

Credible Mechanisms*

Mohammad Akbarpour[†] Shengwu Li[‡]

March 27, 2018

Abstract

Consider an extensive-form mechanism, run by an auctioneer who communicates sequentially and privately with agents. Suppose the auctioneer can make any deviation that no single agent can detect. We study the mechanisms such that it is incentive-compatible for the auctioneer not to deviate – the *credible* mechanisms. Consider the optimal auctions in which only winners make transfers. The first-price auction is the unique credible static mechanism. The ascending auction is the unique credible strategy-proof mechanism.

Keywords: Mechanism design, auction, credible, strategy-proof, sealed-bid.

JEL classification codes: D44, D47, D82

*This paper has been previously distributed with the title “Credible Mechanism Design”. We thank Ben Brooks, Jeremy Bulow, Alex Frankel, Ben Golub, Matthew Jackson, Scott Kominers, David Kreps, Elliot Lipnowski, Erik Madsen, Eric Maskin, Paul Milgrom, Roger Myerson, Michael Ostrovsky, Doron Ravid, Ilya Segal, Andy Skrzypacz, Cameron Taylor, Alex Teytelboym, Bob Wilson, Alex Wolitzky, and several seminar participants for valuable comments and advice. All errors remain our own.

[†]Graduate School of Business, Stanford University. Email: mohamwad@stanford.edu

[‡]Society of Fellows, Harvard University. Email: shengwu_li@fas.harvard.edu

1 Introduction

Auctions are used to sell a wide variety of goods - art, fish, real estate, treasury bonds, internet advertising, wireless spectrum, and, in 193 AD, the entire Roman Empire (Shubik, 2004). The theory of optimal auctions pins down only a few features of the auction format. Under the standard assumptions, if bidders with higher values submit higher bids, and the reserve price is optimally chosen, then the auction is optimal (Myerson, 1981). Despite this, most real-world auctions are variations on just a few canonical formats - the first-price auction, the ascending auction, and (more recently) the second-price auction (Cassady, 1967; McAfee and McMillan, 1987).¹ Why?

In this paper, we shed light on this question by proposing a framework for mechanism design under partial commitment. The standard mechanism design paradigm assumes full commitment on behalf of the auctioneer:

[In the standard paradigm] it is assumed that the organizer of the auction has the ability to commit himself in advance to a set of policies. He binds himself in such a way that all the bidders know that he cannot change his procedures after observing the bids, even though it might be in his interest ex post to renege. In other words, the organizer of the auction acts as the Stackelberg leader or first mover. (McAfee and McMillan, 1987)

In real-world auctions, however, the auctioneer may have various opportunities to deviate that are difficult for bidders to detect. The auctioneer could secretly inspect and tamper with sealed bids. She could hire confederates to pose as bidders to manipulate the outcome. She could call out nonexistent bids in order to give the impression of greater demand - a practice known as “chandelier bidding”:

Under New York City regulations auctioneers can fabricate bids up to an item’s reserve price. Because a reserve price is secret and not listed in the catalog, bidders have no way of knowing which offers are real. (The New York Times, April 24 2000)

Which deviations are feasible plainly depends on the auction format. Instead of considering deviations in an *ad hoc* manner, it would be useful to model the feasible deviations as a function of the format. This paper provides one systematic framework to do so.

Consider any *protocol*; a pair consisting of an extensive-form mechanism and a strategy profile for the agents. We can think of the auctioneer *running* the mechanism by engaging in a sequence of private communication with the agents. Upon encountering an information set, she picks up the telephone and conveys that information to the agent who

¹The Dutch (descending) auction, in which the price falls until one bidder claims the object, is less prevalent (Krishna, 2010, p.2).

is called to play, along with a set of acceptable replies (actions). The agent then chooses a reply. The auctioneer keeps making telephone calls, sending information and receiving replies, until she reaches a terminal history, whereupon she chooses the corresponding outcome and the game ends.

Suppose some utility function for the auctioneer. For instance, assume that the auctioneer wants revenue. Suppose that each agent intrinsically observes certain features of the outcome. For instance, each agent observes whether or not he wins the object, and how much he pays, but not how much other agents pay.

By participating in the protocol, each agent observes a sequence of communication between himself and the auctioneer and some features of the outcome. Even if the auctioneer deviates from her assigned strategy, agent i 's observation could still have an *innocent explanation*. That is, when the auctioneer plays by the rules, there exist types for the other agents that result in that same observation for i .

For any given protocol, there may be some deviations that are *safe*, in the sense that for every type profile, each agent's observation has an innocent explanation. That is, every observation that an agent might have (under the deviation) is also an observation he might have when the auctioneer is running the mechanism. For instance, when a bidder bids \$100 in a second-price auction, receives the object, and is charged \$99, that observation has an innocent explanation - it could be that the second-highest value was \$99. Thus, in a second-price auction, the auctioneer can safely deviate by exaggerating the second-highest bid.²

Instead of just choosing a different outcome, the auctioneer may also alter the way she communicates with agents. For example, consider a protocol in which the auctioneer acts as a middleman between one seller and one buyer. The seller chooses a price for the object, which the auctioneer tells to the buyer. The object is sold to the buyer at that price if and only if the buyer accepts, and the auctioneer takes a 10% commission. The auctioneer has a safe deviation - she can quote a higher price to the buyer, and pocket the difference if the buyer accepts.

