mathsf{tick}(\mathbb{U}) \cdot \left(\frac{k}{n} + \frac{k}{n-1} + \ldots + \frac{k}{1}\right) \leq \mathsf{tick}(\mathbb{U}) \cdot k \cdot (\ln n + O(1))$$

This concludes the proof of the Bayesian case.

Ex post setting. We now consider the ex post case. Redefine $\bar{x}_i(\mathbf{b})$, $\bar{p}_i(\mathbf{b})$, $\bar{\mu}_i(\mathbf{b})$, to be *i*'s probability of getting an item, its expected payment, and the platform's expected revenue under the value vector **b**. Using the same argument as before, we can prove that for an arbitrary \mathbf{b}_{-i} , any b', $\bar{\mu}(\mathbf{b}_{-i}, b') - \bar{\mu}(\mathbf{b}_{-i}, 0) \leq \operatorname{tick}(\mathbb{U}) \cdot \bar{x}_i(\mathbf{b}_{-i}, b')$.

Now, given some value vector $\mathbf{b} = (b_1, \ldots, b_n)$, we have

$$\begin{split} \bar{\mu}(\mathbf{b}) &\leq \bar{\mu}(\mathbf{b}_{-i}, 0) + \mathsf{tick}(\mathbb{U}) \cdot k/n \\ &\leq \bar{\mu}(\mathbf{b}_{-\{i,j\}}, 0, 0) + \mathsf{tick}(\mathbb{U}) \cdot k/n + \mathsf{tick}(\mathbb{U}) \cdot k/(n-1) \end{split}$$

where the first line chooses an *i* such that $\bar{x}_i(\mathbf{b})$ is minimized; and the second line chooses a $j \neq i$ such that $\bar{x}_j(\mathbf{b}_{-i}, 0)$ is minimized. It must be that $\bar{x}_i(\mathbf{b}) \leq k/n$; and $\bar{x}_j(\mathbf{b}_{-i}, 0) \leq k/(n-1)$.

Carrying out the above derivation inductively, we will eventually arrive at a world where everyone bids 0, and in this case, the platform's revenue is at most 0, and thus the conclusion holds. Note that in the ex post case, we do not need the strong computational pIC condition because we no longer need to remove the 0 bids.

The proofs in this section also directly imply the following corollary that says the platform's revenue must be 0 if the value domain is continuous over the non-negative reals.

Corollary 6.3. Let Π be an auction for selling k items over a continuous, real-valued domain. Suppose that Π satisfies Bayesian bIC, Bayesian pIC, and Bayesian 1-pbIC w.r.t. the distribution \mathcal{D} . Then, the platform's expected revenue under a random value vector drawn from \mathcal{D}^n is 0. Further, if Π satisfies bIC and 1-pbIC both in the ex post setting, then, the platform's revenue is 0 for any value vector.

6.2 bIC + 1-pbIC \Longrightarrow Public Outcome Necessary

Proof roadmap. We next show that any auction with non-trivial utility that satisfies bIC and 1-pbIC (even in the Bayesian sense) must make use of a broadcast channel. Informally speaking, the intuition behind our proof is as follows. We will use the expost, information theoretic setting to explain the intuition, but our actual proof later generalizes it to even Bayesian and computational settings. Suppose there is a value vector $\mathbf{b} = (b_1, \ldots, b_n)$ such that user b_i gets u > 0 utility. Imagine that the world consists of not only $\mathbf{b} = (b_1, \ldots, b_n)$ but also many buyers indexed $n+1, n+2, \ldots, n+1$ K with the true value b_i . Since the number of buyers is much greater than the number of items k, one of the buyers with value b_i , denoted U, must get smaller than u expected utility (as long as K is sufficiently large). Now, if U is colluding with the platform, the coalition can adopt the following strategy. To the buyers $[n] \setminus \{i\}$ as well as the seller, the coalition pretends that the world consists of **b** where the *i*-coordinate is replaced with U; to every other buyer $j \in \{i, n+1, \ldots, n+K\} \setminus \{U\}$, the coalition pretends that the world consists of an extremely large number of buyers all with the same value, such that with high probability, none of these buyers get an item. Note that to make this strategy work, the coalition needs to simulate many fake buyers and the seller to the buyers $j \in \{i, n+1, \ldots, n+K\} \setminus \{U\}$. Because bIC and pbIC implies that the platform's revenue must be 0 (Corollary 6.3), this strategy benefits the coalition positively by ensuring that U can get utility u. This violates 1-pbIC.

It is easy to make the above argument work if we assume that 1) the mechanism satisfied *strong* symmetry, i.e., two buyers with the same true value have the same outcome distribution; and 2) the value vector **b** consists of distinct values. A major technical challenge in our subsequent formal proofs is how to make the argument work relying only on *weak symmetry* (see Definition 2.3 of Section 2.1), and without assuming the distinctness of **b**.

Note that having a broadcast channel can make this impossibility goes away, because players can use the broadcast channel to corroborate their views in the protocol, such that the platform-buyer coalition is no longer able to present a different world to different honest players.

Ex post setting. We say that an auction Π_{λ} over the value domain $\mathbb{U}_{\lambda} \subset [0, 1]$ is non-trivial, iff there exists some polynomially bounded function $\mathsf{poly}(\cdot)$, such that for infinitely many λ 's, there exists some value vector $\mathbf{b} \in \mathbb{U}_{\lambda}^*$ whose length is polynomially bounded in λ , and the expected total buyer utility under \mathbf{b} is at least $1/\mathsf{poly}(\lambda)$.

Theorem 6.4 (bIC + pbIC \implies public outcome necessary). Let Π_{λ} be a non-trivial auction over the value domain $\mathbb{U}_{\lambda} \subset [0,1]$. Suppose that there is no public-key infrastructure (PKI) and the players are not allowed to post any public outcome to the blockchain. Further, suppose that there exists some negligible function $\operatorname{negl}(\lambda)$ such that for every λ , $\operatorname{tick}(\mathbb{U}_{\lambda}) \leq \operatorname{negl}(\lambda)$. Then, Π_{λ} cannot simultaneously satisfy strong computational bIC and strong computational 1-pbIC.

In the above, the condition $\mathsf{tick}(\mathbb{U}_{\lambda}) \leq \mathsf{negl}(\lambda)$ is very mild. For example, if \mathbb{U}_{λ} uses λ bits to encode evenly spaced points between [0,1], then $\mathsf{tick}(\mathbb{U}_{\lambda}) = \Theta(1/2^{\lambda})$. Intuitively, this condition wants the encoding of the normalized value domain [0,1] to be sufficiently dense, without leaving large gaps in between.

Proof. (of Theorem 6.4.) Due to non-triviality, there exists some polynomial $q(\cdot)$, such that for infinitely many λ 's, there exists some value vector $\mathbf{b} = (b_1, \ldots, b_n) \in \mathbb{U}^n_{\lambda}$ where n is polynomially bounded in λ , and buyer identities id_1, \ldots, id_n (henceforth called Scenario 1), and under \mathbf{b} and these identities, the *i*-th buyer has expected utility at least $1/q(\lambda)$.

Recall that the number of items is k. Therefore, if the world consists of $K = k^2 \cdot q^4(\lambda)$ bids all at b_i , then for at least one of these K buyers (henceforth denoted buyer U), its expected probability

of getting an item is at most k/K. The choice of this buyer U may depend on auxiliary identity information — in practice, this the identity can contain information like cryptographic keys, timeof-arrival, and any other information. It must be that at least half of the identities in the identity space (denoted the set \mathcal{J}) satisfy the following: for any $j \in \mathcal{J}$, if j participates in an auction with K-1 other bidders with random distinct identities, and everyone bids b_i , then j's probability of getting an item is at most 2k/K.

Let n' be the number of bids in $\mathbf{b} = (b_1, \ldots, b_n)$ that are equal to b_i . Henceforth we use the term b_i -bids to refer to the bids in \mathbf{b} that are equal to b_i , and the term non- b_i -bids to refer to the bids in \mathbf{b} that are not equal to b_i . Pick n' - 1 distinct identities $\mathcal{J}' \subset \mathcal{J}$ and a distinct $j^* \in \mathcal{J}$ that satisfy the following condition: when the world consists of (b_1, \ldots, b_n) where all non- b_i -bids take the corresponding identities of Scenario 1, and the b_i -bids take on the identities $\mathcal{J}' \cup \{j^*\}$, then j^* has at least $1/q(\lambda)$ expected utility. This is possible due to weak symmetry.

Now, imagine the following scenario.

- 1. There are *n* buyers whose true values are b_1, \ldots, b_n respectively. All the non- b_i -bids take the corresponding identities as in Scenario 1, and the b_i -bids take the identities $\mathcal{J}' \cup \{V\}$ where *V* is an arbitrary fixed identity in \mathcal{J} distinct from all the other identities that have been chosen. Henceforth *V* is also called the victim.
- 2. Additionally, there are $K' = kq^2(\lambda)$ buyers whose identities are chosen arbitrarily from \mathcal{J} as long as they are distinct from all other consumed identities; and all of these K' buyers have the true value b_i . Among these K' buyers, at least one of them denoted U must get an item with probability at most k/K'.

Suppose U and the platform are colluding. Then, the coalition, which has the advice string j^* hardwired in its algorithm, can adopt the following polynomial-time strategy:

- The platform not only acts on behalf of itself, but also simulates a buyer with identity j^* using input value b_i . In this way, the platform interacts with the first n buyers except V as well as the seller. The messages from V and the remaining K' 1 honest buyers bidding b_i are not passed into this session. If the simulated buyer j^* gets an item, this item is actually allocated to the colluding user U.
- With each of V and the remaining K' 1 players bidding b_i , the platform forks K' sessions, and each session does the following: the platform creates K 1 random fake buyer identities and all of them use the input b_i . The platform simulates (the honest behavior of) these fake buyers and the seller in its head, and interacts with the buyer.

The execution trace is safe as long as

- 1. in the sessions with V and the K' 1 buyers bidding b_i , the randomly sampled identities do not collide with honest buyers' identities (see Remark 6.5) this happens with negligible probability as long as the identity space is super-polynomial in λ ; and
- 2. V and the K' 1 buyers bidding b_i all fail to get an item in their respective sessions.

Therefore, the probability of a safe trace is at least

$$p_{ ext{safe}} \ge 1 - (2k/K) \cdot K' - \mathsf{negl}(\lambda) \ge 1 - 3/q^2(\lambda)$$

Under this strategy, U's expected gain in utility is at least

$$\frac{1}{q(\lambda)} \cdot p_{\text{safe}} - b_i \cdot \frac{k}{K'} \ge \frac{1}{q(\lambda)} \cdot (1 - 3/q^2(\lambda)) - \frac{1}{q^2(\lambda)} \ge \frac{1}{2q(\lambda)}$$

The platform's expected loss in utility is at most $\operatorname{tick}(\mathbb{U}) \cdot k \cdot (\ln(n+K')+O(1))$ due to Theorem 6.2. Due to the assumption on $\operatorname{tick}(\mathbb{U})$, we have that $1/2q(\lambda) - \operatorname{tick}(\mathbb{U}) \cdot k \cdot (\ln(n+K')+O(1)) \geq 1/\operatorname{poly}(\lambda)$ for sufficiently large λ , i.e., for infinitely many λ 's, the coalition can gain at least $1/\operatorname{poly}(\lambda)$ in expected utility through this strategy, which violates computational 1-pbIC.

Remark 6.5. We need the no PKI assumption in the above proof, since the strategic coalition needs to simulate the seller when interacting with some honest buyers — this requires that the coalition steals the identity of the seller. We stress that the proof is constructed such that the coalition need not steal any honest buyer's identity. This means that the impossibility also holds if the "no PKI" assumption is replaced with the requirement that the seller does not actively send messages in the auction. The latter is also a natural assumptions since the seller's intention is to outsource the auction to the platform.

Bayesian setting. We modify the non-trivial definition above for the Bayesian setting. We say that an auction Π_{λ} over the value domain $\mathbb{U}_{\lambda} \subset [0, 1]$ is non-trivial w.r.t. the distribution \mathcal{D}_{λ} over \mathbb{U}_{λ} , iff there exists some polynomial function $1/\mathsf{poly}(\lambda)$, such that for infinitely many λ 's, there exists $n(\lambda)$ buyer identities denoted \mathcal{I} where $n(\cdot)$ is polynomially bounded in λ , and the expected utility of all buyers in \mathcal{I} is at least $1/\mathsf{poly}(\lambda)$ when every one draws its true value independently from \mathcal{D}_{λ} .

Corollary 6.6. Suppose some auction Π_{λ} over the value domain $\mathbb{U}_{\lambda} \subset [0,1]$ is non-trivial w.r.t. some polynomial-time samplable distribution \mathcal{D}_{λ} over \mathbb{U}_{λ} , and further, there is some negligible function $\mathsf{negl}(\cdot)$ such that for every λ , $\mathsf{tick}(\mathbb{U}_{\lambda}) \leq \mathsf{negl}(\lambda)$. Suppose that there is no public-key infrastructure (PKI) and the players are not allowed to post any public outcome to the blockchain. Then, Π_{λ} cannot simultaneously satisfy strong computational Bayesian bIC, strong computational Bayesian pIC, strong computational Bayesian pbIC w.r.t. \mathcal{D}_{λ} over \mathbb{U}_{λ} .

Proof. The proof is similar to that of Theorem 6.4. Let \mathcal{J} be a set of superpolynomially many identities such that for $i \in \mathcal{J}$, when we sample K distinct identities at random including i, i's probability of getting an item is at most 2k/K assuming everyone samples their true value from \mathcal{D}_{λ} .

Now, sample $n(\lambda) - 1$ identities $\mathcal{J}' \subset \mathcal{J}$ and some fixed j^* such that when $\mathcal{J}' \cup \{j^*\}$ participate in an auction together using input values independently sampled from \mathcal{D}_{λ} , j^* has expected utility at least $1/(\operatorname{poly}(\lambda) \cdot n(\lambda))$.

Imagine a world consisting of the identities \mathcal{J}' , some victim V with an arbitrary distinct identity in \mathcal{J} , and moreover K' additional buyers with arbitrary distinct identities from \mathcal{J} . Among the last K' buyers, some buyer (henceforth called U) must have at most k/K' probability of getting an item if everyone draws their value at random from \mathcal{D}_{λ} .

Now, imagine that the platform colludes with U, and they adopt the following polynomial-time strategy:

- The platform simulates itself and a player j^* , and interacts with \mathcal{J}' ; all real and simulated buyers draw their true value independently from \mathcal{D}_{λ} ;
- With each of the victim V, as well as the K'-1 remaining buyers, the platform forks K' sessions and each session does the following: the platform creates K-1 random fake buyer identities and all of them use a random input from \mathcal{D}_{λ} . The platform simulates (the honest behavior of) these fake buyers and the seller in its head, and interacts with the buyer.

The remaining calculations are similar to that of Theorem 6.4. For the Bayesian version, we additionally need strong computational Bayesian pIC because Theorem 6.2 which we rely on needs it for the Bayesian setting. \Box

6.3 bIC + 2-pbIC \implies Impossibility

Proof roadmap. We want to show that it is not possible to ask for bIC and 2-pbIC at the same time. Our proof is inspired by the techniques of Shi et al. [SCW23]. We describe the informal intuition assuming the information theoretic and ex post settings, and assuming strong symmetry. Our actual proof later generalizes it to computational and Bayesian settings, and replaces the strong symmetry assumption with weak symmetry (see Definition 2.3).

Due to Corollary 6.3, 2-pbIC (which implies 1-pbIC) plus bIC implies that the platform's revenue is always 0. This fact, plus 2-pbIC, implies 2-bIC, that is, honest behavior maximizes the utility for any coalition of 2 buyers. Due to the elegant work of Goldberg and Hartline [GH05], we know that any 2-bIC auction must be utility-equivalent to posted price, that is, what buyer j's true value is has no effect on buyer i's utility (assuming that every one acts honestly). Therefore, we can lower buyer j's true value to 0 without affect buyer i's utility. We then argue that because of 2-pbIC, if we drop a buyer j whose true value is 0, user i's utility must be unaffected too (Lemma 6.7). If we repeat this argument and lower every buyer's true value to 0, and then make the buyer drop, we conclude that buyer i's utility is the same in a world with only buyer i, and in a world with other buyers with arbitrary values. In particular, in a crowded world with many buyers all having the same value as i, due to strong symmetry, buyer i's utility goes to 0. This means that any buyer's utility must be 0 in any scenario.

We now present the formal proofs where we remove the strong symmetry assumption and generalize the argument to computational and Bayesian settings.

Lemma 6.7. Suppose an auction parametrized by the parameter λ satisfies strong computational Bayesian 2-pbIC w.r.t \mathcal{D}_{λ} , where \mathcal{D}_{λ} is a distribution over some finite domain $\mathbb{U}_{\lambda} \subset \mathbb{R}_{\geq 0}$. Then, for any λ , for any n that is polynomially bounded in λ , for any set \mathcal{I} of n + 2 identities, for any distinct $i, j \in \mathcal{I}$ it holds that

$$\mathbf{E}_{\mathbf{b}_{-j} \stackrel{\$}{\leftarrow} \mathcal{D}_{\lambda}^{n+1}} [\mathsf{pbutil}_i(0_j, \mathbf{b}_{-j})] = \mathbf{E}_{\mathbf{b}_{-j} \stackrel{\$}{\leftarrow} \mathcal{D}_{\lambda}^{n+1}} [\mathsf{pbutil}_i(\mathbf{b}_{-j})] \tag{1}$$

where $\mathsf{pbutil}_i(\mathbf{b}'_{\mathcal{J}})$ denotes the expected joint utility of the platform and buyer *i* when bidders with identities from \mathcal{J} take on the value vector $\mathbf{b}'_{\mathcal{I}}$.

In the above lemma, when we write v_i or 0_j , the subscripts indicates the identity of the buyer. We use -j to denote $\mathcal{I} \setminus \{j\}$, and the notation \mathbf{b}_{-j} means the bids in this vector take on the identities in the set $\mathcal{I} \setminus \{j\}$.

Proof. (of Lemma 6.7) The proof is similar to the proof of Lemma 5.3 in Appendix C.3 of Shi et al. [SCW23]. Below we provide a simpler version of the proof adapted to our context.

The \geq **direction.** It is easy to show that the left-hand side is greater than or equal to the right-hand side. If the right-hand side is strictly larger, the following holds for some specific v_i : **E** $\begin{bmatrix} \mathsf{pbutil}_i(v_i, 0_j, \mathbf{b}_{-\{i,j\}}) \end{bmatrix} < \begin{bmatrix} \mathbf{E} & [\mathsf{pbutil}_i(v_i, \mathbf{b}_{-\{i,j\}}) \end{bmatrix}$. This means that if buyer *i* with $\mathbf{b}_{-\{i,j\}} \stackrel{\&}{\leftarrow} \mathcal{D}^n_{\lambda}$

true value v_i and buyer j with true value j both collude with the platform, the coalition is better off having j drop off, which violates 2-pbIC.

The \leq **direction.** Below, we focus on proving that the left-hand side is smaller than or equal to the right-hand side. Suppose the lemma is not true, i.e., there exists a λ , a set of n + 2 identities \mathcal{I} such that for some $i, j \in \mathcal{I}$, the following holds — henceforth, for convenience let $\mathcal{D} := \mathcal{D}_{\lambda}$:

$$\mathbf{E}_{\mathbf{b}_{-j} \stackrel{\$}{\leftarrow} \mathcal{D}^{n+1}} \left[\mathsf{pbutil}_i(0_j, \mathbf{b}_{-j}) \right] > \mathbf{E}_{\mathbf{b}_{-j} \stackrel{\$}{\leftarrow} \mathcal{D}^{n+1}} \left[\mathsf{pbutil}_i(\mathbf{b}_{-j}) \right].$$
(2)

If the coalition can hijack the identity j, then we can easily show that the above Equation (2) is not possible as follows. Given Equation (2), it must be that there exists some specific v_i such that

$$\mathbf{E}_{\mathbf{b}_{-\{i,j\}} \stackrel{\$}{\leftarrow} \mathcal{D}^{n}} \left[\mathsf{pbutil}_{i}(v_{i}, 0_{j}, \mathbf{b}_{-\{i,j\}}) \right] > \mathbf{E}_{\mathbf{b}_{-\{i,j\}} \stackrel{\$}{\leftarrow} \mathcal{D}^{n}} \left[\mathsf{pbutil}_{i}(v_{i}, \mathbf{b}_{-\{i,j\}}) \right]. \tag{3}$$

This means that if the buyer i with value v_i colludes with the platform, and the rest of the world consists of random bids from \mathcal{D} under the identities $-\{i, j\}$, then the coalition should simulate a fake identity j whose true value is 0, which improves their joint utility. This violates strong computational 2-pbIC.

In the remainder of the proof, we will show that the lemma holds even when the strategic coalition cannot hijack the fixed identity j. Suppose that some identity m is under the control of the strategic platform. It must be that

$$\mathbf{E}_{\mathbf{b}_{-j} \stackrel{\$}{\leftarrow} \mathcal{D}^{n+1}} \left[\mathsf{pbutil}_i(0_m, \mathbf{b}_{-\{i, j\}}) \right] \leq \mathbf{E}_{\mathbf{b}_{-j} \stackrel{\$}{\leftarrow} \mathcal{D}^{n+1}} \left[\mathsf{pbutil}_i(\mathbf{b}_{-j}) \right] < \mathbf{E}_{\mathbf{b}_{-j} \stackrel{\$}{\leftarrow} \mathcal{D}^{n+1}} \left[\mathsf{pbutil}_i(0_j, \mathbf{b}_{-j}) \right]$$

where the second inequality is due to Equation (2), and the first inequality can be shown using the same argument as above: if not, we can reach a contradiction by showing that the platform-buyer i coalition can improve their utility by injecting a fake bid 0_m .

Due to weak symmetry, the platform's expected revenue does not change when one 0 bid changes its identity from j to m. Thus, we have that $\mathbf{E}_{\mathbf{b}_{-j} \stackrel{\$}{\leftarrow} \mathcal{D}^{n+1}} [\operatorname{util}_i(0_m, \mathbf{b}_{-j})] < \mathbf{E}_{\mathbf{b}_{-j} \stackrel{\$}{\leftarrow} \mathcal{D}^{n+1}} [\operatorname{util}_i(0_j, \mathbf{b}_{-j})]$ where $\operatorname{util}_i(\mathbf{b}'_{\mathcal{J}})$ denotes buyer i's expected utility when the buyers with identities \mathcal{J} adopt the value vector $\mathbf{b}'_{\mathcal{J}}$. Due to weak symmetry, we claim that there exists some $i' \in \mathcal{I} \setminus \{i, j\}$, such that $\mathbf{E}_{\mathbf{b}_{-j} \stackrel{\$}{\leftarrow} \mathcal{D}^{n+1}} [\operatorname{util}_{i'}(0_m, \mathbf{b}_{-j})] > \mathbf{E}_{\mathbf{b}_{-j} \stackrel{\$}{\leftarrow} \mathcal{D}^{n+1}} [\operatorname{util}_{i'}(0_j, \mathbf{b}_{-j})]$. This is because by weak symmetry, the

expected total buyer utility should be the same when one 0-bid changes its identity from j to m. This implies that there exists some specific choice of $v_{i'}$, such that

$$\mathbf{E}_{\mathbf{b}_{-\{i',j\}} \stackrel{\$}{\leftarrow} \mathcal{D}^n} \left[\mathsf{pbutil}_{i'}(v_{i'}, 0_m, \mathbf{b}_{-\{i',j\}}) \right] > \mathbf{E}_{\mathbf{b}_{-\{i',j\}} \stackrel{\$}{\leftarrow} \mathcal{D}^n} \left[\mathsf{pbutil}_{i'}(v_{i'}, 0_j, \mathbf{b}_{-\{i',j\}}) \right]$$

Now, if the buyer i' has true value $v_{i'}$, the buyer j has true value 0, and both buyers collude with the platform, and the rest of the world consists of identities $\mathcal{I} \setminus \{i', j\}$ and their true values are randomly drawn from \mathcal{D} , then the coalition is better off if j simply drops offline, and the coalition simulates a fake identity m with true value 0, which violates strong computational 2-pbIC.

Lemma 6.8. Let Π_{λ} be an auction over the value domain \mathbb{U}_{λ} which satisfies strong computational Bayesian bIC and strong computational Bayesian 2-pbIC w.r.t \mathcal{D}_{λ} , which is a distribution over \mathbb{U}_{λ} . Then, for any $\lambda \in \mathbb{N}$, for any set \mathcal{I} of n identities where n is polynomially bounded in λ , and for any buyer $i \in \mathcal{I}$, it must be that

$$\mathbf{E}_{\mathbf{b} \stackrel{\$}{\leftarrow} \mathcal{D}_{\lambda}^{n}}[\mathsf{util}_{i}(\mathbf{b})] - \mathbf{E}_{v_{i} \stackrel{\$}{\leftarrow} \mathcal{D}_{\lambda}}[\mathsf{util}_{i}(v_{i})] \le 2 \cdot \mathsf{tick}(\mathbb{U}_{\lambda}) \cdot k \cdot (\ln n + O(1))$$

T

Proof. Throughout the proof, we fix an arbitrary λ , n, and an arbitrary buyer i. For simplicity, we write $\mathcal{D} = \mathcal{D}_{\lambda}$, $\mathbb{U} = \mathbb{U}_{\lambda}$, and $-j = \mathcal{I} \setminus \{j\}$ below in this proof. Given any buyer j, and any bid b_j , we define

$$\bar{x}_{j}(b_{j}) \coloneqq \mathbf{E}_{\mathbf{b}_{-j} \stackrel{\$}{\leftarrow} \mathcal{D}^{n-1}} [x_{j}(b_{j}, \mathbf{b}_{-j})]$$
$$\bar{p}_{j}(b_{j}) \coloneqq \mathbf{E}_{\mathbf{b}_{-j} \stackrel{\$}{\leftarrow} \mathcal{D}^{n-1}} [p_{j}(b_{j}, \mathbf{b}_{-j})]$$
$$\bar{\mu}(b_{j}) \coloneqq \mathbf{E}_{\mathbf{b}_{-j} \stackrel{\$}{\leftarrow} \mathcal{D}^{n-1}} [\mu(b_{j}, \mathbf{b}_{-j})]$$

Claim 6.9. For any i, j and b_j ,

$$\left| \underbrace{\mathbf{E}}_{\mathbf{b}_{-j} \overset{\$}{\leftarrow} \mathcal{D}^{n-1}} [\mathsf{util}_i(b_j, \mathbf{b}_{-j})] - \underbrace{\mathbf{E}}_{\mathbf{b}_{-j} \overset{\$}{\leftarrow} \mathcal{D}^{n-1}} [\mathsf{util}_i(0_j, \mathbf{b}_{-j})] \right| \leq 2 \cdot \mathsf{tick}(\mathbb{U}) \cdot \bar{x}_j(b_j)$$

Proof. Using the same proof as that of Theorem 6.2, due to strong computational bIC and 1-pbIC (which is implied by 2-pbIC), we have that for any b_i ,

$$|\bar{\mu}(b_j) - \bar{\mu}(0_j)| \le \mathsf{tick}(\mathbb{U}) \cdot \bar{x}_i(b_j) \tag{4}$$

Similarly, using the technique as the proof of Theorem 6.2, when the platform, i, and j form a coalition, the increase in the joint utility of the platform and i should not exceed the loss in j's utility when j bids strategically. Therefore, we can show that for any b_i, b_j ,

$$\begin{vmatrix} \mathbf{E}_{\mathbf{b}_{-\{i,j\}} \stackrel{\$}{\leftarrow} \mathcal{D}^{n-2}} \left[\mathsf{pbutil}_i(b_i, 0_j, \mathbf{b}_{-\{i,j\}}) \right] - \mathbf{E}_{\mathbf{b}_{-\{i,j\}} \stackrel{\$}{\leftarrow} \mathcal{D}^{n-2}} \left[\mathsf{pbutil}_i(b_i, b_j, \mathbf{b}_{-\{i,j\}}) \right] \\ \leq \mathsf{tick}(\mathbb{U}) \cdot \mathbf{E}_{\mathbf{b}_{-\{i,j\}} \stackrel{\$}{\leftarrow} \mathcal{D}^{n-2}} \left[x_j(b_i, b_j, \mathbf{b}_{-\{i,j\}}) \right] \end{aligned}$$

where $\mathsf{pbutil}_i(\mathbf{b})$ denotes the joint utility of the platform and buyer *i* under **b**. Thus, taking expectation over b_i , we have

$$\left| \underbrace{\mathbf{E}}_{\mathbf{b}_{-j} \stackrel{\$}{\leftarrow} \mathcal{D}^{n-1}} \left[\mathsf{pbutil}_i(0_j, \mathbf{b}_{-\{i,j\}}) \right] - \underbrace{\mathbf{E}}_{\mathbf{b}_{-j} \stackrel{\$}{\leftarrow} \mathcal{D}^{n-1}} \left[\mathsf{pbutil}_i(b_j, \mathbf{b}_{-\{i,j\}}) \right] \right| \le \mathsf{tick}(\mathbb{U}) \cdot \bar{x}_i(b_j) \tag{5}$$

The claim follows by combining Equations (4) and (5).

We now continue with the proof of Lemma 6.8. By Lemma 6.7 and Claim 6.9, we have that for all b_i , ī.

$$\mathbf{E}_{\mathbf{b}_{-j} \stackrel{\$}{\leftarrow} \mathcal{D}^{n-1}} [\mathsf{util}_i(b_j, \mathbf{b}_{-j})] - \mathbf{E}_{\mathbf{b}_{-j} \stackrel{\$}{\leftarrow} \mathcal{D}^{n-1}} [\mathsf{util}_i(\mathbf{b}_{-j})] \le 2 \cdot \mathsf{tick}(\mathbb{U}) \cdot \bar{x}_j(b_j)$$

Therefore,

T

T

$$\left| \underbrace{\mathbf{E}}_{\mathbf{b} \leftarrow \mathcal{D}^n} [\mathsf{util}_i(\mathbf{b})] - \underbrace{\mathbf{E}}_{\mathbf{b}_{-j} \leftarrow \mathcal{D}^{n-1}} [\mathsf{util}_i(\mathbf{b}_{-j})] \right| \le 2 \cdot \mathsf{tick}(\mathbb{U}) \cdot \underbrace{\mathbf{E}}_{\mathbf{b} \leftarrow \mathcal{D}^n} [x_j(\mathbf{b})]$$

If we pick $j \neq i$ to be the one such that $\mathbf{E}_{\mathbf{b} \leftarrow \mathcal{D}^n}[x_j(\mathbf{b})]$ is minimized, we have that

$$\left| \underbrace{\mathbf{E}}_{\mathbf{b} \xleftarrow{\mathbb{S}} \mathcal{D}^n} [\mathsf{util}_i(\mathbf{b})] - \underbrace{\mathbf{E}}_{\mathbf{b}_{-j} \xleftarrow{\mathbb{S}} \mathcal{D}^{n-1}} [\mathsf{util}_i(\mathbf{b}_{-j})] \right| \le 2 \cdot \mathsf{tick}(\mathbb{U}) \cdot k/(n-1)$$

Now, using the same reasoning inductively on $\mathbf{E}_{\mathbf{b}_{-j} \leftarrow \mathcal{D}^{n-1}}[\mathsf{util}_i(\mathbf{b}_{-j})]$, we get that

$$\begin{split} \mathbf{E}_{\mathbf{b} \leftarrow \mathcal{D}^n}[\mathsf{util}_i(\mathbf{b})] &- \mathbf{E}_{v_i \leftarrow \mathcal{D}}[\mathsf{util}_i(v_i)] \\ \leq 2 \cdot \mathsf{tick}(\mathbb{U}) \cdot \left(\frac{k}{n-1} + \frac{k}{n-1} + \ldots + \frac{k}{1}\right) \\ &\leq 2 \cdot \mathsf{tick}(\mathbb{U}) \cdot k \cdot (\ln n + O(1)) \end{split}$$

Theorem 6.10 (2pbIC + bIC \implies impossible). Let Π_{λ} be an auction for selling k items over the value domain $\mathbb{U}_{\lambda} \subset [0, 1]$. Suppose there exists a negligible function negl such that tick(\mathbb{U}_{λ}) \leq negl(λ) for all sufficiently large λ . We have that for any distribution \mathcal{D}_{λ} over \mathbb{U}_{λ} , if Π_{λ} is non-trivial w.r.t. \mathcal{D}_{λ} , then Π_{λ} cannot simultaneously satisfy strong computational Bayesian bIC and strong computational Bayesian 2-pbIC simultaneously w.r.t. \mathcal{D}_{λ} over \mathbb{U}_{λ} .

In the above, the non-trivial condition is defined in the same way as in Section 6.2.

Proof. (of Theorem 6.10) By weak symmetry, $\mathbf{E}_{v_i \stackrel{\$}{\leftarrow} \mathcal{D}_{\lambda}}[\mathsf{util}_i(v_i)]$ is the same no matter what the identity of the lone buyer takes. Due to Lemma 6.8, the fact that the auction is non-trivial, and the fact that $\mathsf{tick}(\mathbb{U}_{\lambda})$ is negligibly small, it must be that infinitely many λ 's, $\mathbf{E}_{v_i \stackrel{\$}{\leftarrow} \mathcal{D}_{\lambda}}[\mathsf{util}_i(v_i)] \geq 1/\mathsf{poly}(\lambda)$. Now, consider a crowded world with K buyers where K is sufficiently large. Due to Lemma 6.8, everyone's expected utility is at least $1/\mathsf{poly}(\lambda) - \mathsf{negl}(\lambda) \geq 1/\mathsf{poly}_1(\lambda)$. This means that every one has more than $1/\mathsf{poly}_2(\lambda)$ probability of getting an item. By setting K to be greater than $k \cdot \mathsf{poly}_2(\lambda)$, we get that the expected number of items allocated is strictly greater than k, which gives the contradiction.

Corollary 6.11. Let Π_{λ} be a non-trivial auction for selling k items over the value domain $\mathbb{U}_{\lambda} \subset [0, 1]$. Suppose there exists a negligible function negl such that $\operatorname{tick}(\mathbb{U}_{\lambda}) \leq \operatorname{negl}(\lambda)$ for all sufficiently large λ . Then, Π_{λ} cannot simultaneously satisfy strong computational bIC and strong computational 2-pbIC simultaneously (in the ex post setting).

Proof. Observe that in the proof of Theorem 6.10, we in fact only need that all honest buyers' values are sampled *independently*, but not necessarily from the same distribution. The ex post case is directly implied when honest buyers' bids are sampled independently but not necessarily from the same distribution, since when we fix a value vector of the honest buyers, it is the same as sampling the value vector from a deterministic distribution. \Box

6.4 bIC + psIC + 1-Message \implies Small Revenue

We say that an auction is 1-message, iff the only communication in the protocol is for each buyer to send a single message to the platform (not counting the round where the platform informs each player of their private outcome). A 1-message auction need not be a direct revelation mechanism where the buyer reveals its true value, since the message can contain arbitrary information and randomness that may depend on the buyer's true value. **Proof roadmap.** We want to show that any 1-message auction that is bIC and psIC must be revenue dominated by posted price. We now describe the informal proof roadmap. For simplicity, we assume the information theoretic and ex post settings for describing the intuition, but our formal proofs later generalize the argument to computational and Bayesian settings. First, relying on bIC and psIC, we prove that if the world has only one buyer, then the auction is revenue dominated by posted price. Then, we argue that under the value vector $\mathbf{b} = (b_1, \ldots, b_n)$, every buyer *i*'s expected pay must not exceed its expected pay if the world has only buyer *i* and no other buyers. Suppose this is not true, that is, under **b**, there is some buyer *i* who in expectation pays more than when the world has only buyer *i*. Then, the strategic platform-seller coalition can adopt the following strategy if the world consists of only buyer *i*. The coalition can simulate n - 1 fake buyers with values \mathbf{b}_{-i} , and pretend to *i* that the world is **b**. This way, the coalition will get more revenue in expectation than honest behavior, which violates psIC. We now present the formal proofs.

6.4.1 The Ex Post Setting

We begin by proving a limit on revenue for the expost setting.

Definition 6.12 (Revenue dominated by posted price). Let Π_{λ} be an auction parametrized by λ over the value domain \mathbb{U}_{λ} . We say that Π_{λ} is $\delta(\cdot)$ -revenue dominated by posted price iff every λ , there exist some r_{λ} such that for any value vector $\mathbf{b} \in \mathbb{U}_{\lambda}^*$, Π_{λ} 's expected total payment under \mathbf{b} is at most $\delta(\lambda)$ more than "the posted price auction with reserve r_{λ} " under \mathbf{b} . For the special case where $\delta(\lambda) = 0$ for all λ , we also simply say that the auction is (strictly) revenue dominated by posted price.

Theorem 6.13 (bIC + psIC + 1-message \implies small revenue). Consider any 1-message auction Π_{λ} parametrized by λ , over the value domain $\mathbb{U}_{\lambda} \subset [0, 1]$. Suppose Π_{λ} satisfies computational bIC and computational psIC simultaneously. Then, there exists a negligible function negl(·), such that for every λ , Π_{λ} is negl(λ)-revenue dominated by posted price.

Further, if "computational" is replaced with "information theoretic" above, then the auction must be strictly revenue dominated by posted price.

Proof. Consider a world with only one buyer \mathcal{B} . Henceforth let msg be the buyer's single message to the platform, and let $\operatorname{coin}_{\mathcal{P}}$ be the platform's randomness. Let $\mu(\mathsf{msg})$ be the buyer's expected payment conditioned on msg under honest execution. Let $\mu(\mathsf{msg}, \mathsf{coin}_{\mathcal{P}})$ denote the buyer's payment conditioned on msg and the platform's coins being $\operatorname{coin}_{\mathcal{P}}$. Let $\mathsf{msg}(b, *)$ denote the distribution of the buyer's message when it uses the value b as input.

Fact 6.14. Let $\{X_{\lambda}\}_{\lambda}$ a probability ensemble that takes values between [0,1]. If for some polynomially bounded function $poly(\cdot)$, $Pr[|X_{\lambda} - \mathbf{E}[X_{\lambda}]| \ge 1/poly(\lambda)] \ge 1/poly(\lambda)$ for infinitely many λ 's, then there is a polynomially bounded function $poly'(\cdot)$ such that for infinitely many λ 's,

- $\Pr\left[X_{\lambda} \mathbf{E}[X_{\lambda}] \ge 1/\mathsf{poly}'(\lambda)\right] \ge 1/\mathsf{poly}'(\lambda), and$
- $\Pr\left[\mathbf{E}[X_{\lambda}] X_{\lambda} \ge 1/\mathsf{poly}'(\lambda)\right] \ge 1/\mathsf{poly}'(\lambda).$

Claim 6.15. In a world with only one buyer, for sufficiently large λ , for any true value of the buyer drawn from \mathbb{U}_{λ} , under honest execution, except with negligible probability over the choice of msg, it holds that except with negligible probability over the choice $\operatorname{coin}_{\mathcal{P}}$, $|\mu(\operatorname{msg}, \operatorname{coin}_{\mathcal{P}})-\mu(\operatorname{msg})| \leq \operatorname{negl}(\lambda)$.

Proof. For the sake of contradiction, suppose that for infinitely many λ 's, there is some value from $v_{\lambda} \in \mathbb{U}_{\lambda}$ and a set denoted bad that contains at least $1/\operatorname{poly}(\lambda)$ fraction of msg drawn from $\operatorname{msg}(v_{\lambda}, *)$, such that at least $1/\operatorname{poly}(\lambda)$ fraction of $\operatorname{coin}_{\mathcal{P}}$ would cause the buyer's payment to have at least $1/\operatorname{poly}(\lambda)$ difference from $\mu(\operatorname{msg})$. Below, we omit writing "for infinitely many λ 's" when the context is clear. By Fact 6.14, there is some $\operatorname{poly}_1(\cdot)$ such that for any $\operatorname{msg} \in \operatorname{bad}$, at least $1/\operatorname{poly}_1(\lambda)$ fraction of $\operatorname{coin}_{\mathcal{P}}$ would cause the payment to be at least $1/\operatorname{poly}_1(\lambda)$ larger than $\mu(\operatorname{msg})$. We can consider the following platform-seller polynomial-time strategy: whenever a buyer with true value v_{λ} sends msg, the platform samples $\operatorname{coin}_{\mathcal{P}}$ for $\operatorname{poly}_1^2(\lambda)$ times, and chooses the best $\operatorname{coin}_{\mathcal{P}}$ that maximizes its revenue. Conditioned on $\operatorname{msg} \in \operatorname{bad}$, except with negligible probability, this strategy would result in at least $1/\operatorname{poly}_1(\lambda)$ expected payment over the honest strategy. Conditioned on $\operatorname{msg} \notin \operatorname{bad}$, the expected payment is at least as large as the honest case. Overall, the expected gain over the honest strategy is $1/\operatorname{poly}_2(\lambda)$. This violates computational psIC.