A protocol is *credible* if running the mechanism is incentive-compatible for the auctioneer; that is, if the auctioneer prefers playing by the book to any safe deviation. This is a way to think about partial commitment power for any extensive-form mechanism. For instance, if the auctioneer wants revenue, then the second-price auction is not credible.

Our model assumes that the auctioneer is the nexus of communication, exchanging private messages with individual agents. Sometimes, this is literally the case: Large auction houses, such as Sotheby's and Christie's, convey information and solicit bids over multiple

²Such deviations are not purely hypothetical. An auctioneer running second-price auctions in Connecticut admitted, "After some time in the business, I ran an auction with some high mail bids from an elderly gentleman who'd been a good customer of ours and obviously trusted us. My wife Melissa, who ran the business with me, stormed into my office the day after the sale, upset that I'd used his full bid on every lot, even when it was considerably higher than the second-highest bid." (Lucking-Reiley, 2000)

telephone lines and the Internet; most of Christie’s auctions are won by bidders who are not in the room (Grant, 2014). Sometimes, the law forbids bidders from communicating with each other, as in the 2017 US Federal Communications Commission Incentive Auction for wireless spectrum.³ Of course, many auctions involve other communication procedures. We see our definition of a “safe deviation” as a tractable shorthand for a rich set of real-world deviations, such as tampering with sealed bids, altering the relative speed of price clocks, or hiring confederates to influence the outcome.

Having defined the framework, we now turn to our main application. Consider three standard auction formats. The first-price auction is a static (“sealed-bid”) – so each agent is called to play exactly once, and has no information about the history of play when selecting his action. This yields a substantial advantage: Sealed-bid auctions can be conducted rapidly and asynchronously, thus saving logistical costs.⁴ The ascending auction is strategy-proof. Thus, it demands less strategic sophistication from bidders, and does not depend sensitively on bidders’ beliefs (Wilson, 1987; Bergemann and Morris, 2005; Chung and Ely, 2007). The second-price auction is static and strategy-proof; it combines the virtues of the first-price auction and the ascending auction (Vickrey, 1961). However, many real-world auctioneers persist in running first-price auctions and ascending auctions, despite the invention of this (apparently) superior format (Rothkopf et al., 1990). Why is this so?

Credibility illuminates the relationship between these three standard auctions. Suppose the classic assumptions: Private values, with symmetric independent regular distributions (Myerson, 1981). We restrict attention to auctions in which only winning bidders make (or receive) transfers. Thus, the second-price auction (with reserve) is the unique strategy-proof static optimal auction, by the Green-Laffont-Holmström theorem (Green and Laffont, 1977; Holmström, 1979).

The results that follow require us to bridge the discrete world of extensive game forms and the continuous world of optimal auctions. We suppress these technicalities in the introduction, but the reader should be aware that the first result holds ‘in the limit’, as a finite grid type space becomes arbitrarily fine.

Our first result is as follows: The first-price auction (with reserve) is the unique credible static optimal auction. This implies that, in the class of “sealed-bid” mechanisms, we must choose between incentive-compatibility for the auctioneer and dominant strategies for the agents.

Static mechanisms include the direct revelation mechanisms, in which each agent simply reports his type. Thus, when designing credible protocols, restricting attention to

³The Incentive Auction’s rules state that bidders “are prohibited from communicating directly or indirectly any incentive auction applicant’s bids or bidding strategies” (Section 1.2205(b)).

⁴Using data from U.S. Forest Service timber auctions, Athey et al. (2011) find that “sealed bid auctions attract more small bidders, shift the allocation toward these bidders, and can also generate higher revenue”.




Figure 1: An auction trilemma: In the class of optimal auctions in which only the winner makes transfers, no auction is static, strategy-proof, and credible. Picking two out of three properties uniquely characterizes each standard format.

revelation mechanisms loses generality. The problem is that revelation mechanisms reveal too much, too soon. For a bidder to have a dominant strategy, his payment must depend on the other bidders' types. If the auctioneer knows the entire type profile, and the winning bidder's payment depends on the other bidders' types, then the auctioneer can safely deviate to raise revenue. This makes it impossible to run a credible strategy-proof optimal auction. What happens when we look outside the class of revelation mechanisms - when we use the full richness of extensive forms to regulate who knows what, and when?

Our second result is as follows: The ascending auction (with an optimal reserve) is credible. Moreover, it is the *unique* credible strategy-proof optimal auction. No other extensive forms satisfy these criteria.

These results imply an auction trilemma. Static, strategy-proof, or credible: An optimal auction can have any two of these properties, but not all three at once. Moreover, picking two out of three characterizes each of the standard auction formats (first-price, second-price, and ascending). Figure 1 illustrates.

1.1 Related work

In the literature on mechanism design with an informed principal, the principal has exogenous private information and cannot commit to how she uses it. Thus, the mechanism must give the principal incentives to report that information to a neutral mediator (Myerson, 1983; Maskin and Tirole, 1990, 1992). In this paper, we study a case of endogenous private information. Our principal is the mediator - she learns, in the course of running the mechanism, private information about each agent's type, and needs incentives to follow the rules.

There has been substantial prior work on auction design under limited commitment. The main approaches are as follows: Some papers consider models in which an auctioneer runs auctions over multiple periods, but cannot commit today to the mechanism used tomorrow (Milgrom, 1987; McAfee and Vincent, 1997; Caillaud and Mezzetti, 2004; Bester and Strausz, 2001; Liu et al., 2014; Skreta, 2015). (This approach often relies on equilibrium refinements to yield interesting results.) Other papers consider bargaining games in which the auctioneer cannot commit to end the auction - she can continue to make or solicit offers in defiance of the rules (Burguet and Sakovics, 1996; McAdams and Schwarz, 2007; Vartiainen, 2013; Lobel and Paes Leme, 2017). Baliga et al. (1997) consider a setting where agents make simultaneous moves and they know the state, while the planner cannot commit to the outcome and can change it after observing agents' actions. Some papers consider mechanisms in which each agent reports his type to the principal, and the principal's contract with each agent can only directly depend on that bidder's report (Rothkopf and Harstad, 1995; Porter and Shoham, 2005; Dequiedt and Martimort, 2015). In addition, seller's "shill bidding" has been studied in some papers (Chakraborty and Kosmopoulou, 2004; Lamy, 2009). Finally, Loertscher and Marx (2017) consider a clock auction in which the auctioneer chooses whether to keep the clocks running or to end the auction.

Our current approach is distinct in a few ways. Firstly, we consider a non-repeated interaction, and thus do not rely on equilibrium refinements. Secondly, our auctioneer's commitment power depends on what agents can observe, not on the passage of calendar time. We thus rule out flagrant rule-breaking - for instance, the auctioneer cannot make new offers to a bidder who knows that the auction should already be over. Thirdly, instead of assuming that each bidder makes a single report, or assuming that every bidder makes one offer in every period, we study the class of extensive game forms with perfect recall. We show that credibility is a defining characteristic of the first-price auction and the ascending auction. As far as we know, no prior work has produced a desirable criterion that selects these two formats.

Li (2017) introduces a definition of bilateral commitment power, to characterize the choice rules that are implementable in obviously dominant strategies. Our paper builds on that formalism, with two key differences. Firstly, the definition in Li (2017) is only for dominant-strategy mechanisms, whereas credibility allows for Bayes-Nash mechanisms. Secondly, obvious dominance is a stronger incentive requirement for the agents, but is silent on the incentives of the auctioneer. Credibility imposes no novel restrictions on the incentives of agents. Instead, it explicitly requires the mechanism to be incentive-compatible for the auctioneer.

2 The Model

2.1 Definitions

The environment consists of:

1. A finite set of agents, N .
2. A set of outcomes, X .
3. A finite type space, $\Theta_N = \times_{i \in N} \Theta_i$.
4. A joint probability distribution $D : \Theta_N \rightarrow [0, 1]$.
5. Agent utilities $u_i : X \times \Theta_N \rightarrow \mathbb{R}$
6. A partition Ω_i of X for each $i \in N$. (ω_i denotes a cell of Ω_i .)

The partition Ω_i represents what agent i *directly observes* about the outcome. Conceptually, these partitions represent physical facts about the world, which are not objects of design. They capture the bare minimum that each agent observes about the outcome, regardless of the choice of mechanism.⁵

A mechanism is an extensive game form with consequences in X . This is an extensive game form for which each terminal history is associated with some outcome. Formally, a mechanism G is a tuple $(H, \prec, P, A, \mathcal{A}, (\mathcal{I}_i)_{i \in N}, g)$, where each part of the tuple is as specified in Table 1. The full definition of extensive forms is familiar to most readers, so we relegate further detail to Appendix A. Let \mathcal{G} denote the set of all extensive game forms with consequences in X with finitely many histories and perfect recall.

Table 1: Notation for Extensive Game Forms

| Name | Notation | Representative element |
|------------------------------------|------------------|------------------------|
| histories | H | h |
| precedence relation over histories | \prec | |
| initial history | h_\emptyset | |
| terminal histories | Z | z |
| player called to play at h | $P(h)$ | |
| actions | A | a |
| most recent action at h | $\mathcal{A}(h)$ | |
| information sets for agent i | \mathcal{I}_i | I_i |
| outcome resulting from z | $g(z)$ | |
| immediate successors of h | $\sigma(h)$ | |
| actions available at I_i | $A(I_i)$ | |

⁵In the application that follows, we will assume that each bidder in an auction knows how much he paid and whether he receives the object. In effect, this rules out the possibility that the auctioneer could hire pickpockets to raise revenue, or sell the object to multiple bidders by producing counterfeit copies.

S_i denotes a (pure) strategy: For each information set where agent i is called to play and each type of i , S_i chooses an action $S_i(I_i, \theta_i) \in A(I_i)$. $(S_i)_{i \in N} \equiv S_N$ denotes a strategy profile for the agents, for some $G \in \mathcal{G}$. We refer to a pair (G, S_N) as a *protocol*. Let $x^G(S_N, \theta_N)$ denote the outcome in G , when agents play according to S_N and the type profile is θ_N . Let $u_i^G(S_N, \theta_N) \equiv u_i(x^G(S_N, \theta_N), \theta_N)$.