Claim 6.16. Suppose the world has only one buyer with an arbitrary value $b \in \mathbb{U}_{\lambda}$. Except over negligible probability over the choice of $\mathsf{msg} \xleftarrow{\$} \mathsf{msg}(b, *)$, $|\mathsf{butil}(\mathsf{msg}) - \mathsf{butil}(b)| \leq \mathsf{negl}(\lambda)$ where $\mathsf{butil}(\mathsf{msg})$ denotes the buyer's expected utility conditioned on msg , and $\mathsf{butil}(b)$ denotes the buyer's expected utility given value b.

Proof. For the sake of contradiction, suppose for infinitely many λ 's, there is some $b_{\lambda} \in \mathbb{U}_{\lambda}$, such that with at least $1/\operatorname{poly}(\lambda)$ probability over the choice of msg, $|\operatorname{butil}(\operatorname{msg}) - \operatorname{butil}(b_{\lambda})| \geq 1/\operatorname{poly}(\lambda)$. Below, we omit writing "for infinitely many λ 's" whenever the context is clear. By Fact 6.14, it must be that with at least $1/\operatorname{poly}_1(\lambda)$ probability over msg, $\operatorname{butil}(\operatorname{msg}) - \operatorname{butil}(b_{\lambda}) \geq 1/\operatorname{poly}_1(\lambda)$. Now a strategic buyer with value b_{λ} can simply choose a msg such that $\operatorname{butil}(\operatorname{msg}) - \operatorname{butil}(b_{\lambda}) \geq 1/\operatorname{poly}_1(\lambda)$, and such a message can be provided to the buyer as an advice string. This strategy is clearly polynomial time and lets the buyer gain $1/\operatorname{poly}_1(\lambda)$ amount over the honest strategy, which violates computational bIC.

Fix some λ , and $b_{\lambda} \in \mathbb{U}_{\lambda}$. The msgs from $\mathsf{msg}(b_{\lambda}, *)$ can be classified into two types, where $\mathsf{negl}'(\cdot)$ denotes a negligible function whose value is sufficiently large w.r.t. the negligible functions in Claim 6.15 for sufficiently large λ s.

- Type 1 msg: for at least negl'(λ) fraction of coin_P, the buyer does not get an item;
- Type 2 msg: the buyer gets an item with at least $1 \operatorname{negl}'(\lambda)$ over the choice of $\operatorname{coin}_{\mathcal{P}}$.

Lemma 6.17. Consider a world with only one buyer. It must be that for sufficiently large λ , for any $b_{\lambda} \in \mathbb{U}_{\lambda}$, except with negligible probability over the choice of $\mathsf{msg} \overset{\$}{\leftarrow} \mathsf{msg}(b_{\lambda}, *)$, if msg is of type 2, then the buyer's payment conditioned on msg is at most negligibly apart from some constant μ_{λ} that does not depend on b or msg .

Proof. By Claim 6.16, except with $\operatorname{negl}(\lambda)$ probability over the choice of $\operatorname{msg} \xleftarrow{\$} \operatorname{msg}(b_{\lambda}, *)$, if msg is of type 2, then $|\mu(\operatorname{msg}) - \mu(b_{\lambda})| \leq \operatorname{negl}(\lambda)$ where $\mu(\operatorname{msg})$ denotes the buyer's expected payment conditioned on msg , and $\mu(b_{\lambda})$ is a function that depends only on b_{λ} .

Suppose that for infinitely many λ 's, there exists some $b'_{\lambda} \in \mathbb{U}_{\lambda}$ such that type-2 messages happen with $1/\operatorname{poly}(\lambda)$ probability under honest execution. We claim that for any $b_{\lambda} \neq b'_{\lambda}$, either $\operatorname{msg} \overset{\$}{\leftarrow} \operatorname{msg}(b_{\lambda}, *)$ gives a type-2 message with negligible probability, or $\mu(b_{\lambda}) - \mu(b'_{\lambda}) \leq \operatorname{negl}(\lambda)$. Suppose this is not true, that is, for infinitely many λ 's, there is some $b_{\lambda} \neq b'_{\lambda}$ such that $\operatorname{msg} \overset{\$}{\leftarrow} \operatorname{msg}(b_{\lambda}, *)$ gives a type-2 message with $1/\operatorname{poly}(\lambda)$ probability, and $\mu(b_{\lambda}) - \mu(b'_{\lambda}) \geq 1/\operatorname{poly}(\lambda)$. In this case, a buyer with value b_{λ} can adopt the following polynomial-time strategy. Choose a type-2 message $\mathsf{msg}^* \in \mathsf{msg}(b'_{\lambda}, *)$ which minimizes its payment — msg^* can be provided as an advice string to the buyer. By Claim 6.16, if the buyer acts honestly, its expected utility is at most $b_{\lambda} - \mu(b_{\lambda}) + \mathsf{negl}(\lambda)$. Now under the above strategy, its expected utility is at least $b_{\lambda} - \mu(b'_{\lambda}) - \mathsf{negl}(\lambda)$, which is $1/\mathsf{poly}(\lambda) - 2\mathsf{negl}(\lambda)$ larger than acting honestly, which violates computational bIC.

The lemma now follows by letting $\mu_{\lambda} := \mu(b'_{\lambda})$.

Finishing the proof of Theorem 6.13. By Claim 6.15, and because $\operatorname{negl}'(\cdot)$ is sufficiently large, we have except with negligible probability over the choice of $\operatorname{msg} \overset{\$}{\leftarrow} \operatorname{msg}(b_{\lambda}, *)$, if msg is of type 1, then the buyer's expected payment is at most $\operatorname{negl}(\lambda)$. Combined with Lemma 6.17, we conclude that when the world has a single buyer, if its true value is less than $\mu_{\lambda} - \operatorname{negl}_1(\lambda)$, then its expected payment is at most $\operatorname{negl}_2(\lambda)$; else its expected payment cannot exceed $\mu_{\lambda} + \operatorname{negl}_3(\lambda)$. Henceforth, define $\mu_{\text{posted}}(v)$ to be an upper bound on the expected revenue earned from a buyer with value v:

$$\mu_{\text{posted}}(v) = \begin{cases} \mathsf{negl}_2(\lambda) & \text{if } v < \mu_\lambda - \mathsf{negl}_1(\lambda) \\ \mu_\lambda + \mathsf{negl}_3(\lambda) & \text{o.w.} \end{cases}$$

Now, suppose that for infinitely many λ 's, there is value vector $\mathbf{b} = (b_1, \ldots, b_n) \in \mathbb{U}^n_{\lambda}$ such that some buyer *i*'s expected payment is more than $\mu_{\text{posted}}(b_i)+1/\text{poly}(\lambda)$. We have the following scenario and a polynomial-time strategy that violate computational psIC. Suppose the world actually has only a single buyer with true value b_i . The platform-seller pretends that the world is actually \mathbf{b} where the victim buyer is the *i*-th coordinate. In other words, the strategy will simulate n-1 fake players with the true values \mathbf{b}_{-i} and run the protocol with the single victim buyer. This strategy has expected gain at least $1/\text{poly}(\lambda) - \text{negl}(\lambda)$. This violates computational psIC.

For the information-theoretic case, we can force all the negligible functions to be 0. In this case, all the $1/\text{poly}(\lambda)$ slacks would vanish to 0, and the statement follows directly.

6.4.2 The Bayesian Setting

We now extend the result to the Bayesian setting.

Corollary 6.18. Consider any 1-message auction Π_{λ} parametrized by λ , over the finite value domain $\mathbb{U}_{\lambda} \subset [0,1]$. Suppose Π_{λ} satisfies computational Bayesian bIC and computational Bayesian psIC w.r.t. a polynomial-time samplable distribution \mathcal{D}_{λ} over \mathbb{U}_{λ} . Then, for any polynomial function $n(\cdot)$, there exists a negligible function $\operatorname{negl}(\cdot)$, such that for any λ , there is a constant μ_{λ} , for any $n' \leq n(\lambda)$, except with $\operatorname{negl}(\lambda)$ probability over the choice of $\mathbf{v} \stackrel{\$}{\leftarrow} \mathcal{D}_{\lambda}^{n'}$, the auction's expected revenue on \mathbf{v} is at most $\operatorname{negl}(\lambda)$ more than the revenue of a posted price auction with reserve μ_{λ} on \mathbf{v} .

Further, if computational is replaced with "information theoretic" above, then, for any λ , there exists some constant μ_{λ} such that for any n, with probability 1 over the choice of $\mathbf{v} \stackrel{\$}{\leftarrow} \mathcal{D}_{\lambda}^{n}$, the auction's expected revenue on \mathbf{v} is no more than the revenue of a posted price auction with reserve μ_{λ} .

Proof. With Bayesian psIC, the following variant of Claim 6.15 holds.

Claim 6.19. In a world with only one buyer, for sufficiently large λ , for all but negligible fraction of buyer true value drawn from \mathcal{D}_{λ} , under honest execution, except with negligible probability over the choice of msg, it holds that except with negligible probability over the choice $\operatorname{coin}_{\mathcal{P}}$, $|\mu(\mathsf{msg}, \operatorname{coin}_{\mathcal{P}}) - \mu(\mathsf{msg})| \leq \operatorname{negl}(\lambda)$.

Proof. We first prove it for the computational setting. The proof is almost the same as that of Claim 6.15 except that "for any true value" is replaced with "for all but negligible fraction of the values sampled from \mathcal{D}_{λ} ", and the new proof relies on Bayesian psIC rather than ex post psIC. \Box

Claim 6.16 and Lemma 6.17 still hold because when the world consists of only one buyer, Bayesian bIC is the same as ex post bIC.

Now, the remainder of the proof works as follows. Similar to the proof in the ex post setting, combining Claim 6.19 and Lemma 6.17, we have that when the world consists of a single buyer, except with negligible probability over the choice of its true value $v \stackrel{\$}{\leftarrow} \mathcal{D}_{\lambda}$, the expected revenue is upper bounded by $\mu_{\text{posted}}(v)$.

Suppose that for some polynomial $n(\cdot)$, for infinitely many λ 's, there exist $1/\operatorname{poly}(\lambda)$ and some $n' \leq n(\lambda)$ such that over $1/\operatorname{poly}(\lambda)$ probability over the choice of $\mathbf{v} = (v_1, \ldots, v_{n'}) \stackrel{\$}{\leftarrow} \mathcal{D}_{\lambda}^{n'}$, the first buyer's expected payment exceeds $\mu_{\operatorname{posted}}(v_1)+1/\operatorname{poly}(\lambda)$. Then, a strategic platform-seller coalition can take the following polynomial-time strategy. Upon receiving the message msg_1 from the first buyer, it samples n'-1 buyers whose true values drawn at random from $\mathcal{D}_{\lambda}^{n'-1}$. It then samples all the random coins of the n'-1 fake buyers as well as $r_{\mathcal{P}}$ and computes the payment of the first buyer. It repeats the above sampling for a sufficiently polynomially many times and chooses a scenario that maximizes the first buyer's payment (including the honest case without the fake buyers), and pretends that the world is that particular scenario to the first buyer. This strategy allows the strategic platform-seller coalition to gain at least $1/\operatorname{poly}'(\lambda)$ amount in expectation.

Finally, for the information theoretic setting, the negligible functions and $1/\text{poly}(\lambda)$ terms all vanish to 0 in the above argument, so the claimed conclusion holds.

6.5 Information Theoretic bIC + $psIC \implies Small$ Revenue

In this section, we prove the following theorem.

Theorem 6.20 (Information theoretic bIC + psIC \implies small revenue). Let Π be a possibly multiround auction that satisfies information theoretic Bayesian bIC and ex post psIC. Then, Π is revenue-dominated by posted price.

In comparison with the proof in Section 6.4, the main challenge here is how to still prove that with only one buyer, the auction is revenue-dominated by posted price, now that the auction can be multi-round. We present our proof below.

Proof. (of Theorem 6.20.) We first analyze a world with only one buyer.

Single buyer setting. We will consider a world with only one buyer with true value v. Since Bayesian bIC is equal to expost bIC in a single-buyer setting, we will simply say bIC for short in this part of the proof. We use the following notations:

- $bmsg_{-1}$: the buyer's last message in the protocol,
- tr_{-1} : all messages exchanged between the buyer and the platform-seller coalition till right before the buyer is about to send its last message $bmsg_{-1}$,
- pcoin: additional private coin tosses made by the platform-seller coalition that did not contribute to tr_{-1} .

Fact 6.21. Let x and y be random variables sampled from an arbitrary joint distribution \mathcal{D} over $\mathbb{D}_x \times \mathbb{D}_y$. Let $f(x, y) : \mathbb{D}_x \times \mathbb{D}_y \to \mathbb{R}$ be some real-valued function, and let f(x) be the expectation of f(x, y) conditioned on x. If with non-zero probability, $f(x, y) \neq f(x)$, then the expectation of f(x, y) in the first process is greater (or smaller) than in the second process below:

- 1. Sample x from the marginal distribution over x induced by \mathcal{D} , choose y to maximize (or minimize) f(x, y);
- 2. Sample x and y at random from \mathcal{D} .

Proof. It suffices to show that with non-zero probability, $f(x) < \max_y f(x, y)$. Suppose this is not true, that is, with probability 1, $f(x) = \max_y f(x, y)$. Then, with probability 1, $f(x, y) \leq f(x)$. Because $f(x) = \int_y f(x, y) \cdot \mathsf{pdf}(y)$ where $\mathsf{pdf}(y)$ denotes the probability density function over y, it must be that with probability 1, f(x) = f(x, y) which contradicts our assumption.

Claim 6.22. Suppose that the world has only one buyer with an arbitrary value. Let $r(tr_{-1}, bmsg_{-1})$ denote the auction's expected revenue conditioned on $(tr_{-1}, bmsg_{-1})$. Then, with probability 1, the auction's revenue is $r(tr_{-1}, bmsg_{-1})$, i.e., the revenue does not depend on pcoin.

Proof. Suppose that the claim is not true, then a strategic platform-seller coalition can adopt the following strategy: upon observing $(tr_{-1}, bmsg_{-1})$, choose a pcoin that maximizes its revenue. Due to Fact 6.21, the platform-seller coalition can gain a positive amount with this strategy which violates ex post psIC.

Claim 6.22 implies that with probability 1, one of the following must be true:

- Case 1: $r(tr_{-1}, bmsg_{-1}) > 0$, the buyer gets an item and pays $r(tr_{-1}, bmsg_{-1})$;
- Case 2: $r(tr_{-1}, bmsg_{-1}) = 0$, and the buyer pays 0 if it gets an item.

Claim 6.23. Fix an arbitrary v of the lone buyer. With probability 1, if $(tr_{-1}, bmsg_{-1})$ belongs to Case 1, it must be that $r(tr_{-1}, bmsg_{-1}) = r(tr_{-1})$ where $r(tr_{-1}, bmsg_{-1})$ is the expected payment of the buyer conditioned on $(tr_{-1}, bmsg_{-1})$ and $r(tr_{-1})$ is the expected payment of the buyer conditioned on tr_{-1} and sampling some $bmsg_{-1}$ such that we land in Case 1.

Proof. We may assume that Case 1 happens with non-zero probability, since otherwise the claim trivially holds.

Suppose the claim is not true, then a strategic buyer can adopt the following strategy: it follows the protocol honestly, until it is about to send the last message $bmsg_{-1}$. At this moment, the buyer first samples $bmsg_{-1}$ honestly and checks if we are in Case 1. If so, it replaces the honestly sampled $bmsg_{-1}$ with a $bmsg_{-1}^*$ such that we are in Case 1 and $r(tr_{-1}, bmsg_{-1}^*)$ is minimized. Due to Fact 6.21, this strategy brings positive gain to the buyer which violates bIC.

Let $\mu(\mathsf{tr}_{-1}) := r(\mathsf{tr}_{-1}) \cdot \Pr[\text{Case 1}|\mathsf{tr}_{-1}]$ denote the expected revenue conditioned on tr_{-1} .

Claim 6.24. Fix an arbitrary v. Then, with probability 1 over the choice of tr_{-1} , the following holds: if $\mu(tr_{-1}) > 0$, then, conditioned on tr_{-1} , the buyer's expected utility is $v - r(tr_{-1})$ where $r(tr_{-1})$ is defined as in Claim 6.23.

Proof. Claim 6.23 shows that with probability 1, if we are in Case 1, the buyer's utility must be $v - r(tr_{-1})$. It suffices to show that with probability 1 over the choice of $(tr_{-1}, bmsg_{-1})$, if $\mu(tr_{-1}) > 0$ and we are in Case 2, then the buyer's expected utility conditioned on $(tr_{-1}, bmsg_{-1})$ must still be $v - r(tr_{-1})$. We may assume that with non-zero probability, $\mu(tr_{-1}) > 0$ and we are in Case 2, since otherwise the claim trivially holds. Suppose the above claim is not true, i.e., there is some non-zero probability such that $\mu(tr_{-1}) > 0$ and we are in Case 2, but the buyer's conditional expected utility is not $v - r(tr_{-1})$. Then, a strategic buyer can adopt the following strategy: follow the protocol honestly until it is about to send its last message $bmsg_{-1}$, and pick a $bmsg_{-1}^*$ that maximizes its expected utility. Similar to the proof of Fact 6.21, it is easy to show that this strategy can increase the buyer's expected utility which violates bIC.

Claim 6.25. Fix an arbitrary v. With probability 1 over the choice of tr_{-1} , the conditional expected revenue $\mu(tr_{-1})$ depends only on the buyer's coin tosses that contributed to tr_{-1} .

Proof. Consider that the lone buyer and the platform-seller coalition engage in a protocol to generate tr_{-1} . Suppose right before the platform-seller coalition sends the first message $pmsg_1$ to the buyer, the transcript so far is denoted tr_1 . Let $\mu(tr_1)$ be the conditional expectation of μ given tr_1 . Let $\mu(tr_1, pmsg_1)$ be the conditional expectation of μ given $(tr_1, pmsg_1)$. We claim that with probability 1 over the choice of $(tr_1, pmsg_1)$, it must be that $\mu(tr_1) = \mu(tr_1, pmsg_1)$. Suppose this is not true. Then, a strategic platform-seller coalition can adopt the following strategy: when it is about to send its first message, sample a $pmsg_1$ that maximizes $\mu(tr_1, pmsg_1)$. Due to Fact 6.21, it can gain positively based on this strategy which violates ex post psIC.

Repeating this argument for every message sent from the platform-seller coalition to the buyer, we conclude the following. Suppose we sample tr_{-1} at random, and let **bcoin** be the buyer's coins that contributed to the sampling of tr_{-1} , we can compute $\mu(tr_{-1})$ from **bcoin** as follows: sample the platform-seller coalition's coin tosses at random, and combined with **bcoin**, we get some tr'_{-1} . With probability 1 over the choice of tr_{-1} and tr'_{-1} , it holds that $\mu(tr'_{-1}) = \mu(tr_{-1})$.

Let bcoin be the buyer's coin tosses in generating tr_{-1} , and suppose bcoin leads to $\mu > 0$. In this case, let $\bar{r}(bcoin)$ be the expected payment of the buyer conditioned on bcoin and landing in Case 1. Let \bar{r} denote the expectation of $\bar{r}(bcoin)$ over the choice of bcoin subject to $\mu > 0$.

Claim 6.26. Fix an arbitrary v. With probability 1 over the choice of bcoin, if bcoin leads to a positive μ , then $\bar{r}(bcoin) = \bar{r}$.

Proof. Henceforth, we may assume that given v, with non-zero probability over the choice of bcoin, it holds that $\mu > 0$. Otherwise, the claim holds trivially.

Suppose the claim does not hold, then, a strategic buyer can adopt the following strategy: if it happens to sample bcoin such that $\mu > 0$, then replace bcoin with a choice bcoin^{*} that minimizes $\bar{r}(bcoin^*)$. Due to Claim 6.24 and Fact 6.21 this allows the buyer to gain which violates bIC.

Given the above claim, given any v such that the expected revenue is positive, we can define the buyer's pay in Case 1 as $\bar{r} = \bar{r}(v)$, which is a function only of v. We next argue that in fact, \bar{r} cannot even depend on v.

Claim 6.27. Given v_1 and v_2 such that the expected revenue is non-zero under v_1 and v_2 , $\bar{r}(v_1) = \bar{r}(v_2)$.

Proof. Suppose $\bar{r}(v_1) < \bar{r}(v_2)$. Then, a buyer with true value v_2 can adopt the following strategy. It first samples bcoin₂ honestly taking on the true value v_2 . If bcoin₂ leads to $\mu > 0$, it will instead pretend that its true value is v_1 , and sample some corresponding bcoin₁ which leads to $\mu > 0$, and

its pay in Case 1 would be $\bar{r}(v_1)$. It will then participate in the protocol using v_1 and bcoin₁ as input, until the buyer is about to send its last message bmsg_{-1} . At this moment, the buyer samples bmsg_{-1} such that we land in Case 1.

Due to Claim 6.24, conditioned on bcoin₂ leading to $\mu > 0$, if the buyer continued to behave honestly, its expected utility would be $v_2 - \bar{r}(v_2)$. If it adopted the above strategy, its expected utility would be $v_2 - \bar{r}(v_1)$ which is strictly better than the honest case. This violates bIC.

Given the above claim, we can let $\bar{r} = \bar{r}(v_1) = \bar{r}(v_2)$ be a universal constant. By individual rationality, if the buyer's true value $v < \bar{r}$, then $\mu > 0$ happens with probability 0. In other words, if $v < \bar{r}$, the auction's revenue must be 0. Otherwise, its expected revenue is at most \bar{r} . Therefore, when the world has a single buyer with value v, the auction's revenue is dominated by

$$\mu_{\text{posted}}(v) = \begin{cases} 0 & \text{if } v < \bar{r} \\ \bar{r} & \text{o.w.} \end{cases}$$

Completing the proof of Theorem 6.20. Suppose that there exists some value vector $\mathbf{v} = (v_1, \ldots, v_n)$ and some $i \in [n]$ such that *i*'s expected payment is more than $\mu_{\text{posted}}(v_i)$. Then, when the world has a single buyer v_i , the platform-seller coalition can simulate n - 1 fake buyers with values \mathbf{v}_{-i} . This strategy allows the coalition to gain in expectation.

7 Utility-Dominated Emulation

We will eventually show that by considering computationally bounded agents, we can obtain a constant-round auction that satisfies computational bIC, pIC, 1-pbIC, Bayesian sIC and Bayesian psIC simultaneously. To this end, we introduce a new design paradigm. We first define an ideal world in which the seller, the platform, and the buyers all interact with a trusted ideal functionality $\mathcal{F}_{auction}$ which helps to implement the auction. If in this ideal world, we can design a direct-revelation auction that satisfies the desired incentive compatibility properties in an *information theoretic* sense, then we can compile it to a real-world protocol that replaces the ideal functionality with cryptography, and the real-world protocol satisfies the same set of incentive compatibility properties *against computationally bounded players*.

As mentioned in Section 1, the most straightforward approach is to compile the ideal auction using generic multi-party computation (MPC). However, all known generic approaches (including threshold FHE) incur at least quadratic overhead, partly because in an auction, every buyer obtains a different outcome. Instead, we formulate a weaker notion of simulation called *utility-dominated emulation*. With this new notion, we show an efficient compiler where every buyer and seller has only $\tilde{O}_{\lambda}(1)$ cost and the platform has $\tilde{O}_{\lambda}(n)$ cost where n is the number of buyers.

7.1 Model for the Ideal Auction

Because we will later compile the ideal-world auction into a real-world cryptographic protocol, we will also parametrize the ideal auction with λ , even though it provides information theoretic guarantees. In natural designs of the ideal auction, the rules of the auction are uniform algorithms that work for all λ .

We will consider a direct revelation mechanism in the ideal world. Since the ideal world has a trusted functionality $\mathcal{F}_{auction}$, it is not hard to extend the direct revelation principle to cover all of the incentive compatibility properties we care about. Thus, considering a direct revelation mechanism in the ideal world is without loss of generality. Therefore, We can specify an ideal-world auction for

selling k identical items by specifying the following possibly randomized algorithms. Henceforth, we assume that the value domain is \mathbb{U}_{λ} given λ .

- Allocation rule $\mathbf{x}(1^{\lambda}, \mathbf{b})$: takes as input 1^{λ} and a bid vector $\mathbf{b} = (b_1, \ldots, b_n) \in \mathbb{U}^n_{\lambda}$, and outputs a vector (x_1, \ldots, x_n) where $x_i \in \{0, 1\}$ specifies whether the buyer *i* is allocated an item.
- Payment rule $\mathbf{p}(1^{\lambda}, \mathbf{b})$: takes as input 1^{λ} and a bid vector $\mathbf{b} = (b_1, \ldots, b_n) \in \mathbb{U}^n_{\lambda}$, and outputs (p_1, \ldots, p_n) where $p_i \ge 0$ specifies buyer *i*'s payment.
- Seller revenue rule $\mu_{\mathcal{S}}(1^{\lambda}, \mathbf{b})$: takes as input 1^{λ} and a bid vector $\mathbf{b} = (b_1, \ldots, b_n) \in \mathbb{U}^n_{\lambda}$, and outputs $\mu_{\mathcal{S}} \geq 0$ specifying the revenue that the seller gets.

For compiling the ideal auction to a real-world protocol, we additionally require that the above rules can be computed in polynomial time.

Ideal auction. Given these rules $(\mathbf{x}, \mathbf{p}, \mu_{\mathcal{S}})$, we define the following ideal auction protocol denoted $\Pi_{\lambda}^{\text{ideal}}[\mathbf{x}, \mathbf{p}, \mu_{\mathcal{S}}]$. Henceforth let $\mathcal{F}_{\text{auction}}$ denote a trusted ideal functionality.

- 1. Every buyer sends zero to multiple bids to $\mathcal{F}_{auction}$.
- 2. $\mathcal{F}_{\text{auction}}$ sends the number of bids received to the platform.
- 3. The platform and the seller may each send zero to multiple bids to $\mathcal{F}_{auction}$.
- 4. $\mathcal{F}_{auction}$ computes the outcome for the seller, the platform, and all buyer identities based on the rules $(\mathbf{x}, \mathbf{p}, \mu_{\mathcal{S}})$, using all bids it has received. If the auction's rules are randomized, $\mathcal{F}_{auction}$ tosses the random coins. $\mathcal{F}_{auction}$ informs the platform and seller of their respective outcomes. Further, for each bid, $\mathcal{F}_{auction}$ informs its outcome to the player who sent the bid earlier.
- 5. The platform sends $ok \text{ or } \perp \text{ to } \mathcal{F}_{auction}$. The auction is considered successful if the platform sends ok; else we treat the auction as a failure and every one's utility is defined to be 0.

Honest behavior. The honest behavior is defined as follows. Every honest buyer sends its true value in Step 1. An honest platform and an honest seller do not send any bid in Step 3. An honest platform always sends ok in Step 5.

Strategic behavior. All possible strategic behavior is already made evident through the description of the ideal auction itself.

7.2 Definition: Utility-Dominated Emulation

Let $\Pi_{\lambda}^{\text{ideal}}$ be an ideal auction over some a family of finite value domains $\mathbb{U}_{\lambda} \subset \mathbb{R}$, and suppose $\mathcal{F}_{\text{auction}}$ satisfies information theoretic incentive compatibility guarantees. We want to get rid of the ideal functionality $\mathcal{F}_{\text{auction}}$ in the ideal auction and instantiate it with cryptography, resulting in a real-world cryptographic protocol denoted $\Pi_{\lambda}^{\text{real}}$. We want to argue that the information theoretic incentive compatibility guarantees in the ideal world translate to computational counterparts of the same properties in the real world. Since we only need to argue game-theoretic properties, it turns out that we do not require a notion as strong as full simulation [Can00, Can01]. In this section, define a weaker notion of simulation called utility-dominated emulation which suffices for proving the game theoretic properties we care about. With this new relaxed notion, we can construct cryptographic protocols that are asymptotically faster than known generic MPC from standard assumptions.

Definition 7.1 (Utility-dominated emulation). Let $\Pi_{\lambda}^{\text{real}}$ and $\Pi_{\lambda}^{\text{ideal}}$ be auctions over the finite value domain $\mathbb{U}_{\lambda} \subset \mathbb{R}$. We say that $\Pi_{\lambda}^{\text{real}}$ is a utility-dominated emulation of $\Pi_{\lambda}^{\text{ideal}}$ w.r.t. the strategic individual or coalition \mathcal{C} , iff the following properties hold:

- 1. Utility equivalence under honest execution: for any $\lambda \in \mathbb{N}$, for any $\mathbf{v} \in \mathbb{U}_{\lambda}^{*}$, and for any player p (buyer, seller or platform), $\mathsf{UReal}_{p}(1^{\lambda}, \mathbf{v})$ and $\mathsf{UIdeal}_{p}(1^{\lambda}, \mathbf{v})$ are identically distributed, where $\mathsf{UReal}_{p}(1^{\lambda}, \mathbf{v})$ is the random variable representing the utility of player p under an honest execution of $\Pi_{\lambda}^{\text{real}}$ parametrized with security parameter 1^{λ} , and with buyers taking on the value vector \mathbf{v} , and $\mathsf{UIdeal}_{p}(1^{\lambda}, \mathbf{v})$ is similarly defined but for $\Pi_{\lambda}^{\text{ideal}}$.
- 2. Utility dominance for strategic players: for any polynomial $n(\cdot)$, for any real-world PPT strategy S adopted by C, there exist a negligible function $\operatorname{negl}(\cdot)$, and an ideal-world strategy $S'(\cdot)$ which may additionally depend on the coalition's true value vector, such that for any λ , any $n_H \leq n(\lambda)$, any true value vector $\mathbf{v}_{-C} \in \mathbb{U}_{\lambda}^{n_H}$, any true value vector $\mathbf{v}_{\mathcal{C}} \in \mathbb{U}_{\lambda}^{n_C}$ where n_C denotes the number of buyers in C, it holds that

$$\mathbf{E}\left[\mathsf{UReal}_{\mathcal{C}}^{S}(1^{\lambda},\mathbf{v}_{\mathcal{C}},\mathbf{v}_{-\mathcal{C}})\right] \leq \mathbf{E}\left[\mathsf{UIdeal}_{\mathcal{C}}^{S'(\mathbf{v}_{\mathcal{C}})}(1^{\lambda},\mathbf{v}_{\mathcal{C}},\mathbf{v}_{-\mathcal{C}})\right] + \mathsf{negl}(\lambda)$$

where

- UReal^S_C $(1^{\lambda}, \mathbf{v}_{\mathcal{C}}, \mathbf{v}_{-\mathcal{C}})$ is the random variable representing the utility of the coalition \mathcal{C} when we execute $\Pi_{\lambda}^{\text{real}}$ with security parameter 1^{λ} , honest buyers taking the true values $\mathbf{v}_{-\mathcal{C}}$, and the coalition \mathcal{C} taking the true values $\mathbf{v}_{\mathcal{C}}$ and the strategy S; and
- $\mathsf{Uldeal}_{\mathcal{C}}^{S'}(1^{\lambda}, \mathbf{v}_{\mathcal{C}}, \mathbf{v}_{-\mathcal{C}})$ is similarly defined but for the ideal protocol Π_{λ}^{ideal} .

Note that in general, the utility-dominated emulation notion is well-defined for any (possibly multi-round) ideal-world auction, even though in this paper, we shall focus on ideal auctions in the model specified in Section 7.1.

Theorem 7.2 (Ideal-real design paradigm). Let $\Pi_{\lambda}^{\text{ideal}}$ and $\Pi_{\lambda}^{\text{real}}$ be auctions parametrized by λ over a family of finite value domains \mathbb{U}_{λ} . Suppose that $\Pi_{\lambda}^{\text{real}}$ is a utility-dominated emulation of $\Pi_{\lambda}^{\text{ideal}}$ w.r.t. the strategic individual or coalition C, and $\Pi_{\lambda}^{\text{ideal}}$ satisfies information-theoretic (Bayesian) incentive compatibility w.r.t. C. Then, $\Pi_{\lambda}^{\text{real}}$ satisfies computational (Bayesian) incentive compatibility w.r.t. C.

Proof. Suppose for the sake of contradiction that $\Pi_{\lambda}^{\text{real}}$ does not satisfy computational incentive compatibility w.r.t. \mathcal{C} . This means that there is some polynomial $n(\cdot)$, some real-world PPT strategy S, some polynomial function $\text{poly}(\cdot)$ such that for infinitely many λ 's, there exist $n_H \leq n(\lambda)$, some value vector $\mathbf{v}_{-\mathcal{C}} \in \mathbb{U}_{\lambda}^{n_H}$ corresponding to honest buyers, some value vector $\mathbf{v}_{\mathcal{C}} \in \mathbb{U}_{\lambda}^{n_C}$ corresponding to the buyers in \mathcal{C} , such that

$$\mathbf{E}\left[\mathsf{util}_{\mathcal{C}}^{S}(1^{\lambda}, \mathbf{v}_{-\mathcal{C}}, \mathbf{v}_{\mathcal{C}})\right] \geq \mathbf{E}\left[\mathsf{util}_{\mathcal{C}}^{H}(1^{\lambda}, \mathbf{v}_{-\mathcal{C}}, \mathbf{v}_{\mathcal{C}})\right] + 1/\mathsf{poly}(\lambda)$$

By the definition of utility-dominated emulation, and the fact that real and ideal executions are utility equivalent in an all-honest execution, we conclude that there is some strategy S' that outperforms the honest strategy in the ideal world by $1/\text{poly}(\lambda) - \text{negl}(\lambda)$. This violates the fact that the ideal-world auction satisfies incentive compatibility w.r.t. C. The proof for the Bayesian notion is similar, except that \mathbf{v}_{-C} will be sampled at random from $\mathbb{U}_{\lambda}^{n_{H}}$ instead.

Remark 7.3 (An alternative definition of computational incentive compatibility). In fact, Theorem 7.2 also gives an alternative way of defining computational incentive compatibility, by requiring it to be a utility-dominated emulation of some ideal auction that is incentive compatible in the information theoretic sense. This alternative definition implies our earlier formulations in Definition 4.1 and Definition 4.2.

8 Compiling an Ideal Auction to a Real-World Protocol

Given an ideal auction where we need a trusted entity $\mathcal{F}_{auction}$ to enforce the auction's rules, we can compile it into real-world cryptographic protocol that does not rely on $\mathcal{F}_{auction}$. Ideally, we want the protocol to satisfy an additional robustness property defined below.

Definition 8.1 (Robustness). Given an auction protocol parameterized by λ , we say that it is *robust*, iff the following holds: as long as the platform is honest and not part of the coalition C, then for any non-uniform PPT strategy adopted by C, there exists a negligible function $negl(\cdot)$ such that for every $\lambda \in \mathbb{N}$, with at least $1 - negl(\lambda)$ probability, all honest players accept the auction.

Although our robustness notion requires the platform to be honest, it should not be viewed as a limitation in our model. Because our platform-assisted model assumes that we want to have a platform to facilitate the auction by providing services such as rendezvous, product search, recommender systems, and payment processing. In other words, our model already assumes that the platform is needed for availability. Moreover, because of our incentive compatibility guarantees, it is in the platform's best interest to act honestly and provide availability to users.

8.1 Compiler

Cryptographic building blocks. We will make use of the following cryptographic building blocks.

- Vector commitment VC. We rely on a vector commitment scheme denoted VC = (Gen, Digest, Open, Vf). We can instaniate the VC scheme with a Merkle tree [Mer89].
- *Erasure code* RS. We rely on an erasure code that can correctly reconstruct as long as at least 2/3 fraction of the locations in the codeword are correct. We can employ a Reed-Solomon code denoted RS = (Encode, Recons).
- *Publicly verifiable non-malleable timed commitments* NITC. We rely on a publicly verifiable non-malleable timed commitments denoted (Gen, Com, ComVf, DecVf, FDec, FDecVf). We can instantiate the scheme with Chvojka and Jager [CJ23].
- Succinct argument of knowledge AoK. We additionally make use of a succinct argument of knowledge henceforth denoted AoK = (Gen, P, V). We will instantiate the AoK scheme using Kilian [Kil92] using a quasilinear Probabilistic Checkable Proof (PCP) [BSS08].

The formal definitions of these cryptographic building blocks are presented in Appendix A.

NP language. We specify the following NP language which will be used in our protocol. Since the terms NITC.crs and VC.crs will be clear from the context, we do not explicitly carry them in the statement below. A statement (digest, digest', $n, r_{\mathcal{P}}$) is in the language \mathcal{L}_{λ} , iff there exists a witness of the form $\left(\mathsf{code}, \mathcal{I}, \{c_j, \pi_j^{\mathrm{com}}, \pi_j^*, v_j, r_j, \mathsf{out}_j\}_{j \in \mathcal{I}}\right)$ such that the following claims hold:

- all identities in \mathcal{I} are distinct, and $|\mathcal{I}| = n$;
- code = RS.Encode(c) where $\mathbf{c} := \{j, c_j, \pi_j^{\text{com}}\}_{j \in \mathcal{I}};$
- digest is the first output of VC.Digest(VC.crs, c);
- digest' is the first output of VC.Digest(VC.crs, $\{i, c_i, \pi_i^{\text{com}}, \text{out}_i\}_{i \in \mathcal{I}}$);
- $\forall i \in \mathcal{I}$, NITC.ComVf(NITC.crs, $c_i, \pi_i^{\text{com}}) = 1$;
- $\forall i \in \mathcal{I}$, either NITC.DecVf(NITC.crs, $c_i, i \| v_i \| r_i, \pi_i^*$) = 1 or NITC.FDecVf(NITC.crs, $c_i, i \| v_i \| r_i, \pi_i^*$) = 1.
- $\{\operatorname{out}_i\}_{i\in\mathcal{I}}$ are the correctly computed outcomes by applying the rules $(\mathbf{x}, \mathbf{p}, \mu_S)$ over randomness $r = \left(\bigoplus_{j\in\mathcal{I}} r_j\right) \oplus r_P$ and input $\mathbf{v} = \{v_j\}_{j\in\mathcal{I}\setminus\{0\}}$.

Compiler. We assume that every buyer has a distinct identity *i*. The strategic players can choose fake identities and pretend to be multiple buyers. Given an ideal auction $\Pi_{\lambda}^{\text{ideal}}$ over the value domain $\mathbb{U}_{\lambda} \subset [0, 1]$ specified by the rules $(\mathbf{x}, \mathbf{p}, \mu_{\mathcal{S}})$ which are PPT algorithms, we describe a new compiler that compiles it to a robust cryptographic protocol.

Robust cryptographic auction $\Pi_{\lambda}^{\text{real}}(\mathbf{x},\mathbf{p},\mu_{\mathcal{S}})$

Inputs: All players have the security parameter 1^{λ} . Every buyer with identity *i* has a true value $v_i \in \mathbb{U}_{\lambda}$.

Setup: Let NITC.crs $\stackrel{\$}{\leftarrow}$ NITC.Gen $(1^{\lambda}, T)$, VC.crs $\stackrel{\$}{\leftarrow}$ VC.Gen (1^{λ}) , and AoK.crs \leftarrow AoK.Gen (1^{λ}) . Publish NITC.crs, VC.crs, and AoK.crs.

Main auction protocol:

(a) **Time renamed to** 0 (protocol start): Each buyer *i* samples $r_i \leftarrow \{0, 1\}^{\lambda}$, and $\operatorname{coin}_i \leftarrow \{0, 1\}^{\lambda}$, and (1, 2, 1\}^{\lambda}, and (1, 2, 1)^{\lambda}, and (1, 2,

The seller samples $r_0 \leftarrow \{0, 1\}^{\lambda}$ at random, computes $(c_0, \pi_0^{\text{com}}, \pi_0^{\text{dec}}) \leftarrow \mathsf{NITC.Com}(\mathsf{NITC.crs}, 0 \| \perp \| r_0; \operatorname{coin}_0)$, and sends $(0, c_0, \pi_0^{\text{com}})$ to the platform. Without loss of generality, we may assume that the seller has the identity 0 that is distinct from all buyers.