Definition 1. (G, S_N) is **Bayesian incentive-compatible (BIC)** if, for all $i \in N$,

$$S_i \in \operatorname{argmax}_{S'_i} \mathbb{E}_{\theta_N} [u_i^G(S'_i, S_{-i}, \theta_N)] \quad (1)$$

We explicitly model the auctioneer⁶ as a player (denoted 0). The auctioneer has utility $u_0 : X \times \Theta_N \rightarrow \mathbb{R}$.

2.2 Pruning

We restrict attention to the class of *pruned* protocols.⁷ This technique allows us to remove redundant parts of the game tree, and implies cleaner definitions for the theorems that follow. In words, a pruned protocol has three properties.

1. For every history h , there exists some type profile such that h is on the path of play.
2. At every information set, there are at least two actions available (equivalently, every non-terminal history has at least two immediate successors).
3. If agent i is called to play at history h , then there are two types of i compatible with his actions so far, that could lead to different eventual outcomes.

Let $z(S_N, \theta_N)$ denote the terminal history that results from (S_N, θ_N) . Formally,

Definition 2. (G, S_N) is **pruned** if, for any history h :

1. There exists θ_N such that $h \preceq z(S_N, \theta_N)$
2. If $h \notin Z$, then $|\sigma(h)| \geq 2$.
3. If $h \notin Z$, then for $i = P(h)$, there exist $\theta_i, \theta'_i, \theta_{-i}$ such that

$$(a) \ h \prec z(S_N, (\theta_i, \theta_{-i}))$$

$$(b) \ h \prec z(S_N, (\theta'_i, \theta_{-i}))$$

$$(c) \ x^G(S_N, (\theta_i, \theta_{-i})) \neq x^G(S_N, (\theta'_i, \theta_{-i}))$$

⁶We use the term ‘auctioneer’ to refer to the mediator, but this could be any mediator who runs a mechanism, such as a school choice authority or the National Resident Matching Program.

⁷This is stronger than the notion of pruning used in Li (2017), which includes only the first requirement.

By the next proposition, when our concern is BIC implementation, it is without loss of generality to consider only pruned protocols.

Proposition 1. *If (G, S_N) is BIC, then there exists (G', S'_N) such that (G', S'_N) is pruned and BIC and for all θ_N , $x^{G'}(S'_N, \theta_N) = x^G(S_N, \theta_N)$.*

Hence, from this point onwards we restrict attention to pruned (G, S_N) . Since every information set is reached with positive probability, any Bayes-Nash equilibrium in a pruned protocol survives equilibrium refinements that restrict off-path beliefs.

2.3 A messaging game

Our goal is to study surreptitious deviations by the auctioneer. To do so, we must first take a stand on what each player knows at each point in the mechanism. Consider the messaging game, defined thus:

1. The auctioneer chooses to:
 - (a) Either: Select $x \in X$ and end the game.
 - (b) Or: Go to step 2.
2. The auctioneer chooses some agent $i \in N$ and sends a message along with a set of acceptable replies (m, R) .
3. Agent i privately observes (m, R) and chooses $r \in R$.
4. The auctioneer privately observes r .
5. Go to step 1.

Assume the auctioneer has an arbitrarily rich message space M . At any prior round k , the auctioneer messaged $i^k \in N$ with (m^k, R^k) , and received reply $r^k \in R^k$.

Let \mathcal{S}_0 denote the set of auctioneer pure strategies. A strategy for the auctioneer specifies what to do next, as a function of the entire history of communications: $S_0((i^k, m^k, R^k, r^k)_{k=1}^t) \in (N \times M \times (2^M \setminus \{\emptyset\})) \cup X$. We restrict these to send finitely many messages and (for each message) to allow finitely many replies.

A strategy for agent i specifies what reply to give, as a function of the previous communications between agent i and the auctioneer, and the agent's type. Let m_i^k denote the k th message that the auctioneer sent to agent i , and similarly for R_i^k and r_i^k . Let t_i denote the total number of messages sent to agent i so far. $S_i((m_i^k, R_i^k, r_i^k)_{k=1}^{t_i-1}, m^{t_i}, R^{t_i}, \theta_i) \in R^{t_i}$.⁸

⁸Note the lack of calendar time: The agent observes the sequence of past communications between himself and the auctioneer, *not* a sequence of periods in which he either sees some communication or none.

For any $S_0 \in \mathcal{S}_0$ and any S_N , (S_0, S_N) results in some sequence of communication between the auctioneer and agent i , $o_i^c \equiv (m_i^k, R_i^k, r_i^k)_{k=1}^{T_i}$ and some outcome x . Let o_i^x denote $\omega_i \in \Omega_i \mid x \in \omega_i$. $o_i = (o_i^c, o_i^x)$ denotes an *observation* for agent i . $\phi_i(S_0, S_N, \theta_N)$ denotes the unique observation resulting from (S_0, S_N) , when the type profile is θ_N .