- (b) **Time** $T_1(\lambda)$: The platform ignores all tuples $(i, c_i, \pi_i^{\text{com}})$ where NITC.ComVf(NITC.crs, $c_i, \pi_i^{\text{com}}) = 0$, and ignores any tuple of the form (0, ..., ..) that is not sent by the seller. Among the remaining tuples, if there are multiple tuples containing the same identity i, the platform randomly keeps one of the tuples and ignores the rest. For each identity ignored, the platform notifies the corresponding buyer, and the corresponding buyer rejects the auction.
- (c) Let $\mathbf{c} = \{(i, c_i, \pi_i)\}_{i \in \mathcal{I}}$ be the remaining tuples where \mathcal{I} denotes the identities they are associated with. The platform samples randomness $r_{\mathcal{P}}$, and computes (digest, _) = VC.Digest(VC.crs, code) where code := RS.Encode(c). It sends $n = |\mathcal{I}|$, digest and $r_{\mathcal{P}}$ to all buyers and the seller.

- (d) **Time** $T_2(\lambda)$: If a buyer or seller still has not received a tuple $(n, \text{digest}, r_{\mathcal{P}})$ from the platform, reject the auction. Else, each buyer and the seller challenges the platform as follows. Sample $\kappa = \omega(\log \lambda)$ random indices $Q \subset [n']$, where n' is the length of an RS codeword for an input of length n, and send Q to the platform. The platform responds with the opening of the indices Q along with a membership proof w.r.t. digest.
- (e) **Time** $T_3(\lambda)$: If a buyer or the seller still has not received from the platform a response for its challenge, or if the platform's response fails to verify, reject the auction.

If a buyer $i \in \mathcal{I} \setminus \{0\}$ did not abort in the above, then send the opening $v_i || r_i$ along with the opening proof π_i^{dec} to the platform. If the seller did not abort, then send $\perp || r_0$ along with the opening proof π_0^{dec} to the platform.

- (f) **Time** $T_4(\lambda)$: Let $(v_i || r_i, \pi_i^{\text{dec}})$ denote the opening received by the platform from identity *i*. If there exists some $i \in \mathcal{I}$, such that NITC.DecVf(NITC.crs, $c_i, i || v_i || r_i, \pi_i^{\text{dec}}) = 0$ or *i* did not respond with an opening, then the platform performs forced decryption $(i || v_i || r_i, \pi_i^{\text{fdec}}) \leftarrow \text{NITC.FDec}(\text{NITC.crs}, c_i, \pi_i^{\text{com}})$. If the forced decryption does not produce a valid decryption proof, the platform posts \perp to the blockchain, and all players reject the auction.
- (g) The platform runs the auction's rules $(\mathbf{x}, \mathbf{p}, \mu_{\mathcal{S}})$ using the bid vector $\{v_i\}_{i \in \mathcal{I} \setminus \{0\}}$ and the randomness $r = \left(\bigoplus_{i \in \mathcal{I}} r_i\right) \oplus r_{\mathcal{P}}$ as input. Let out_i be *i*'s private outcome for $i \in \mathcal{I}$. The platform then computes $(\mathsf{digest}', -) = \mathsf{VC}.\mathsf{Digest}(\mathsf{VC}.\mathsf{crs}, \{i, c_i, \pi_i^{\mathrm{com}}, \mathsf{out}_i\}_{i \in \mathcal{I}})$ and sends digest' to all buyers and the seller.
- (h) Each buyer $i \in \mathcal{I} \setminus \{0\}$ and the seller invokes a separate AoK instance with the platform using AoK.crs, where the platform proves that (digest, digest', $n, r_{\mathcal{P}}$) is in the language \mathcal{L}_{λ} using the witness

$$(\mathsf{code}, \mathcal{I}, \{c_j, \pi_j^{\mathrm{com}}, \pi_j^*, v_j, r_j, \mathsf{out}_j\}_{j \in \mathcal{I}})$$

where $\pi_j^* = \pi_j^{\text{dec}}$ if j provided a valid opening, else $\pi_j^* = \pi_j^{\text{fdec}}$. If for any buyer or the seller, the AoK proof fails to verify, it rejects the auction.

- (i) For each $i \in \mathcal{I}$ that responded with a valid opening earlier, the platform sends to i its private outcome out_i , as well as a membership proof for the tuple $(i, c_i, \pi_i^{\mathrm{com}}, \mathsf{out}_i)$ w.r.t. digest' where $(i, c_i, \pi_i^{\mathrm{com}})$ is the message i sent the platform in Step (a). The player i rejects if the verification fails.
- (j) The platform posts (digest, digest', $n, r_{\mathcal{P}}$) to the blockchain. Everyone reads the blockchain and verifies that they have received the same values from the platform earlier in the protocol. Reject if the verification fails. Otherwise, accept the private outcome they have received.

Intuition. Informally, the protocol works as follows — for simplicity, we describe the intuition for the case when the auction's rules $(\mathbf{x}, \mathbf{p}, \mu_S)$ are deterministic. In Step (a), all buyers commit to their bids using a non-malleable timed commitment scheme. In Step (c), the platform computes a Reed-Solomon encoding of all buyers' commitments, and commits to the resulting code using a hash digest (denoted digest). In Steps (d) and (e), all buyers and the seller challenge the platform to open a small number of locations committed under digest and verify that the openings are correct.

Then, all buyers open their commitments to the platform. In Step (f), if any buyer fails to open, the platform force-decrypts its timed commitment. In Step (g), the platform computes all buyers' and the seller's private outcomes, and commits to a hash digest (denoted digest') of all original timed commitments as well as every one's private outcome. In Step (h), the platform invokes a succinct argument-of-knowledge to prove every one that 1) digest and digest' encode the same vector of committed bids denoted \mathbf{c} ; 2) the vector \mathbf{c} does not contain multiple commitments from the same identity; and 3) the outcomes committed under digest' are computed correctly based on valid openings of \mathbf{c} . In Step (i), the platform opens to each buyer and the seller its private outcome. To each buyer, it provides a succinct proof that its timed commitment as well as private outcome belong to the set represented by digest'. To the seller, it provides a succinct proof that its private outcome belongs to the set represented by digest'. Finally, in Step (j), the platform posts digest and digest' to the blockchain and all other players only accept if they agree with the digest and digest' they have received from the platform earlier.

It may be helpful to think of Steps (c), (d) and (e) as a proof-of-retrievability (PoR) protocol [BJO09, SW13]. If the platform can respond correctly to the PoR challenges, it is possible to rewind the platform and extract an underlying witness, i.e., a vector of timed commitments. Jumping ahead, later in our proof, we will use the extracted witness to construct a reduction that breaks the IND-CCA security of the timed commitment scheme. The depth of the extractor will contribute to the depth of the reduction — fortunately, the PoR's extractor indeed has small depth. Because we have a full AoK later in Step (h), we do not need the platform to provide a proximity proof that the committed code is close to a valid Reed-Solomon codeword in the PoR stage.

Parameter choices. We will choose the NITC scheme's parameters such that it satisfies $(W(\cdot), \epsilon)$ -IND-CCA security for any super-polynomial function $W(\cdot)$. Further, its difficulty parameter $T(\lambda)$ and the gap parameter ϵ should satisfy the following constraint:

$$T^{\epsilon}(\lambda) > T_2(\lambda) + T_3(\lambda) + D_{\text{auction}}(\lambda^{\alpha(\lambda)}) + \operatorname{poly}\log(\lambda)$$

where $\alpha(\lambda)$ is any super-constant function, and $D_{\text{auction}}(\lambda^{\alpha(\lambda)})$ is an upper bound on the depth for executing the auction's rules and computing the coalition's utility for any input whose length is an unknown polynomial in λ — as long as is $D_{\text{auction}}(L) = L^{o(1)}$, $D_{\text{auction}}(\lambda^{\alpha(\lambda)})$ is polynomially bounded for a suitably small super-constant function $\alpha(\lambda)$.

The above specifies the parameter constraints needed for proving the incentive compatibility guarantees. Besides the above, we should appropriately set T_1 such that buyers have enough time to compute and submit their commitments (Step (a)); set $T_2 - T_1$ such that the platform has enough time to compute and send the tuple $(n, \text{digest}, r_{\mathcal{P}})$ to every one (Step (c)); set $T_3 - T_2$ such that the platform has enough time to compute and send a response to the PoR opening challenge (Steps (d) and (e)), and set $T_4 - T_3$ such that buyers have enough time to send their openings (Step (e)).

Asymptotical efficiency. It is not hard to see that the seller and each buyer's computational and bandwidth cost are bounded by $\tilde{O}_{\lambda}(1)$ where $\tilde{O}_{\lambda}(\cdot)$ hides $poly(\lambda, \log n)$ factors where n is the total number of buyers. The platform's bandwidth is upper bounded by the sum of the bandwidth costs of all buyers and the seller, that is, $\tilde{O}_{\lambda}(n)$. The platform's computation is dominated by Step (h) where the platform needs to run an AoK instance to convince each buyer and the seller. Notice that the platform is proving the same statement to all other players. So if we use Kilian to instantiate the AoK, the platform only needs to compute the PCP proof once as well as the vector commitment of the PCP proof. Then, it can answer each verifier's challenge with the same PCP proof and vector commitment. Therefore, if we use a quasilinear PCP [BSS08], the platform's total computation is bounded by $\widetilde{O}_{\lambda}(n)$. Last but not the last, the platform posts only $\widetilde{O}_{\lambda}(1)$ bits to the broadcast channel at the end of the auction.

More practical variants. In practice, we can use a Succinct Non-interactive ARgument of Knowledge (SNARK) [GGPR13, Gro16, BBHR19, XZZ⁺19, ZXZS20, CHM⁺20, GWC19, CBBZ23, BSCR⁺19, Set20] to replace the AoK. This would result in more a practical instantiation but at the price of worsening the cryptographic assumptions needed, since known SNARK constructions typically rely on either the random oracle model or knowledge-type assumptions. Additionally, for the PoR challenge and response steps (Steps (d) and (e)), we can use the Fiat-Shamir paradigm to remove the interaction. Specifically, after computing digest, the server can use H(digest, i) to compute the random challenges for player i where $H(\cdot)$ denotes a random oracle.

Suppose we instantiate the protocol with a SNARK proof in lieu of the AoK, and using Fiat-Shamir to make the PoR challenge and response steps non-interactive, then, the only remaining message for the seller to send is the commitment of its random coins $coin_0$ in Step (a). We can remove even this message from the protocol, and get a version that supports an *offline server*. Specifically, the offline seller can query the platform a-posteriori to get its private outcome along with the corresponding SNARK and VC proofs (i.e., Steps (h) and (i)), and it can check that the outcome is correct against the messages that have been posted to the blockchain. It is not hard to mechanically go through our proofs and verify that this offline-server variant is also a utilitydominated emulation of the ideal auction, as long as one of the following is true: 1) the auction's rules ($\mathbf{x}, \mathbf{p}, \mu_S$) are deterministic, 2) there is at least one honest buyer, or 3) the platform is honest.

Finally, if the SNARK is instantiated with modern schemes [GWC19,CBBZ23,XZZ⁺19,ZXZS20, BSCR⁺19, Set20, CHM⁺20] that combine polynomial commitments and interactive oracle proofs (IOPs), it is typically desirable to instantiate the vector commitment using the polynomial commitment scheme native to the SNARK construction.

Theorem 8.2 (Efficient cryptographic auction). Suppose that the NITC scheme is computationally binding against quasi-polynomially sized adversaries, IND-CCA secure against quasi-polynomially sized and T^{ϵ} -depth adversaries, and satisfies soundness of forced decryption (for polynomially-sized adversaries). Additionally, suppose that VC satisfies collision resistance against quasi-polynomially sized adversaries, and AoK satisfies adaptive knowledge soundness for polynomially-sized adversaries. Let C be either a strategic individual buyer, the strategic seller alone, the strategic platform alone, the strategic platform-seller coalition, or a coalition between the platform and one or more buyers. Suppose also that the auction's rules $(\mathbf{x}, \mathbf{p}, \mu_S)$ can be computed in $\ell^{o(1)}$ depth on any input of length ℓ .

Then, the real-world auction $\pi_{\lambda}^{\text{real}}$ satisfies robustness, and is a utility dominated emulation of the ideal auction $\pi_{\lambda}^{\text{ideal}}$ w.r.t. C.

We provide the robustness proof in Section 8.2, and the proofs for utility-dominated emulation in Sections 8.3 and 8.4.

Remark 8.3 (Basing the protocol on standard assumptions when robustness is not needed). If we do not need the robustness property, we can replace the non-malleable timed commitments with a computationally hiding and perfectly binding commitment, plus a simulation extractable Non-Interactive Zero-Knowledge Proof (NIZK) proof that the commitment is computed correctly. With this change, the proof of Claim 8.9 would use the extraction key of the NIZK scheme to extract the bid under the commitment rather than calling the Dec oracle of the IND-CCA challenger. Thus, Claim 8.9 can be proven using the simulation extractability of the NIZK and the computational hiding property of the commitment against quasi-polynomially sized adversaries.

8.2 **Proofs of Robustness**

We prove the following lemma showing that our real-world protocol satisfies robustness.

Lemma 8.4. Suppose that NITC satisfies soundness of forced decryption for polynomially-sized adversaries, and that strategic players cannot hijack honest players' identities. Then, the real-world protocol $\Pi_{\lambda}^{\text{real}}$ satisfies robustness as defined at the beginning of Section 8.

In the above, the assumption that strategic players cannot hijack honest players' identities is mild and reflects realistic deployment scenarios. For example, this assumption is respected if the players' identities include their email addresses, or government-issued identities. Alternatively, if honest players choose their cryptographic identities at random by sampling public and private key pairs, then the probability that honest keys are hijacked by strategic players is also negligibly small.

Proof. (of Lemma 8.4) Suppose that the platform is honest. Since the strategic coalition cannot hijack honest players' identities, they cannot cause honest players to reject due to their identities being suppressed. Another potential way for the strategic coalition to cause honest players to reject is for some strategic buyer or seller denoted *i* to submit an NITC commitment, fail to provide the opening, and somehow the honest platform fails to get a valid π_i^{fdec} during the forced decryption. Due to the soundness of forced decryption of the underlying NITC, this happens with only negligible probability.

8.3 Proofs of Utility-Dominated Emulation: When the Platform is Strategic

Throughout the proof, we assume that $\mathbb{U}_{\lambda} \subset [0,1]$ is a discretized subset of the normalized value domain [0,1] — see also Section 4.1.

Experiment Hyb₁. Hyb₁ is almost the same as the real-world execution except with the following modifications. There must be one honest verifier V (either a buyer or seller) in the Step (h) of the protocol. Therefore, we can run the AoK's extractor along the side and extract a witness of the form $\left(\widetilde{\text{code}}, \widetilde{\mathcal{I}}, \{\widetilde{c}_j, \widetilde{\pi}_j^{\text{com}}, \widetilde{\pi}_j, \widetilde{v}_j, \widetilde{r}_j, \widetilde{\text{out}}_j\}_{j \in \widetilde{\mathcal{I}}}\right)$ for the AoK instance where V is the verifier. Abort and treat the coalition's utility as 0 if any of the following bad events happens:

- 1. the witness is not valid;
- 2. some honest $i \in \mathcal{I}$ is not included in the extracted $\widetilde{\mathcal{I}}$;
- 3. for some $i \in \widetilde{\mathcal{I}}$, NITC.FDec(NITC.crs, $\widetilde{c}_i, \widetilde{\pi}_i^{\text{com}}) \neq i \| \widetilde{v}_i \| \widetilde{r}_i;$
- 4. there exists some honest $i \in \mathcal{I}$ that accepts a private outcome out_i that is different from the extracted $\widetilde{\mathsf{out}}_i$.
- 5. for some $i \in \mathcal{I}$ that is honest, the message *i* sent to the platform in Step (a) is not contained in $\widetilde{\mathbf{c}}$ where $\widetilde{\mathbf{c}} = \{j, \widetilde{c}_j, \widetilde{\pi}_j^{\text{com}}\}_{j \in \widetilde{\mathcal{I}}}$ comes from the extracted values;
- 6. for some $i \in \mathcal{I}$ that is honest, the actual v_i and r_i values it used in the protocol differ from \tilde{v}_i and \tilde{r}_i from the extracted witness here we assume that $v_0 = \bot$ if the seller is honest.

Claim 8.5. Suppose that the AoK scheme satisfies knowledge soundness, the NITC scheme is computationally binding against quasi-polynomially sized adversaries, and VC is collision resistant against quasi-polynomially sized adversaries, then the coalition's utility in the real world has negligible statistical distance from its utility in Hyb_1 . *Proof.* We will go over each bad event one by one, and show that except with negligible probability, if all honest players accept in the end, the bad event should not happen.

- 1. Using a straightforward reduction to the knowledge soundness of AoK, we can show that the extracted witness is invalid with only negligible probability. Henceforth, we may assume that the extracted witness is valid.
- 2. If the extracted witness is valid, it must be digest' = VC.Commit(VC.crs, $\{i, \tilde{c}_i, \tilde{\pi}_i^{\text{com}}, \widetilde{\operatorname{out}}_i\}_{i \in \widetilde{\mathcal{I}}}$). If all honest players accepted at the end of the protocol, it must be that the platform showed a valid membership proof for some tuple (i, ..., ..., ...) w.r.t. digest' to every honest $i \in \mathcal{I}$. Thus, if not all honest $i \in \mathcal{I}$ is contained in $\widetilde{\mathcal{I}}$, we can construct a quasi-polynomial time reduction that outputs a collision for VC.
- 3. If the extracted witness is valid, then, for all $i \in \tilde{\mathcal{I}}$, either NITC.DecVf(NITC.crs, $\tilde{c}_i, i \| \tilde{v}_i \| \tilde{r}_i, \tilde{\pi}_i) = 1$ or NITC.FDecVf(NITC.crs, $\tilde{c}_i, i \| \tilde{v}_i \| \tilde{r}_i, \tilde{\pi}_i) = 1$ which is checked by the language \mathcal{L}_{λ} . Suppose for some $i \in \mathcal{I}$, the message output by NITC.FDec(NITC.crs, $\tilde{c}_i, \tilde{\pi}_i^{\text{com}})$ is not equal to $i \| \tilde{v}_i \| \tilde{r}_i$, we can construct a quasi-polynomial time adversary that breaks the computationally binding property of NITC.
- 4. If the extracted witness is valid, it must be $\mathsf{digest}' = \mathsf{VC.Commit}(\mathsf{VC.crs}, \{i, \tilde{c}_i, \tilde{\pi}_i^{\mathrm{com}}, \widetilde{\mathsf{out}}_i\}_{i \in \tilde{\mathcal{I}}})$. Thus, if for some honest *i*, the platform can convince it to accept some different $\mathsf{out}_i \neq \widetilde{\mathsf{out}}_i$ by producing a valid membership proof for some tuple $(i, _, _, \mathsf{out}_i)$, then we can build a quasipolynomial time reduction that breaks the collision resistance of the VC scheme.
- 5. The proof is similar to (4) above.
- 6. If the extracted witness is valid, then, either NITC.DecVf(NITC.crs, $\tilde{c}_i, i \| \tilde{v}_i \| \tilde{r}_i, \tilde{\pi}_i) = 1$ or NITC.FDecVf(NITC.crs, $\tilde{c}_i, i \| \tilde{v}_i \| \tilde{r}_i, \tilde{\pi}_i) = 1$. Given (5), we know that for an honest $i, \tilde{c}_i = c_i, \tilde{\pi}_i^{\text{com}} = \pi_i^{\text{com}}$ where c_i and π_i^{com} are what i sent in the protocol. By the correctness of NITC, if all honest players accept in the end, i must have also sent π_i^{dec} such that NITC.DecVf(NITC.crs, $c_i, i \| v_i \| r_i, \pi_i^{\text{dec}}) = 1$. If $\tilde{v}_i \| \tilde{r}_i \neq v_i \| r_i$, we can construct a quasi-polynomial time reduction that breaks the computationally binding property of NITC.

Experiment Hyb_2 . Hyb_2 is almost the same as Hyb_1 except the following modifications. We know that there is at least one honest verifier V (either a buyer or seller) for the challenge-response protocol in Steps (d) and (e), and we additionally run the following extractor along the side.

- Let code' be an array of length n' where all elements are initialized to \perp .
- Repeat $\lambda^{\alpha(\lambda)}$ times in parallel where $\alpha(\cdot)$ is an arbitrary super-constant function:
 - rewind the coalition's algorithm to right before it receives the challenge;
 - sample κ fresh random indices denoted Q on behalf of V and send Q to the coalition;
 - if the coalition responds with openings $\{z_q\}_{q \in Q}$ with valid proofs within in $T_3(\lambda)$ time, then populate the entry $\operatorname{code}'[q] := z_q$ for $q \in Q$.

We define the following bad event in Hyb_2 : all honest players accept in the end, but one of the following happens:

1. fewer than 2/3 fraction of the entries of code' have been populated; or

2. RS.Recons(code') $\neq \tilde{\mathbf{c}}$ where $\tilde{\mathbf{c}} = \{(j, \tilde{c}_j, \tilde{\pi}_j^{\text{com}})\}_{j \in \tilde{\mathcal{I}}}$ is part of the witness extracted by the AoK extractor.

Claim 8.6. Suppose that VC satisfies collision resistance against quasi-polynomial time adversaries. Then, the coalition's utility in Hyb_2 is identically distributed as in Hyb_1 , and moreover, the probability that the bad event happens is negligibly small.

Proof. Hyb_2 does not modify the way the outcomes are computed, it only runs an additional extractor for the challenge-response protocol corresponding to Steps (d) and (e) on the side. We now prove that the probability of encountering a bad event is negligibly small.

Since the witness output by the AoK extractor is valid, if all honest players accept at the end, they must all have the correct n, that is, digest is a vector commitment of a length n'-vector where n' is the length of the RS encoding of a length-n vector.

Applying Lemma 1 of Chiesa et al. $[CDG^+24]$, we can conclude that except with negligible probability, the above extractor extracts at least 2/3 locations, along with valid membership proofs w.r.t. digest. We know that digest = VC.Commit(VC.crs, code) and code = RS.Encode(\tilde{c}) since the witness output by the AoK extractor is valid. If there exists some location $i \in [n']$ such that $code'[i] \neq \bot$ and $code'[i] \neq code[i]$, then we can construct a quasi-polynomial time reduction that produces a collision for VC. If the above bad event does not happen, by the property of RS, it must be that RS.Recons(code') = \tilde{c} .

Experiment Hyb_3 . In Hyb_3 , if all honest players accept the auction, we instead compute the coalition's utility as follows:

- reconstruct $\mathbf{c} = \{(i, c_i, \pi_i^{\text{com}})\}_{i \in \mathcal{I}}$ from the code' extracted from the challenge-response protocol in Steps (d) and (e);
- for $i \in \mathcal{I} \setminus \mathcal{H}$ where \mathcal{H} denotes the honest players, let $(i || v_i || r_i, _) \leftarrow \mathsf{NITC.FDec}(\mathsf{NITC.crs}, c_i, \pi_i^{\mathrm{com}});$
- compute the coalition's utility by running the rules $(\mathbf{x}, \mathbf{p}, \mu_{\mathcal{S}})$ on the following inputs and randomness:
 - the honest buyer's true values $\{v_i\}_{i \in \mathcal{H} \setminus \{0\}}$, and the strategic bids $\{v_i\}_{i \in \mathcal{I} \setminus (\mathcal{H} \cup \{0\})}$ extracted above;
 - randomness $r = \begin{pmatrix} \oplus \\ i \in \mathcal{I} \end{pmatrix} \oplus r_{\mathcal{P}}$ where $\{r_i\}_{i \in \mathcal{H}}$ are the random coins chosen by honest players, $\{r_i\}_{i \in \mathcal{I} \setminus \mathcal{H}}$ are the strategic coins extracted above, and $r_{\mathcal{P}}$ denotes the $r_{\mathcal{P}}$ value the platform has sent to an arbitrary honest player.

If the above computation fails (e.g., due to failure of RS reconstruction), simply treat the coalition's utility as 0.

Claim 8.7. Under the same assumptions as Claim 8.6, the coalition's utility in Hyb_3 has negligible statistical difference from Hyb_2 .

Proof. If the bad events defined in Hyb_2 do not happen, then the two ways of computing the auction's outcomes are equivalent by observing that 1) the witness extracted by the AoK extractor is valid, i.e., the conditions checked by the language \mathcal{L}_{λ} all hold, and 2) the additional checks introduced by Hyb_1 all hold.

Experiment Hyb_4 . Hyb_4 is almost the same as Hyb_3 , except that we stop running the experiment after Step (e). We compute the coalition's utility like in Hyb_3 . If the utility is less than 0, then simply treat the coalition's utility as 0.

Claim 8.8. The coalition's utility in Hyb_4 stochastically dominates Hyb_3 .

Proof. Let tr be an execution trace of Hyb_4 , let tr' be a continuation of this trace for Hyb_3 . If not all honest players accept in tr', then the coalition's utility is 0 in Hyb_3 , and its utility has to be at least 0 in Hyb_4 . If all honest players accept in tr', then either the coalition's utility in Hyb_3 given tr' is the same as its utility in Hyb_4 given tr, or the coalition's utility is negative in Hyb_3 but rounded up to 0 in Hyb_4 .

Experiment Hyb₅. Hyb₅ is almost the same as Hyb_4 , except that when the experiment needs to compute NITC commitments on behalf of honest players, we replace the commitments with commitments of 0.

Claim 8.9. Suppose that the NITC scheme satisfies IND-CCA security for quasi-polynomially sized and depth- T^{ϵ} adversaries, then, the coalition's expected utility in Hyb₅ is only negligibly apart from its utility in Hyb₄.

Proof. We may assume that the number of buyers is exactly n-1 since we an apply the same proof for every choice of n. Without loss of generality, we will index the seller as 0, and index all honest buyers using $1, 2, \ldots, n-1$ below. The *i*-th buyer has identity id_i . Let $id_0 = 0$, and let $v_0 = \bot$ if the seller is honest.

We will prove this claim through a sequence of inner hybrid experiments denoted H_0, H_1, \ldots, H_n . H_0 is the same as Hyb_4 , and for $i \in [n]$, H_i is the same as H_{i-1} except for changing the (i-1)-th honest player's NITC commitment to a commitment of 0. Clearly, $H_n = Hyb_5$.

Due to Fact 3.2, it suffices to show that no polynomially sized, depth- $C \log(\cdot)$ distinguisher can effectively distinguish the coalition's utility in each pair of adjacent hybrids, where C is a suitably large constant.

Suppose that for some i, there is some polynomially sized, depth- $C \log(\cdot)$ distinguisher \mathcal{B} that can distinguish the coalition's utilities in H_{i-1} and H_i with non-negligible probability. We will reach a contradiction by constructing the following reduction \mathcal{R} that wants to break the IND-CCA security of NITC.

- 1. \mathcal{R} first prepares the commitments for all honest players except the (i-1)-th player. Specifically, for any honest buyer j < i 1, it will compute a commitment of 0. For any honest buyer ranked $j \ge i$, it will compute an honest commitment of its true value v_j and some random r_j . Let the resulting commitments be $(c_0, \pi_0^{\text{com}}), \ldots, (c_{i-2}, \pi_{i-2}^{\text{com}}), (c_i, \pi_i^{\text{com}}), \ldots, (c_{n-1}, \pi_{n-1}^{\text{com}})$. This preprocessing only needs to be quasi-polynomial in total work and we need not worry about its depth.
- 2. Next, \mathcal{R} sends a challenge to the IND-CCA challenger on messages $(0, v_{i-1}||r_{i-1})$ where v_{i-1} is player (i-1)'s true value and r_{i-1} is some random string. It obtains some commitment $(c_{i-1}, \pi_{i-1}^{\text{com}})$ from the IND-CCA challenger.
- 3. \mathcal{R} invokes the coalition's algorithm, and sends $\{id_i, c_i, \pi_i^{\text{com}}\}_{i \in \{0, 1, \dots, n-1\}}$ to the coalition \mathcal{C} . \mathcal{C} responds with $(n, \text{digest}, r_{\mathcal{P}})$. \mathcal{R} now runs the extractor defined in Hyb_2 with \mathcal{C} , and extracts some code' which reconstructs to **c** using the reconstruction algorithm of the RS code. \mathcal{R} now computes the coalition's utility using **c** just like in Hyb_4 , but where the calls to FDec are replaced with calls to the Dec oracle provided by the IND-CCA challenger.

4. \mathcal{R} now forwards the resulting utility to \mathcal{B} and outputs whatever it outputs.

Clearly, the work of \mathcal{R} is quasi-polynomially bounded. We now analyze the depth of \mathcal{R} in Steps 3 and 4. The depth of Step 3 is upper bounded by

$$T_2(\lambda) + T_3(\lambda) + D_{rs}(\lambda^{\alpha(\lambda)}) + D_{auction}(\lambda^{\alpha(\lambda)}) + \mathsf{poly}\log\lambda$$

where

- $T_2(\lambda)$ is the maximum time for the coalition to produce the tuple $(n, \text{digest}, r_{\mathcal{P}})$;
- $T_3(\lambda)$ is the maximum time for the coalition to respond to an opening challenge for VC;
- $D_{\rm rs}(\lambda^{\alpha(\lambda)})$ is an upper bound on the depth of the reconstruction algorithm of RS for an input of unbounded polynomial length;
- $D_{\text{auction}}(\lambda^{\alpha(\lambda)})$ is an upper bound on the depth of the algorithm that runs the auction's rules $(\mathbf{x}, \mathbf{p}, \mu_{\mathcal{S}})$ and then computes the coalition's utility from the coalition's true value vector and the outcomes of the auction for any unbounded polynomial length input;
- and poly' $\log(\lambda)$ captures the depth of all other operations, e.g., for a machine that does not support concurrent writes, concurrent writes to the same array code' may take $O(\log(\lambda^{\alpha(\lambda)}))$ depth.

The depth of Step 4 is upper bounded by the maximum depth of the distinguisher \mathcal{B} whose depth is logarithmic in the input length. Because the input length can be an unbounded polynomial, we can upper bound the depth of \mathcal{B} by $\log^2 \lambda$ for sufficiently large λ . Because we do not need any error correction, the reconstruction algorithm of RS simply computes a linear combination, and its depth is logarithmic in the input length. Therefore, for a suitably small $\alpha(\cdot)$, and sufficiently large λ , $D_{\rm rs}(\lambda^{\alpha(\lambda)})$ is upper bounded by $\log^2 \lambda$. Summarizing the above, the total depth of \mathcal{R} in Steps 3 and 4 is upper bounded by

$$T_2(\lambda) + T_3(\lambda) + D_{auction}(\lambda^{\alpha(\lambda)}) + \text{poly} \log \lambda$$

Recall that we choose the parameters T and ϵ of the NITC to satisfy

$$T^{\epsilon}(\lambda) > T_2(\lambda) + T_3(\lambda) + D_{\text{auction}}(\lambda^{\alpha(\lambda)}) + \operatorname{\mathsf{poly}}\log(\lambda)$$

Therefore, if \mathcal{B} can effectively distinguish the coalition's utility in H_{i-1} and H_i , then \mathcal{R} would break the IND-CCA security of the underlying NITC.

Ideal auction and ideal strategy. Hyb_5 is almost equivalent to the following ideal auction in terms of the coalition's utility, except that when NITC.FDec fails, the coalition may have positive utility in the ideal auction below below, whereas in Hyb_5 it has 0 utility:

- 1. Honest buyers send their true values to $\mathcal{F}_{auction}$;
- 2. $\mathcal{F}_{\text{auction}}$ sends the number of honest bids n_H to \mathcal{C} ;
- 3. C computes n_H NITC commitments of 0 and sends them to the real-world coalition's algorithm; when the real-world coalition's algorithm responds with some (n, digest, r_P) , it runs the extractor algorithm defined in Hyb₂ with the real-world coalition's algorithm, to extract and reconstruct $\mathbf{c} = \{i, c_i, \pi_i^{\text{com}}\}_{i \in \mathcal{I}}$. Then, for every $i \in \mathcal{I}$ that is not an honest identity, C calls NITC.FDec(NITC.crs, c_i, π_i^{com}) to decrypt $v_i || r_i$. If FDec failed, simply send 0 bids on behalf of all strategic buyers to $\mathcal{F}_{\text{auction}}$; otherwise, send the extracted bids corresponding to strategic buyers to $\mathcal{F}_{\text{auction}}$.

4. $\mathcal{F}_{auction}$ chooses the coins r at random, and computes the outcome of the auction using the honest players' values, the bids submitted by C, the randomness r. It sends the outcome vector to C. If the coalition's utility is less than 0, the platform sends \perp to $\mathcal{F}_{auction}$; else it sends ok to $\mathcal{F}_{auction}$.

Recall that at least one buyer or the seller is honest (henceforth denoted i). Given that the

honest player *i* chooses its random coins r_i at random, the joint coin toss result $r = \begin{pmatrix} \oplus \\ j \in \mathcal{I} \end{pmatrix} \oplus r_{\mathcal{P}}$ is also a uniform random string. Therefore, it is equivalent for $\mathcal{F}_{\mathsf{auction}}$ to just choose *r* at random in the above.

Summarizing the above, we have that the coalition's expected utility in the real world cannot be negligibly more than its expected utility in the ideal world.

8.4 Proofs of Utility Dominated Emulation: When the Platform is Honest

This is the easier case. The coalition C consists of either one or more buyers, or just the seller. Consider the following sequence of hybrid experiments. Throughout the proof below, we assume that the coalition's algorithm is PPT, and all the cryptographic building blocks only need to be secure against PPT adversaries.

Experiment Hyb₁. Hyb₁ is the same as the real-world protocol except that when \mathcal{C} sends NITC commitments of the strategic bids $\{(j, c_j, \pi_j^{\text{com}})\}_{j \in \mathcal{J}'}$, we ignore all the tuples where π_j^{com} does not verify or tuples with identity 0 but not sent from the seller, perform duplicate suppression on the identities and notify the suppressed identities as before, but additionally, we call the NITC.FDec function to force-open $\{(v_j, r_j)\}_{j \in \mathcal{J}}$ where \mathcal{J} is the set of strategic identities after duplicate suppression.

If all honest players accept at the end, we compute the coalition's utility using 1) the honest buyers' true values $\{v_i\}_{i \in \mathcal{H} \setminus \{0\}}$ where \mathcal{H} denotes the honest players (not including the platform), as well as honest randomness $\{v_i\}_{i \in \mathcal{H}}$ and $r_{\mathcal{P}}$; and 2) the force-opened bids $\{v_j\}_{j \in \mathcal{J} \setminus \{0\}}$ and forcedopened randomness $\{r_j\}_{j \in \mathcal{J}}$ corresponding to strategic players.

Claim 8.10. Suppose that the NITC scheme is computationally binding. Then, the coalition's utility in Hyb_1 has negligible statistical distance from the real world.

Proof. The only reasons that would cause the coalition's utility to differ is if some strategic $j \in \mathcal{J}$ successfully opened its NITC commitment to a different tuple than the forced opened tuple. This happens with only negligible probability as long as NITC is computationally sound.

Experiment Hyb_2 . In Hyb_2 , we stop executing the protocol after the coalition submits their NITC commitments. If no honest identity has been suppressed, we use the approach of Hyb_1 to compute the coalition's identity. Otherwise, we treat the coalition's identity as 0.

Claim 8.11. Suppose that NITC satisfies soundness of forced decryption. The coalition's utility in Hyb_1 cannot be negligibly more than its utility in Hyb_2 .

Proof. The only way for the coalition to gain more utility in Hyb_1 than Hyb_2 is if the utility computed through the forced-opened strategic inputs is negative, but in Hyb_1 , the coalition manages to make the execution abort by not submitting an opening, and FDec fails to produce a valid proof. Note that we implicitly assume that FDec always outputs a message even if it does not output a valid proof. By the soundness of forced decryption of NITC, this happens with negligible probability.

Ideal auction and strategy. Note that Hyb_2 can be equivalently viewed as an ideal auction with the following coalition strategy. The coalition simply invokes the real-world C's algorithm, and wait for it to send the NITC commitments. It then verify the commitment proofs, suppresses duplicate identities as well as any buyer that claims the identity of 0, and performs force decryption as mentioned above. It sends the forced-opened strategic bids to $\mathcal{F}_{auction}$. $\mathcal{F}_{auction}$ then tosses random coins, and computes the outcome of the auction honestly. We then determine the coalition's utility based on the outcome and the true value of the buyers in C.

Summarizing the above, the coalition's expected utility in the real world cannot be negligibly more than its utility in the ideal world.

9 Ideal Auction: 2nd-Price with Reserve and Fixed Platform Fees

9.1 Ideal Auction

We assume that the value domain is discrete like in Section 5. However, in this section, the auction needs to be parametrized by the security parameter λ since we want to compile it to a real-world cryptographic protocol using the compiler in Section 8. More formally, suppose the value domain \mathbb{U}_{λ} is finite and $0 \in \mathbb{U}_{\lambda}$ for every λ . Suppose \mathbb{U}_{λ} consists of the values $0 = \theta_{\lambda}^{1} < \theta_{\lambda}^{2} < \ldots < \theta_{\lambda}^{T(\lambda)}$. Just like in Section 5, given some distribution \mathcal{D}_{λ} over \mathbb{U}_{λ} with the cumulative distribution function F_{λ} and the probability density function f_{λ} , the virtual value $\phi(\theta_{\lambda}^{i}) := \theta_{\lambda}^{i} - \frac{1 - F_{\lambda}(\theta_{\lambda}^{i})}{f_{\lambda}(\theta_{\lambda}^{i})} (\theta_{\lambda}^{i+1} - \theta_{\lambda}^{i})$, and $\phi(\theta_{\lambda}^{T}) := \theta_{\lambda}^{T}$. We say that the distribution \mathcal{D}_{λ} is *regular*, iff $\phi(\cdot)$ is a strictly increasing function. Let the reserve price r_{λ} be the smallest θ_{λ}^{i} such that $\phi(\theta_{\lambda}^{i}) \geq 0$.

We describe the ideal auction below — recall that to specify an ideal auction, all we need is to specify the allocation, payment, and revenue rules:

Ideal auction: 2nd price with reserve and fixed platform fees

- Discard any bid that is less than the reserve r_{λ} . Rank the remaining bids in order from large to small, and break ties randomly. Let $b_1 \ge b_2 \ge \ldots \ge b_n$ be the resulting vector of bids, and rename corresponding buyers as $1, 2, \ldots, n$.
- Buyers $1, 2, \ldots, k'$ are allocated an item, where $k' = \min(k, n)$ henceforth these are called the confirmed buyers and confirmed bids.
- Define $b_{k'+1} := r_{\lambda}$ if $n \leq k$. Define the payment as follows where A be the total number of bids that are equal to $b_{k'+1}$, and α is the number of confirmed bids that are equal to $b_{k'+1}$:
 - Any confirmed buyer whose bid is equal to $b_{k'+1}$ pays $b_{k'+1}$;
 - Any confirmed buyer whose bid is strictly greater than $b_{k'+1}$ pays $b_{k'+1} \cdot q + \mathsf{NextTick}(b_{k'+1}) \cdot (1-q)$ where $q = \frac{\alpha+1}{A+1}$, and $\mathsf{NextTick}(v)$ means the smallest value larger than v in \mathbb{U}_{λ} .
- The platform gets nothing (or a fixed fee independent of the auction), and the seller gets all the payment.

Theorem 9.1 (Ideal auction: second price with reserve and fixed platform fees). Suppose that the above second-price auction with reserve and fixed platform fees is executed under the ideal model described in Section 7.1. Then, the resulting ideal auction has the following properties:

• It satisfies information-theoretic bIC, pIC, 1-pbIC regardless of \mathcal{D}_{λ} ;

- Suppose that the distribution \mathcal{D}_{λ} is regular for every λ , then the auction additionally satisfies; Bayesian psIC, and Bayesian sIC.
- Suppose that the distribution \mathcal{D}_{λ} is regular for every λ , then the auction is the revenue maximizing auction among all bIC auctions;
- The auction's rules can be computed depth $O(\log^2(L))$ where L is the total input length.

The proof of Theorem 9.1 is presented in Section 9.3.

9.2 Applying the Ideal-to-Real Compiler

The above Theorem 9.1, combined with the ideal-to-real compiler described in Theorem 8.2 of Section 8.1, as well as Theorem 7.2, immediately leads to the following corollary:

Corollary 9.2 (Achieving bIC, pIC, pbIC, Bayesian psIC and Bayesian sIC in O(1) rounds). Suppose that the strong repeated squaring assumption and the hardness of the DDH and DCR problems in suitable groups hold against quasi-polynomially sized adversaries. Then, there exists an auction that satisfies computational bIC, pIC, pbIC regardless of the value distribution \mathcal{D} ; further, assuming that the value distribution \mathcal{D} is regular, then it additionally satisfies computational Bayesian psIC, and Bayesian sIC. Moreover, the auction satisfies the following properties:

- It completes in constant number of rounds with $\widetilde{O}_{\lambda}(n)$ total communication and computation, and each buyer and the seller's computation and communication is bounded by $\widetilde{O}_{\lambda}(1)$.
- When all buyers' values are sampled independently from a regular distribution \mathcal{D} , the auction maximizes the expected revenue among all bIC auctions.