Definition 3. Take any $G = \langle H, \prec, P, A, \mathcal{A}, (\mathcal{I}_i)_{i \in N}, g \rangle$. S_0 **runs** G if there exists a one-to-one function $\lambda : (\cup_{i \in N} \mathcal{I}_i) \cup A \rightarrow M$ such that S_0 is described by the following algorithm, where we initialize $h := h_\emptyset$.

1. If $h \in Z$, terminate and select $x = g(h)$.
2. Else:
 - (a) Choose agent $P(h)$ and send $(m, R) = \lambda(I_i, A(I_i))$ for I_i such that $h \in I_i$.
 - (b) Upon receiving $r \in R$, choose h' such that $\mathcal{A}(h') = \lambda^{-1}(r)$ and $h' \in \sigma(h)$. Set $h := h'$ and go to step 1.

We use S_0^G to denote an auctioneer strategy that runs G .

Given S_0^G , for any S_i , we can define an equivalent strategy \tilde{S}_i for agent i in the messaging game. Given (m, R) and θ_i , \tilde{S}_i selects reply $\lambda(S_i(\lambda^{-1}(m), \theta_i))$.⁹ We abuse notation and use S_i to denote both a strategy for agent i in G , and the equivalent strategy for agent i in the messaging game.

2.4 Credible mechanisms

In formulating the idea of incentive compatibility, [Hurwicz \(1972\)](#) writes:

In effect, our concept of incentive compatibility merely requires that no one should find it profitable to “cheat,” where cheating is defined as behavior that can be made to look “legal” by a misrepresentation of a participant’s preferences or endowment, with the proviso that the fictitious preferences should be within certain “plausible” limits.

Here we modify Hurwicz’s seminal idea to include incentive compatibility for the auctioneer. First, we say that some observation o_i has an “innocent explanation” (with respect to S_0) if there exist some types of other agents such that if they had those types and the auctioneer played S_0 , o_i was the observation of agent i .

Definition 4. Suppose the auctioneer promises to play S_0 , the agents play S_N , and the type profile is θ_N . i ’s observation o_i has an **innocent explanation** if there exists θ'_{-i} such that $o_i = \phi_i(S_0, S_N, (\theta_i, \theta'_{-i}))$.

⁹We specify \tilde{S}_i arbitrarily for communication sequences that are never observed under S_0^G .




Figure 2: A mechanism and a deviation. If agent 1 cannot distinguish outcomes a and b , then the deviation is safe.

Definition 5. Suppose the auctioneer promises to play S_0 and the agents play S_N . Then, an auctioneer strategy $S'_0 \in \mathcal{S}_0$ is **safe** if for all agents $i \in N$ and all type profiles $\theta_N \in \Theta_N$, $\phi_i(S'_0, S_N, \theta_N)$ has an innocent explanation.

Let $\mathcal{S}_0^*(S_0, S_N) \equiv \{S'_0 \mid S'_0 \text{ is safe given promise } S_0 \text{ and agent strategies } S_N\}$. The function $\mathcal{S}_0^*(\cdot, \cdot)$ takes a strategy for the auctioneer and a strategy profile for the agents as inputs and outputs the set of all strategies that the auctioneer can deviate to without being detected by a single agent.

Let $u_0(S_0, S_N, \theta_N)$ denote the auctioneer's utility in the messaging game, when the strategy profile is (S_0, S_N) and the type profile is θ_N . A protocol is credible if playing 'by the book' maximizes the auctioneer's expected utility.

Definition 6. (G, S_N) is **credible** if:

$$S_0^G \in \operatorname{argmax}_{S_0 \in \mathcal{S}_0^*(S_0^G, S_N)} \mathbb{E}_{\theta_N}[u_0(S_0, S_N, \theta_N)] \quad (2)$$

Definition 6 permits the auctioneer to misrepresent agents' actions to each other mid-way through the mechanism. The following example illustrates.

Example 1. Consider the mechanism on the left side of Figure 2. Each agent has one information set, two moves (left and right), and two types (l_i and r_i) that play the corresponding moves. By assumption, agent 1 observes whether the outcome is in the set $\{a, b\}$ or in $\{c\}$. Agents 2 and 3 perfectly observe the outcome.

The right side of Figure 2 illustrates a safe deviation: If agent 1 plays left, then the auctioneer plays 'by the book'. If agent 1 plays right, then instead of querying agent 2, the auctioneer queries agent 3. If agent 3 then plays left, the auctioneer chooses outcome a . If agent 3 plays right, only then does the auctioneer query agent 2, choosing c if 2 plays left and b if 2 plays right.

For every type profile, each agent’s observation has an innocent explanation. The most interesting case is when the type profile is (r_1, l_2, l_3) . In this case, playing by the book results in outcome b , but the deviation results in outcome a . Agent 1 cannot distinguish between a and b , so (l_2, l_3) is an innocent explanation for 1. (l_1, l_3) is an innocent explanation for 2, and (l_1, l_2) is an innocent explanation for 3. Thus, if the auctioneer prefers outcome a to any other outcome, then the mechanism is not credible.

Notably, this deviation involves not just choosing different outcomes, but communicating differently even before a terminal history is reached. Indeed, when the type profile is (r_1, l_2, l_3) , the auctioneer can only get outcome a by deviating midway. If she waited until the end and then deviated to choose a , then agent 2’s observation would not have an innocent explanation. Once agent 2 is called to play, he knows that outcome a should not occur.

Definition 6 takes the expectation of θ_N with respect to the *ex ante* distribution D , but it implicitly requires the auctioneer to best-respond to her *updated* beliefs in the course of running G . Recall that a strategy for the auctioneer is a complete contingent plan. Suppose that in the course of running G , the auctioneer discovers new information about agents’ types, such that she can profitably change her continuation strategy. There exists a deviating strategy that adopts this new course of action contingent on the auctioneer discovering this information, and plays by the book otherwise.

When our concern is credible implementation, it is also without loss of generality to consider only pruned mechanisms.

Proposition 2. *If (G, S_N) is credible and BIC, then there exists (G', S'_N) such that (G', S'_N) is pruned, credible, and BIC, and for all θ_N , $x^{G'}(S'_N, \theta_N) = x^G(S_N, \theta_N)$.*

Observation 1. *(G, S_N) is credible and BIC if and only if (S_0^G, S_N) is a Bayes-Nash equilibrium of the messaging game in which the auctioneer is constrained to play strategies in $\mathcal{S}_0^*(S_0^G, S_N)$.*

Credibility restricts attention to ‘promise-keeping’ equilibria of the messaging game. However, any equilibrium can be turned into a promise-keeping equilibrium by altering the promise.

Observation 2. *If $S'_0 \in \mathcal{S}_0^*(S_0, S_N)$, then $\mathcal{S}_0^*(S'_0, S_N) \subseteq \mathcal{S}_0^*(S_0, S_N)$. Thus, if (S'_0, S_N) is a Bayes-Nash equilibrium given promise S_0 , then (S'_0, S_N) is a Bayes-Nash equilibrium given promise S'_0 .*

Credibility is neither weaker nor stronger than the definition of bilateral commitment studied in Li (2017). That definition essentially requires that each agent’s strategy S_i is weakly dominant, even when the auctioneer can deviate to strategies that produce identical observations for i . By contrast, credibility requires S_0^G to be incentive-compatible for the auctioneer, when the auctioneer can make any safe deviation.

3 Credible Optimal Auctions

We consider an auction for a single object. We study the independent private values model of Myerson (1981), but with discrete type spaces. The use of discrete type spaces allows us to bypass the known paradoxes of extensive game forms with continuous time or infinite actions (Simon and Stinchcombe, 1989; Myerson and Reny, 2016).

We restrict attention to mechanisms in which losing bidders neither make nor receive transfers.¹⁰ An outcome $x = (y, t)$ consists of a winner $y \in N \cup \{0\}$ and a payment $t \in \mathbb{R}$ by the winner, so $X = (N \cup \{0\}) \times \mathbb{R}$. Assume there are at least two bidders.

Type spaces are discrete: $\Theta_i = \{\theta_i^0, \theta_i^1, \dots, \theta_i^K\}$. We associate each type with a real number $v(\theta_i^k)$. Assume $v(\theta_i^K) > 0$. For all i, j, k , $v(\theta_i^k) = v(\theta_j^k)$, and $v(\theta_i^{k+1}) - v(\theta_i^k) = \epsilon > 0$. We will abuse notation slightly, and use θ_i^k to refer both to i 's k th type, and to the real number associated with that type. It will be clear from context which is intended. Types are independently and identically distributed, with probability mass function $p : \Theta_i \rightarrow [0, 1]$. Every type has positive probability. Let $F(\theta^k) \equiv \sum_{l=0}^k p(\theta^l)$ and $f(\theta^k) \equiv \frac{p(\theta^k)}{\epsilon}$.

Agents have private values, that is:

$$u_i((y, t), \theta_N) = 1_{i=y}(\theta_i - t) \tag{3}$$

Ω_i is as follows: Each bidder observes whether he wins the object and observes his own payment. That is, $(y, t), (y', t') \in \omega_i$ if and only if:

1. $y \neq i$ and $y' \neq i$
2. or $y = y' = i$ and $t = t'$.

The auctioneer desires revenue¹¹:

$$u_0((y, t), \theta_N) = 1_{y \in N} t \tag{4}$$

Let $\pi(G, S_N)$ denote the expected revenue of (G, S_N) .

The virtual values machinery in Myerson (1981) applies *mutatis mutandis* to the discrete setting. Suppose we choose (G, S_N) to maximize expected revenue subject to incentive compatibility and (interim) individual rationality. That is $\max_{G, S_N} \pi(G, S_N)$ subject to:

1. (G, S_N) is BIC.

¹⁰This restriction matters for the results that follow. Otherwise, the set of credible static auctions includes all-pay auctions, and the set of strategy-proof static auctions includes second-price auctions in which each agent also pays some mean-zero function of the other agents' bids.

¹¹The results that follow would require only small modifications if the auctioneer's payoff was a convex combination of revenue and social welfare.