9.3 Proof of Theorem 9.1

Additional preliminaries. We will rely on the following technical lemma proven by Elkind [Elk07].

Lemma 9.3 (Technical lemma from [Elk07]). Suppose that the value domain \mathbb{U} is a finite set consisting of $0 = \theta^0 < \theta^1 < \ldots < \theta^T$. Suppose each buyer's true value is sampled independently from some distribution \mathcal{D} over \mathbb{U} . Given any $\mathbf{b}_{-i} \in \mathbb{U}^*$, any non-decreasing allocation rule \mathbf{x} , suppose that the payment rule is defined as follows:

$$p_i(\theta^{\tau}, \mathbf{b}_{-i}) = \theta^{\tau} \cdot x_i(\theta^{\tau}, \mathbf{b}_{-i}) - \sum_{j=1}^{\tau} \left(\theta^j - \theta^{j-1}\right) \cdot x_i\left(\theta^{j-1}, \mathbf{b}_{-i}\right), \tag{6}$$

Then, the resulting auction is incentive compatible for any individual buyer that is restricted to input replacement strategies. Moreover, when all buyers' bids are sampled independently from \mathcal{D} , it is the revenue-maximizing auction subject to bIC and the allocation rule **x**. Specifically, given an arbitrary \mathbf{b}_{-i} , the expected payment from any buyer i is the following:

$$\mathbf{E}_{b_i \stackrel{\$}{\leftarrow} \mathcal{D}}[p_i(b_i, \mathbf{b}_{-i})] = \mathbf{E}_{b_i \stackrel{\$}{\leftarrow} \mathcal{D}}[x_i(b_i, \mathbf{b}_{-i}) \cdot \phi(b_i)].$$

We now prove the claimed properties one by one.

bIC. It is not hard to check that the allocation rule is monotone, and the auction's payment rule agrees with Equation (6) specified in Lemma 9.3. Thus, by Lemma 9.3, any input replacement strategy cannot benefit an individual buyer. It remains to prove that injecting fake bids does not benefit the buyer. Let $b^* = r_{\lambda}$ if $n \leq k$ or $b^* = b_{k+1}$ otherwise. Any fake bid that is less than b^* does not affect the outcome of the auction. Any fake bid that is at least b^* can never decrease the payment if the buyer gets an item. Therefore, injecting fake bids does not help the buyer.

pIC. Since the platform's revenue is always zero, pIC is trivially satisfied.

1-pbIC. Directly implied by bIC and the fact that the platform's revenue is always 0.

Bayesian psIC. The only possible strategies for a platform-seller coalition in the ideal world is 1) to abort the auction causing every one's utility to be 0, and 2) to inject fake bids after learning the number of honest bids. Clearly, aborting will not benefit the coalition. Therefore, it suffices to show that injecting fake bids does not increase the coalition's expected utility.

Injecting any bid that is less than r_{λ} makes no difference to the auction. Therefore, let \mathbf{b}^* be the fake bids that are at least r_{λ} injected by the coalition. Let n be the number of honest buyers. By Lemma 9.3, the coalition's expected revenue from a fixed honest buyer i is

$$\mathbf{E}_{\mathbf{b} \stackrel{\$}{\leftarrow} \mathcal{D}^n}[p_i(\mathbf{b}, \mathbf{b}^*)] = \mathbf{E}_{\mathbf{b} \stackrel{\$}{\leftarrow} \mathcal{D}^n}[x_i(\mathbf{b}, \mathbf{b}^*) \cdot \phi(b_i)] \le \mathbf{E}_{\mathbf{b} \stackrel{\$}{\leftarrow} \mathcal{D}^n}[x_i(\mathbf{b}) \cdot \phi(b_i)]$$

where the expression on the right-hand side $\mathbf{E}_{\mathbf{b} \leftarrow \mathcal{D}^n} [x_i(\mathbf{b}) \cdot \phi(b_i)]$ corresponds to the expected revenue from buyer *i* when the coalition behaves honestly. Therefore, injecting fake bids does not increase the coalition's expected revenue.

Bayesian sIC. A strategic seller's only possible strategy is to inject fake bids. Recall that the platform gets nothing and all revenue goes to the seller. Therefore, due to the same argument as the above proof of Bayesian psIC, injecting fake bids does not increase the seller's revenue.

Revenue optimality. Let *n* be the number of buyers. By Lemma 9.3, for any bIC auction with a

monotonic allocation rules \mathbf{x} , the maximum possible expected revenue is $\sum_{i \in [n]} \left(\mathbf{E}_{\mathbf{b} \leftarrow \mathcal{D}^n} \left[x_i(\mathbf{b}) \cdot \phi(b_i) \right] \right) = \mathbf{E}_{\mathbf{b} \leftarrow \mathcal{D}^n} \left[\sum_{i \in [n]} x_i(\mathbf{b}) \cdot \phi(b_i) \right]$. It suffices to argue that our choice of the allocation rule \mathbf{x} maximizes $\mathbf{E}_{\mathbf{b} \leftarrow \mathcal{D}^n} \left[\sum_{i \in [n]} x_i(\mathbf{b}) \cdot \phi(b_i) \right].$ For any **b**, we argue that $\sum_{i \in [n]} x_i(\mathbf{b}) \cdot \phi(b_i)$ is maximized under our choice of **x**. This is because our **x** allocates one item to those with the largest virtual value $\phi(b_i)$. subject to allocating at most k items; moreover, it only allocates an item to those with non-negative virtual values.

Small depth. The auction's rules can be computed by a sorting network consisting of comparators. Each comparator can be implemented in logarithmic depth in its input length, and the sorting network needs only logarithmic layers [AKS83]. Therefore, it follows that the auction's rules can be computed in $O(\log^2(L))$ depth.

Acknowledgments

Elaine Shi would like to thank Dakshita Khurana, Binyi Chen, and Sacha Servan-Schreiber for helpful technical discussions. This work is in part supported by NSF awards under the grant numbers 2212746 and 2044679, a Packard Fellowship, a generous gift from the late Nikolai Mushegian, an an ACE center grant from the Algorand Foundation.

References

- [AJLA⁺12] Gilad Asharov, Abhishek Jain, Adriana López-Alt, Eran Tromer, Vinod Vaikuntanathan, and Daniel Wichs. Multiparty computation with low communication, computation and interaction via threshold fhe. In Proceedings of the 31st Annual International Conference on Theory and Applications of Cryptographic Techniques, EURO-CRYPT'12, page 483–501, Berlin, Heidelberg, 2012. Springer-Verlag.
- [AKS83] M. Ajtai, J. Komlós, and E. Szemerédi. An O(n log n) sorting network. In *STOC*, 1983.
- [AL20] Mohammad Akbarpour and Shengwu Li. Credible auctions: A trilemma. *Econometrica, Econometric Society*, 2020.
- [AN20] Ramiro Alvarez and Mehrdad Nojoumian. Comprehensive survey on privacy-preserving protocols for sealed-bid auctions. *Computers & Security*, 88:101502, 2020.
- [BBHR19] Eli Ben-Sasson, Iddo Bentov, Yinon Horesh, and Michael Riabzev. Scalable zero knowledge with no trusted setup. In Advances in Cryptology - CRYPTO 2019 - 39th Annual International Cryptology Conference, Santa Barbara, CA, USA, August 18-22, 2019, Proceedings, Part III, volume 11694 of Lecture Notes in Computer Science, pages 701– 732. Springer, 2019.
- [BCD⁺09] Peter Bogetoft, Dan Lund Christensen, Ivan Damgård, Martin Geisler, Thomas Jakobsen, Mikkel Krøigaard, Janus Dam Nielsen, Jesper Buus Nielsen, Kurt Nielsen, Jakob Pagter, Michael Schwartzbach, and Tomas Toft. Secure multiparty computation goes live. In *Financial Cryptography and Data Security*, page 325–343, Berlin, Heidelberg, 2009. Springer-Verlag.
- [BGG⁺18] Dan Boneh, Rosario Gennaro, Steven Goldfeder, Aayush Jain, Sam Kim, Peter M. R. Rasmussen, and Amit Sahai. Threshold cryptosystems from threshold fully homomorphic encryption. In Advances in Cryptology – CRYPTO 2018: 38th Annual International Cryptology Conference, Santa Barbara, CA, USA, August 19–23, 2018, Proceedings, Part I, page 565–596, Berlin, Heidelberg, 2018. Springer-Verlag.
- [BJO09] Kevin D. Bowers, Ari Juels, and Alina Oprea. Proofs of retrievability: theory and implementation. In Proceedings of the 2009 ACM Workshop on Cloud Computing Security, CCSW '09, page 43–54, New York, NY, USA, 2009. Association for Computing Machinery.
- [BK18] Erik-Oliver Blass and Florian Kerschbaum. Strain: A secure auction for blockchains. In Computer Security: 23rd European Symposium on Research in Computer Security, ESORICS 2018, Barcelona, Spain, September 3-7, 2018, Proceedings, Part I, page 87–110, Berlin, Heidelberg, 2018. Springer-Verlag.
- [BSCR⁺19] Eli Ben-Sasson, Alessandro Chiesa, Michael Riabzev, Nicholas Spooner, Madars Virza, and Nicholas P. Ward. Aurora: Transparent succinct arguments for r1cs. In Advances in Cryptology – EUROCRYPT 2019, pages 103–128, Cham, 2019. Springer International Publishing.
- [BSS08] Eli Ben-Sasson and Madhu Sudan. Short pcps with polylog query complexity. *SIAM Journal on Computing*, 38(2):551–607, 2008.

- [Can00] Ran Canetti. Security and composition of multiparty cryptographic protocols. *Journal* of Cryptology, 2000.
- [Can01] R. Canetti. Universally composable security: A new paradigm for cryptographic protocols. In *FOCS*, 2001.
- [CBBZ23] Binyi Chen, Benedikt Bünz, Dan Boneh, and Zhenfei Zhang. Hyperplonk: Plonk with linear-time prover and high-degree custom gates. In Advances in Cryptology – EUROCRYPT 2023, pages 499–530, Cham, 2023. Springer Nature Switzerland.
- [CDG⁺24] Alessandro Chiesa, Marcel Dall'Agnol, Ziyi Guan, Nicholas Spooner, and Eylon Yogev. Untangling the security of kilian's protocol: Upper and lower bounds. In *TCC*, 2024.
- [CFK23] Tarun Chitra, Matheus VX Ferreira, and Kshitij Kulkarni. Credible, optimal auctions via blockchains. *arXiv preprint arXiv:2301.12532*, 2023.
- [CHM⁺20] Alessandro Chiesa, Yuncong Hu, Mary Maller, Pratyush Mishra, Psi Vesely, and Nicholas P. Ward. Marlin: Preprocessing zksnarks with universal and updatable SRS. In Advances in Cryptology - EUROCRYPT 2020 - 39th Annual International Conference on the Theory and Applications of Cryptographic Techniques, Zagreb, Croatia, May 10-14, 2020, Proceedings, Part I, volume 12105 of Lecture Notes in Computer Science, pages 738–768. Springer, 2020.
- [CJ23] Peter Chvojka and Tibor Jager. Simple, fast, efficient, and tightly-secure non-malleable non-interactive timed commitments. In Public-Key Cryptography – PKC 2023: 26th IACR International Conference on Practice and Theory of Public-Key Cryptography, Atlanta, GA, USA, May 7–10, 2023, Proceedings, Part I, page 500–529, Berlin, Heidelberg, 2023. Springer-Verlag.
- [CK04] Indranil Chakraborty and Georgia Kosmopoulou. Auctions with shill bidding. *Economic Theory*, 24(2):271–287, 2004.
- [CRS24] Hao Chung, Tim Roughgarden, and Elaine Shi. Collusion-resilience in transaction fee mechanism design. In *EC*, 2024.
- [CS23] Hao Chung and Elaine Shi. Foundations of transaction fee mechanism design. In *SODA*, 2023.
- [DGLS22] S. Dov Gordon, Feng-Hao Liu, and Elaine Shi. Constant-round mpc with fairness and guarantee of output delivery. In Advances in Cryptology – CRYPTO 2015, page 63–82, Berlin, Heidelberg, 2022. Springer-Verlag.
- [Elk07] Edith Elkind. Designing and learning optimal finite support auctions. 2007.
- [FW20] Matheus V. X. Ferreira and S. Matthew Weinberg. Credible, truthful, and two-round (optimal) auctions via cryptographic commitments. In *EC*, 2020.
- [GGH⁺13] Sanjam Garg, Craig Gentry, Shai Halevi, Mariana Raykova, Amit Sahai, and Brent Waters. Candidate indistinguishability obfuscation and functional encryption for all circuits. In *IEEE Symposium on Foundations of Computer Science (FOCS)*, 2013.

- [GGPR13] Rosario Gennaro, Craig Gentry, Bryan Parno, and Mariana Raykova. Quadratic span programs and succinct nizks without pcps. In Advances in Cryptology - EUROCRYPT 2013, 32nd Annual International Conference on the Theory and Applications of Cryptographic Techniques, Athens, Greece, May 26-30, 2013. Proceedings, volume 7881 of Lecture Notes in Computer Science, pages 626–645. Springer, 2013.
- [GH05] Andrew V. Goldberg and Jason D. Hartline. Collusion-resistant mechanisms for singleparameter agents. In SODA 2005, pages 620–629, 2005.
- [GLM⁺24] Vipul Goyal, Junru Li, Ankit Kumar Misra, Rafail Ostrovsky, Yifan Song, and Chenkai Weng. Dishonest majority constant-round MPC with linear communication from DDH. In Asiacrypt, 2024.
- [Gro16] Jens Groth. On the size of pairing-based non-interactive arguments. In Proceedings, Part II, of the 35th Annual International Conference on Advances in Cryptology — EUROCRYPT 2016 - Volume 9666, page 305–326, Berlin, Heidelberg, 2016. Springer-Verlag.
- [GWC19] Ariel Gabizon, Zachary J. Williamson, and Oana Ciobotaru. PLONK: Permutations over lagrange-bases for oecumenical noninteractive arguments of knowledge. Cryptology ePrint Archive, Paper 2019/953, 2019.
- [HW15] Pavel Hubacek and Daniel Wichs. On the communication complexity of secure function evaluation with long output. In Proceedings of the 2015 Conference on Innovations in Theoretical Computer Science, ITCS '15, page 163–172, New York, NY, USA, 2015. Association for Computing Machinery.
- [JLS21] Aayush Jain, Huijia Lin, and Amit Sahai. Indistinguishability obfuscation from wellfounded assumptions. In *Proceedings of the 53rd Annual ACM SIGACT Symposium on Theory of Computing*, STOC 2021, page 60–73. Association for Computing Machinery, 2021.
- [Kil92] Joe Kilian. A note on efficient zero-knowledge proofs and arguments (extended abstract). In Proceedings of the Twenty-Fourth Annual ACM Symposium on Theory of Computing, STOC '92, page 723–732, New York, NY, USA, 1992. Association for Computing Machinery.
- [KLX20] Jonathan Katz, Julian Loss, and Jiayu Xu. On the security of time-lock puzzles and timed commitments. In *Theory of Cryptography: 18th International Conference, TCC* 2020, Durham, NC, USA, November 16–19, 2020, Proceedings, Part III, page 390–413, Berlin, Heidelberg, 2020. Springer-Verlag.
- [KMS⁺16] Ahmed Kosba, Andrew Miller, Elaine Shi, Zikai Wen, and Charalampos Papamanthou. Hawk: The blockchain model of cryptography and privacy-preserving smart contracts. In 2016 IEEE Symposium on Security and Privacy (SP), pages 839–858, 2016.
- [Mer89] Ralph C. Merkle. A certified digital signature. In G. Brassard, editor, *Proc. CRYPTO* '89, volume 435 of *LNCS*, pages 218–238. Springer-Verlag, 1989.
- [MHAG22] Jason Milionis, Dean Hirsch, Andy Arditi, and Pranav Garimidi. A framework for single-item nft auction mechanism design. In Proceedings of the 2022 ACM CCS Workshop on Decentralized Finance and Security, DeFi'22, page 31–38, New York, NY, USA, 2022. Association for Computing Machinery.

- [Mon23] Karina Montova. How google manipulated digital prices ad and hurt publishers, per doj. https://www.techpolicy.press/ how-google-manipulated-digital-ad-prices-and-hurt-publishers-per-doj/, 2023.
- [MT19] Giulio Malavolta and Sri Aravinda Krishnan Thyagarajan. Homomorphic time-lock puzzles and applications. In Advances in Cryptology – CRYPTO 2019: 39th Annual International Cryptology Conference, Santa Barbara, CA, USA, August 18–22, 2019, Proceedings, Part I, page 620–649, Berlin, Heidelberg, 2019. Springer-Verlag.
- [Mye81] Roger B. Myerson. Optimal auction design. Math. Oper. Res., 6(1), 1981.
- [NPS99] Moni Naor, Benny Pinkas, and Reuban Sumner. Privacy preserving auctions and mechanism design. In Proceedings of the 1st ACM Conference on Electronic Commerce, EC '99, page 129–139, New York, NY, USA, 1999. Association for Computing Machinery.
- [NRTV07] Noam Nisan, Tim Roughgarden, Eva Tardos, and Vijay V. Vazirani. Algorithmic Game Theory. Cambridge University Press, USA, 2007.
- [oJ23] US Department of Justice. Justice department sues google for monopolizing digital advertising technologies, 2023.
- [Ope23] OpenSea. How to sell an nft, 2023.
- [Pan22] Ravi Panguluri. Taylor swift ticket sales should be decentralized. blockchain could make that possible. https://dbknews.com/2022/11/21/ ticketmaster-antitrust-potential-solution-is-blockchain-technology/, 2022.
- [Rou20] Tim Roughgarden. Transaction fee mechanism design for the Ethereum blockchain: An economic analysis of EIP-1559. Manuscript, https://timroughgarden.org/papers/ eip1559.pdf, 2020.
- [Rou21] Tim Roughgarden. Transaction fee mechanism design. In EC, 2021.
- [SCW23] Elaine Shi, Hao Chung, and Ke Wu. What can cryptography do for decentralized mechanism design. In *ITCS*, 2023.
- [Set20] Srinath Setty. Spartan: Efficient and general-purpose zksnarks without trusted setup. In Advances in Cryptology – CRYPTO 2020: 40th Annual International Cryptology Conference, CRYPTO 2020, Santa Barbara, CA, USA, August 17–21, 2020, Proceedings, Part III, page 704–737, Berlin, Heidelberg, 2020. Springer-Verlag.
- [SW13] Hovav Shacham and Brent Waters. Compact proofs of retrievability. J. Cryptol., 26(3):442–483, July 2013.
- [TAF⁺23] Nirvan Tyagi, Arasu Arun, Cody Freitag, Riad Wahby, Joseph Bonneau, and David Mazières. Riggs: Decentralized sealed-bid auctions. In Proceedings of the 2023 ACM SIGSAC Conference on Computer and Communications Security, CCS '23, page 1227–1241, New York, NY, USA, 2023. Association for Computing Machinery.
- [Vic61] William Vickrey. Counterspeculation, Auctions, And Competitive Sealed Tenders. Journal of Finance, 16(1):8–37, March 1961.

- [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.
- [XZZ⁺19] Tiacheng Xie, Jiaheng Zhang, Yupeng Zhang, Charalampos Papamanthou, and Dawn Song. Libra: Succinct zero-knowledge proofs with optimal prover computation. In Advances in Cryptology – CRYPTO 2019: 39th Annual International Cryptology Conference, Santa Barbara, CA, USA, August 18–22, 2019, Proceedings, Part III, page 733–764, Berlin, Heidelberg, 2019. Springer-Verlag.
- [ZXZS20] Jiaheng Zhang, Tiancheng Xie, Yupeng Zhang, and Dawn Song. Transparent polynomial delegation and its applications to zero knowledge proof. In 2020 IEEE Symposium on Security and Privacy (S&P), pages 859–876, 2020.

A Preliminaries: Cryptographic Building Blocks

In this section, we formally define the cryptographic primitives used in 8.1. Recall that we use λ to denote the security parameter. $\{X_{\lambda}\}_{\lambda \in \mathbb{N}} \equiv_{C} \{Y_{\lambda}\}_{\lambda \in \mathbb{N}}$ implies that the two distribution ensembles $\{X_{\lambda}\}_{\lambda \in \mathbb{N}}$ and $\{Y_{\lambda}\}_{\lambda \in \mathbb{N}}$ are computationally indistinguishable. We use PPT as an abbreviation for probabilistic polynomial time.

A.1 Vector Commitment

A vector commitment scheme is a tuple of algorithms (Gen, Digest, Open, Vf):

- crs \leftarrow Gen (1^{λ}) : on input the security parameter 1^{λ} , output a common reference string crs;
- (digest, aux) \leftarrow Digest(crs, m): given crs and a message m, output a digest digest and some auxiliary information aux we may assume that aux contains the message length $\ell := |m|$;
- Open(crs, aux, Q): on input crs, auxiliary information aux (assumed to contain the message length ℓ), and a query set $Q \subseteq [\ell]$, output an opening proof π that m[Q] is a restriction of m to the indices Q;
- $(0,1) \leftarrow \mathsf{Vf}(\mathsf{crs}, \ell, \mathsf{digest}, Q, \mathsf{ans}, \pi)$: on input crs , message length ℓ , digest, a query set $Q \subseteq [\ell]$, a purported answer ans , and a proof π , outputs either 0 or 1 indicating reject or accept.

Correctness. Correctness requires that for any $\lambda \in \mathbb{N}$, any ℓ , any message $m \in \{0,1\}^{\ell}$, any $Q \subseteq [\ell]$, the following holds with probability 1: let $\operatorname{crs} \leftarrow \operatorname{Gen}(1^{\lambda})$, (digest, aux) $\leftarrow \operatorname{Digest}(\operatorname{crs}, m)$, $\pi \leftarrow \operatorname{Open}(\operatorname{crs}, \operatorname{aux}, Q)$, then it holds that $\operatorname{Vf}(\operatorname{crs}, \ell, \operatorname{digest}, Q, m[Q], \pi) = 1$.

Collision resistance. We say that a vector commitment scheme satisfies collision resistance against size- $W(\cdot)$ adversaries, iff for any non-uniform probabilistic machine $\mathcal{A}(1^{\lambda}, *)$ whose total work is bounded by $W(\lambda)$, there exists a negligible function $\mathsf{negl}(\cdot)$ such that for every $\lambda \in \mathbb{N}$, the probability that the following experiment outputs 1 is at most $\mathsf{negl}(\lambda)$:

- crs \leftarrow Gen (1^{λ}) ;
- $(\ell, \text{digest}, \text{ans}, \text{ans}', Q, Q', \pi, \pi') \leftarrow \mathcal{A}(1^{\lambda}, \text{crs}) \text{ where } Q, Q' \subseteq [\ell];$
- Output 1 if $Vf(crs, \ell, digest, Q, ans, \pi) = Vf(crs, \ell, digest, Q', ans', \pi')$; however, there is some $i \in Q \cap Q'$ such that ans and ans' contain different answers for the index *i*.

Merkle [Mer89] showed how to build such a vector commitment scheme secure against polynomially sized adversaries (or quasi-polynomially sized adversaries resp.) assuming the existence of a collision resistant hash family secure against polynomially sized adversaries (or quasi-polynomially sized adversaries).

A.2 Argument of Knowledge

Let \mathcal{L}_{λ} denote an NP language paramtrized by λ . An argument of knowledge system consists of the following algorithms:

• crs \leftarrow Gen $(1^{\lambda}, \mathcal{L}_{\lambda})$: a randomized algorithm that samples and outputs a common reference string crs.

• $b \leftarrow \langle P(\mathsf{crs}, \mathsf{stmt}, w), V(\mathsf{crs}, \mathsf{stmt}) \rangle$: a pair of randomized interactive algorithms where where P denotes the prover and V denotes the verifier. Both P and V receive the common reference string crs and the statement stmt , and P additionally receives a witness w as input. The two then engage in an interactive protocol at the end of which the verifier outputs $b \in \{0, 1\}$ indicating reject or accept.

Perfect completeness. We say an argument of knowledge scheme satisfies *perfect completeness* if for any $\lambda \in \mathbb{N}$, for any NP language \mathcal{L}_{λ} whose corresponding NP relation is denoted \mathcal{R}_{λ} , for any statement (stmt, w) $\in \mathcal{R}_{\lambda}$, the following holds with probability 1: let crs $\leftarrow \text{Gen}(1^{\lambda}, \mathcal{L}_{\lambda})$, then $\langle P(\text{crs}, \text{stmt}, w), V(\text{crs}, \text{stmt}) \rangle = 1$.

Adaptive knowledge soundness. We say that an argument of knowledge scheme satisfies adaptive knowledge soundness, iff for any NP language \mathcal{L}_{λ} with the corresponding NP relation \mathcal{R}_{λ} , there exists a probabilistic quasi-polynomial-time extractor \mathcal{E} , a quasi-polynomially bounded function $q(\cdot)$, and some negligible function $\mathsf{negl}(\cdot)$, such that for every $\lambda \in \mathbb{N}$, any auxiliary distribution \mathcal{D} , any deterministic P^* that runs in time at most t,

$$\Pr\left[(\mathsf{stmt}, w) \notin \mathcal{R}_{\lambda} \land b = 1 \middle| \begin{array}{c} \mathsf{crs} \leftarrow \mathsf{Gen}(1^{\lambda}, \mathcal{L}_{\lambda}) \\ \eta \leftarrow \mathcal{D} \\ (\mathsf{stmt}, \mathsf{aux}) \leftarrow P^{*}(\mathsf{crs}, \eta) \\ b \xleftarrow{tr} \langle P^{*}(\mathsf{aux}), V(\mathsf{crs}, \mathsf{stmt}) \rangle \\ w \leftarrow \mathcal{E}^{P^{*}(\mathsf{aux})}(\mathsf{crs}, \mathsf{stmt}, tr) \end{array} \right] \leq \mathsf{negl}(\lambda)$$

and moreover, \mathcal{E} 's running time is upper bounded by $q(\lambda, t)$. In the above, $b \stackrel{tr}{\leftarrow} \langle P^*(\mathsf{aux}), V(\mathsf{crs}, \mathsf{stmt}) \rangle$ means that tr is the transcript of the execution $\langle P^*(\mathsf{aux}), V(\mathsf{crs}, \mathsf{stmt}) \rangle$ which includes the crs and messages exchanged between P^* and V, and the notation \mathcal{E}^{P^*} means that \mathcal{E} has blackbox access to each next-message function of P^* .

Succinctness. An argument of knowledge system is said to be succinct iff the verifier's runtime is upper bounded by $poly(\lambda, \log |\mathcal{R}_{\lambda}|)$ where $|\mathcal{R}_{\lambda}|$ denotes the size of the circuit that checks the NP relation.

Kilian [Kil92] showed how to construct a succinct argument of knowledge scheme that satisfies the aforementioned adaptive knowledge soundness notion as well as succinctness requirement, assuming the existence of a collision resistant hash family secure against quasi-polynomially sized adversaries.

A.3 Publicly Verifiable Non-Malleable Timed Commitments

We use the definitions from Chvojka and Jager [CJ23]. A publicly verifiable, non-interactive timed commitment (NITC) with message space \mathcal{M} is a tuple of algorithms (Gen, Com, ComVf, DecVf, FDec, FDecVf) with the following syntax:

- $\operatorname{crs} \leftarrow \operatorname{Gen}(1^{\lambda}, T)$: a probabilistic algorithm that takes as input the security parameter 1^{λ} , a difficulty parameter T, and output a common reference string crs.
- $(\mathsf{cm}, \pi_{\mathrm{com}}, \pi_{\mathrm{dec}}) \leftarrow \mathsf{Com}(\mathsf{crs}, m)$: a probabilistic algorithm that takes as input a common reference string crs and a message $m \in \mathcal{M}$, and outputs a commitment cm along with proofs π_{com} and π_{dec} .

- $0/1 \leftarrow \text{ComVf}(\text{crs}, \text{cm}, \pi_{\text{com}})$: a deterministic algorithm that takes as input a common reference string crs, a commitment cm and a proof π_{com} , and outputs a bit indicating whether to reject or accept.
- 0/1 ← DecVf(crs, cm, m, π_{dec}): a deterministic algorithm that takes as input a common reference string crs, a commitment cm, a message m, and a proof π_{dec}, and outputs a bit indicating whether to reject or accept.
- $(m, \pi_{\text{fdec}}) \leftarrow \mathsf{FDec}(\mathsf{crs}, \mathsf{cm}, \pi_{\text{com}})$: a deterministic algorithm that takes as input a common reference string crs , a commitment cm , and outputs $m \in \mathcal{M} \cup \{\bot\}$ in time at most $T \cdot \mathsf{poly}(\lambda)$.
- 0/1 ← FDecVf(crs, cm, m, π_{fdec}): a deterministic algorithm that takes as input a common reference string crs, a commitment cm, a message m, and a proof π_{fdec}, and outputs a bit indicating whether to reject or accept.

Note that when the NITC is used in an actual protocol, if the committer is adversarial, it may not provide π_{dec} in the opening phase. This is why we need FDec to introduce a forced opening proof π_{fdec} and have a separate FDecVf algorithm for the public verifiability of forced openings.

Soundness of forced decryption. We say that an NITC scheme satisfies soundness of forced decryption, iff for any non-uniform PPT adversary \mathcal{A} , for any polynomial function $T(\cdot)$, there exists a negligible function $negl(\cdot)$ such that the probability that the following experiment outputs 1 is at most $negl(\lambda)$:

- crs \leftarrow Gen $(1^{\lambda}, T)$;
- $(\mathsf{cm}, \pi_{\mathrm{com}}) \leftarrow \mathcal{A}(1^{\lambda}, \mathsf{crs});$
- $(m, \pi_{\text{fdec}}) \leftarrow \mathsf{FDec}(\mathsf{crs}, \mathsf{cm}, \pi_{\text{com}});$
- ComVf(crs, cm, $\pi_{com} = 1$ but FDecVf(crs, cm, $m, \pi_{fdec}) \neq 1$.

In other words, except with negligible probability, if an adversarially produced commitment has a valid commitment proof, then FDec should be able to extract a message along with a valid forced-decryption proof.

Correctness. We say that an NITC scheme is correct if for all $\lambda, T \in \mathbb{N}$, and all $m \in \mathcal{M}$, the following holds with probability 1: let $\operatorname{crs} \leftarrow \operatorname{Gen}(1^{\lambda}, T)$, $(\operatorname{cm}, \pi_{\operatorname{com}}, \pi_{\operatorname{dec}}) \leftarrow \operatorname{Com}(\operatorname{crs}, m)$, $(m', \pi_{\operatorname{fdec}}) \leftarrow \operatorname{FDec}(\operatorname{crs}, \operatorname{cm})$, then it holds that $\operatorname{ComVf}(\operatorname{crs}, \operatorname{cm}, \pi_{\operatorname{com}}) = 1$, $\operatorname{DecVf}(\operatorname{crs}, \operatorname{cm}, m, \pi_{\operatorname{dec}}) = 1$, m' = m, and $\operatorname{FDecVf}(\operatorname{crs}, \operatorname{cm}, m, \pi_{\operatorname{dec}}) = 1$.

Definition A.1 (IND-CCA against depth-bounded adversaries). Given some function $W(\cdot)$ and some $\epsilon \in (0, 1)$, we say that an NITC scheme satisfies $(W(\cdot), \epsilon)$ -IND-CCA (short for indistinguishability under chosen-ciphertext-attack) security iff there exists a polynomial $T_{\emptyset}(\cdot)$, such that for all polynomials $T(\cdot) \geq T_{\emptyset}(\cdot)$, and every non-uniform adversary $\mathcal{A} = (\mathcal{A}_1, \mathcal{A}_2)$ where $\mathcal{A}(1^{\lambda}, \cdot)$ is bounded by $W(\lambda)$ in total work and $\mathcal{A}_2(1^{\lambda}, \cdot)$ is bounded $T^{\epsilon}(\lambda)$ in depth, there exists a negligible function $\mathsf{negl}(\cdot)$ such that for all $\lambda \in \mathbb{N}$, it holds that $|\Pr[\mathsf{Expt}_0^{\mathcal{A}}(1^{\lambda}) = 1] - \Pr[\mathsf{Expt}_1^{\mathcal{A}}(1^{\lambda}) = 1] \leq \mathsf{negl}(\lambda)$ where $\mathsf{Expt}_b^{\mathcal{A}}(1^{\lambda})$ is defined as follows:

- crs \leftarrow Gen $(1^{\lambda}, T(\lambda));$
- $(m_0, m_1, \mathsf{st}) \leftarrow \mathcal{A}_{1,\lambda}^{\mathsf{Dec}(\cdot, \cdot)}(\mathsf{crs});$
- $(\operatorname{cm}^*, \pi_{\operatorname{com}}^*, \pi_{\operatorname{dec}}^*) \leftarrow \operatorname{Com}(\operatorname{crs}, m_b);$

- $b' \leftarrow \mathcal{A}_2^{\mathsf{Dec}(\cdot,\cdot)}(\mathsf{cm}^*, \pi^*_{\mathrm{com}}, \mathsf{st});$
- output b'.

where the oracle $\mathsf{Dec}(\mathsf{cm}, \pi_{\mathrm{com}})$ returns the result of $\mathsf{FDec}(\mathsf{crs}, \mathsf{cm})$ if $\mathsf{ComVf}(\mathsf{crs}, \mathsf{cm}, \pi_{\mathrm{com}}) = 1$; otherwise it returns \bot ; and further, \mathcal{A} must respect the following restrictions: $|m_0| = |m_1|$, and moreover, it is not allowed to query the oracle $\mathsf{Dec}(\cdot, \cdot)$ on $(\mathsf{cm}^*, \pi^*_{\mathrm{com}})$.

Definition A.2 (Computationally binding). We say that an NITC scheme is computationally binding for size- $W(\cdot)$ adversaries, iff for any non-uniform PPT adversary $\mathcal{A}(1^{\lambda}, \cdot)$ whose total work is bounded by $W(\lambda)$, there exists a negligible function $negl(\cdot)$ such that for every λ , the probability that the following experiment outputs 1 is at most $negl(\lambda)$:

- crs \leftarrow Gen $(1^{\lambda}, T(\lambda));$
- $(m, \operatorname{cm}, \pi_{\operatorname{com}}, \pi) \leftarrow \mathcal{A}(\operatorname{crs});$
- output 1 iff all of the following hold:
 - ComVf(crs, cm, π_{com}) = 1,
 - $\operatorname{DecVf}(\operatorname{crs}, \operatorname{cm}, m, \pi) = 1, \text{ or } \operatorname{FDecVf}(\operatorname{crs}, \operatorname{cm}, m, \pi) = 1,$
 - but the message output by FDec(crs, cm) is not equal to m.

Chvojka and Jager [CJ23] constructed an NITC scheme satisfying the aforementioned requirements, assuming that the following assumptions hold for quasi-polynomially sized adversaries: the strong sequential squaring assumption [MT19, KLX20], the Decisional Diffie-Hellman (DDH) and the Decisional Composite Residuosity (DCR) assumptions in suitable groups, and the existence of a simulation sound extractable Non-Interactive Zero Knowledge (NIZK) system (which is implied by the hardness of DDH against quasi-polynomially sized adversaries).

B Deferred Proofs for the Ascending Auction

We now prove Theorem 5.1.

Proof. We prove each property individually.

bIC. Fix an arbitrary buyer *i* with true value v_i , and let τ be the number that the honest platform posts on the blockchain in the honest execution. It is easy to see that when the buyer *i* underbids (sending \perp when $v_i > \theta_{\tau}$) or overbids (keep sending ok even when $v_i < \theta_{\tau}$), it cannot increase its utility compared to behaving honestly. Moreover, injecting fake bids either does not affect or increases the payment a winning buyer needs to pay. Thus, behaving honestly is the best response for a buyer.

pIC. Since τ is posted on the blockchain, the platform only gets θ_{τ} from each winner and has to pay θ_{τ} to the seller for each item sold in any safe execution trace. Because \tilde{k} must be at least the number of buyers who win items in a safe execution trace, the platform's utility is at most 0.

1-pbIC. As we argued for pIC, the coalition cannot extract any profit from buyers outside the coalition, so the only way the coalition can increase their utility is to lower the price that the colluding buyer i has to pay, or to bias the random tie-breaking. However, if the coalition lowers the price (i.e. stop ascending the price earlier), there are more than k buyers who should win items. In this case, either buyer i has to give up winning an item which leads to zero utility, or the platform has to sell items to more than k buyers where the execution is not safe. On the other hand, if the colluding buyer is chosen from \mathcal{I} at the final round τ , it must be that buyer i's utility is zero no matter it wins an item or not. Thus, the auction satisfies 1-pbIC.

Bayesian psIC. Let $\tau \in {\tau_0, \ldots, T}$ denote the number that the platform would post on the blockchain in the honest execution. Let \mathcal{R}_v be the set of buyers whose true values are at least v. If the platform stops the auction early and posts some $\tau^* < \tau$ on the blockchain, then by definition, it must be that $|\mathcal{R}_{\theta_{\tau^*}}| > k$. The resulting execution traces cannot be safe. Therefore, we focus on the case where $\tau^* \geq \tau$.

If the platform-seller coalition follows the protocol, the revenue they gain from each winner is θ_{τ} . If they keep ascending the price to $\theta_{\tau^*} > \theta_{\tau}$, the coalition can only earn revenue from a buyer if the buyer's true value is still larger than or equal to θ_{τ^*} . Thus, the expected revenue they gain from each buyer i in $\mathcal{R}_{\theta_{\tau^*}}$ is $\theta_{\tau^*} \cdot \Pr_{v_i \stackrel{\$}{\leftarrow} \mathcal{D}} [v_i \geq \theta_{\tau^*} \mid v_i \geq \theta_{\tau}]$. To make this deviation profitable, it

must be

$$\theta_{\tau^*} \cdot \Pr_{\substack{v_i \leftarrow \mathcal{D}}} [v_i \ge \theta_{\tau^*} \mid v_i \ge \theta_{\tau}] > \theta_{\tau}.$$

$$\tag{7}$$

Re-arranging Equation (7), we obtain $\theta_{\tau^*}(1 - F(\theta_{\tau^*})) > \theta_{\tau}(1 - F(\theta_{\tau}))$. Now, consider any *i*, we have

$$\theta_{i+1}(1 - F(\theta_{i+1})) - \theta_i(1 - F(\theta_i)) = (\theta_{i+1} - \theta_i)(1 - F(\theta_i)) - \theta_{i+1}f(\theta_i)$$

$$\leq (\theta_{i+1} - \theta_i)(1 - F(\theta_i)) - \theta_i f(\theta_i)$$

$$= -\phi(\theta_i) \cdot f(\theta_i).$$

For any $i \ge \tau_0$, we have $\phi(\theta_i) \ge 0$. Thus, $\theta(1 - F(\theta))$ is a non-increasing function for any $\theta \ge \theta_{\tau_0}$. Because $\theta_{\tau^*} > \theta_{\tau} \ge \theta_{\tau_0}$, we have $\theta_{\tau^*}(1 - F(\theta_{\tau^*})) \le \theta_{\tau}(1 - F(\theta_{\tau}))$, which implies keeping ascending the price above θ_{τ} is not profitable in expectation. Thus, the auction satisfies Bayesian psIC.

Bayesian sIC. If there exists a scenario where the strategic seller can profit with some strategy S, the seller colludes with the platform, and the seller adopts strategy S while the platform behaves honestly. Because the platform revenue is always zero, the joint utility of the coalition increases since the seller itself can increase its utility, and it violates Bayesian psIC. Thus, the auction must satisfy Bayesian sIC.

Approximate revenue maximization. In Theorem 9.1, we showed that the ideal-world secondprice satisfies revenue optimality. For any fixed bid vector **b**, it is easy to see that the total number of items allocated is the same in the ascending auction and the second price auction of Section 9. Further, each confirmed buyer pays at most tick(\mathbb{U}) less in the ascending auction than the maximum pay of any confirmed buyer in the second price auction. Therefore, $(k \cdot \text{tick}(\mathbb{U}))$ -approximate revenue optimality directly follows.