2. For all i, θ_i , $\mathbf{E}_{\theta_{-i}}[u_i^G(S_N, \theta_N) \mid \theta_i] \geq 0$.

Definition 7. (G, S_N) is **optimal** if it is a solution to the above maximization problem. (G, S_N) is ϵ -**optimal** if it satisfies the constraints and the difference between $\pi(G, S_N)$ and the optimal expected revenue is no more than ϵ .

We use $\check{u}_i^{G, S_N}(k, k')$ to denote the expected utility of agent i in G when his type is θ_i^k and he plays as though his type is $\theta_i^{k'}$. In order for a protocol to be optimal, certain individual rationality (IR) and incentive compatibility (IC) constraints must bind.

Proposition 3. (Elkind, 2007) If (G, S_N) is optimal, then:

1. IR-0 binds: $\forall i : \check{u}_i^{G, S_N}(0, 0) = 0$
2. IC binds locally downward: $\forall i : \forall k \geq 1 : \check{u}_i^{G, S_N}(k, k) = \check{u}_i^{G, S_N}(k, k-1)$

Proposition 4. If IR-0 binds and IC binds locally downward, then:

$$\pi(G, S_N) = \mathbf{E}_{\theta_N} \left[\sum_{i \in N} y_i^{G, S_N}(\theta_N) \left(\theta_i - \frac{1 - F(\theta_i)}{f(\theta_i)} \right) \right] \quad (5)$$

where $y_i^{G, S_N}(\theta_N)$ is an indicator equal to 1 if i wins the object when the type profile is θ_N (under (G, S_N)), and 0 otherwise.

We use $\eta(\theta_i) \equiv \theta_i - \frac{1 - F(\theta_i)}{f(\theta_i)}$ to denote the **virtual value** of type θ_i . By Proposition 4, if IR-0 binds and IC binds locally downward, then $\pi(G, S_N)$ is equal to the expected virtual value of the winning bidder. Even when the local IC constraints are slack, the expected revenue is within ϵ of the expected virtual value of the winning bidder.

Proposition 5. If (G, S_N) is BIC, then:

$$0 \leq \underbrace{\mathbf{E}_{\theta_N} \left[\sum_{i \in N} y_i^{G, S_N}(\theta_N) \eta(\theta_i) \right]}_{\text{expected virtual value}} - \pi(G, S_N) - \underbrace{\sum_{i \in N} \check{u}_i^{G, S_N}(0, 0)}_{\text{slack in IR-0}} \leq \epsilon \quad (6)$$

Definition 8. F is **regular** if $\eta(\theta_i)$ is strictly increasing in θ_i .

For regular F , let ρ^* denote an **optimal reserve price**, with the property that if $\eta(\theta^k) < 0$ then $\theta^k < \rho^*$ and if $\eta(\theta^k) > 0$ then $\theta^k \geq \rho^*$.

Definition 9. Consider a strict tie-breaking order \triangleright on N . This generates a strict total order on all agent types, as follows: $\theta_i \triangleright \theta_j$ if and only if $\theta_i \geq \theta_j$ and either $\theta_i > \theta_j$ or $i \triangleright j$. We also include a reserve ρ in this total order: $\theta_i \triangleright \rho$ if and only if $\theta_i \geq \rho$. We use \min_{\triangleright} to denote the minimum of a set with respect to this \triangleright , and \max_{\triangleright} similarly.

(G, S_N) is **orderly** if, for some strict total order \triangleright on N and some reserve price $\rho \leq \theta_i^K$, i wins the object if and only if $\theta_i \triangleright \max_{j \neq i} \theta_j$ and $\theta_i \triangleright \rho$.

3.1 Credible and static ϵ -optimal auctions

We now characterize credible and static ϵ -optimal auctions. Our characterization uses ϵ -optimality rather than optimality, because the Revenue Equivalence Theorem does not hold in our setting. There may not exist optimal first-price auctions, though there exist ϵ -optimal first-price auctions.

Definition 10. (G, S_N) is **static** if every agent has exactly one information set and for every terminal history z , there exists $h \prec z$ such that $P(h) = i$.

Definition 11. (G, S_N) is a **quasi-first-price auction** if (G, S_N) is static, and each agent i chooses either to place a bid in a finite feasible set $B_i \subset \mathbb{R}$ or to decline¹², and:

1. Some bidder wins the object if and only if some bidder places a bid.
2. For every $b_i \in B_i$ there exists a profile of opponent actions such that that i wins after bidding b_i .
3. If i wins the object, then i pays his bid.
4. If i wins the object then:
 - (a) Either: i has the highest bid, which is non-negative.
 - (b) Or: i has type θ_i^K , has the highest tie-breaking priority according to \triangleright , and for all $j \in N$, i 's bid is greater than or equal to the bid that would be placed by type $\theta_j^K - \epsilon$ (if it exists).

We represent a reserve price by restricting the feasible bids B_i for each agent i . Clause 3(b) deals with a corner case: The highest-priority bidder can win the object for sure by placing a very high bid, even if this is slightly less than the highest bid. However, this bid must be *almost* the highest bid, in the sense that it must exceed the bid placed by type $\theta_j^K - \epsilon$. This anomaly vanishes as ϵ goes to 0.

Now we can state the main result of this section:

Theorem 1. Assume F is regular and (G, S_N) is orderly and BIC. (G, S_N) is credible and static if and only if (G, S_N) is a quasi-first-price auction. There exists a quasi-first-price auction that is ϵ -optimal.

Proof. Suppose (G, S_N) is a quasi-first-price auction. (G, S_N) is static by definition. If the highest-priority bidder has bid as though his type is θ_i^K , then the auctioneer has no discretion; every safe deviation sells the object to that bidder at his bid. Otherwise, no safe deviation can sell to a bidder who declined, and every safe deviation that sells the

¹²There can be multiple actions that decline, all having the same consequence.

object involves charging some bidder his bid, so it is optimal to sell the object by the rules. Thus, (G, S_N) is credible.

Suppose (G, S_N) is credible and static. Recall that (G, S_N) is pruned. If after some action by bidder i , there are two prices that i might pay, the auctioneer can safely deviate to charge the higher price. Thus, for each action, if i might win after playing that action then there is a unique price that i might pay (his *bid*). If i never wins after playing some action, that action is to *decline*. Let $b_i(\theta_i)$ denote the bid placed by bidder i with type θ_i , if it exists.

Suppose that i has the highest priority, but does not play the action taken by θ_i^K . Since (G, S_N) is orderly, for every agent $j \in N$, if j placed a bid, there is an innocent explanation if j wins and an innocent explanation if j loses. Thus, if j wins then j has placed the highest bid, or the auctioneer can safely deviate to sell the object to the highest bidder for strictly more revenue. Moreover, if j wins then j 's bid is non-negative, since the auctioneer can safely deviate to keep the object.

Suppose i has the highest priority, and plays the action assigned to θ_i^K . Since (G, S_N) is orderly, i must win for sure. If type $\theta_i^K - \epsilon$ declines, then clause 3(b) holds vacuously and we are done. Otherwise, $b_i(\theta_i^K) \geq b_i(\theta_i^K - \epsilon)$, since (G, S_N) is BIC.¹³ (G, S_N) is orderly, so when i 's type is $\theta_i^K - \epsilon$ and the highest opponent type is $\theta_j^K - \epsilon$, i wins the object. (G, S_N) is credible, so $b_i(\theta_i^K - \epsilon) \geq b_j(\theta_j^K - \epsilon)$, if $\theta_j^K - \epsilon$ places a bid. Thus, (G, S_N) is a quasi-first-price auction.

Finally, to show that there exists an ϵ -optimal quasi-first-price auction, choose some optimal reserve ρ^* , and for each i and θ_i , let $b_i(\theta_i)$ be θ_i 's expected payment *conditional on winning* under a second-price auction that is orderly and has reserve ρ^* .¹⁴ Such a mechanism respects all incentive constraints, and has the property that if $\theta_i \triangleright \theta_j \triangleright \rho^*$, then $b_i(\theta_i) \geq b_j(\theta_j) \geq \rho^*$. By Proposition 5, this is ϵ -optimal. \square

3.2 Credible and strategy-proof optimal auctions

We now characterize credible and strategy-proof optimal auctions.

Definition 12. (G, S_N) is **strategy-proof** if, for all $i \in N$, for all S'_{-i}

$$S_i \in \operatorname{argmax}_{S'_i} \mathbb{E}_{\theta_N} [u_i(S'_i, S'_{-i}, \theta_N)] \quad (8)$$

¹³Suppose $b_i(\theta_i^K) < b_i(\theta_i^K - \epsilon)$. By interim IR, $\theta_i^K - \epsilon - b_i(\theta_i^K - \epsilon) \geq 0$. $\theta_i^K - \epsilon$ could profitably imitate θ_i^K , to increase his probability of winning and decrease his payment when he wins, so (G, S_N) is not BIC.

¹⁴Formally, for some optimal reserve $\rho^* \geq \theta_i^0$,

$$b_i(\theta_i) = \begin{cases} \mathbb{E}[\max\{\rho^*, \max_{j \neq i} \theta_j\} \mid \forall j \neq i : \theta_i \triangleright \theta_j] & \text{if } \theta_i \triangleright \rho^* \\ \text{decline} & \text{otherwise} \end{cases} \quad (7)$$

Definition 13. (G, S_N) is an **ascending auction** (with reserve price ρ) if:

1. All bidders start as **active**, with initial bids $(b_i)_{i \in N} := (\theta_i^0)_{i \in N}$.
2. The **high bidder** is the active bidder with the highest bid that is weakly above ρ (breaking ties according to \triangleright).
3. At each non-terminal history, some active bidder i (other than the high bidder) is called to play, and he chooses between actions that place a **bid** in Θ_i and actions that **quit**.
 - (a) Each bid is no less than the last bid that i placed.
 - (b) Each bid is no more than is necessary for i to become the high bidder.
 - (c) If i quits, then he is no longer active.
 - (d) At each information set, there is a unique action that places a bid, with one exception: If the reserve has not yet been met, and there is exactly one active bidder left, there may be multiple actions that place bids.
4. i 's strategy specifies:
 - (a) If i 's type is strictly below a bid, he does not place that bid.
 - (b) If i 's type is weakly above ρ and there is no high bidder, he places a bid.
 - (c) If i 's type is above the current high bid (breaking ties with \triangleright), he places a bid.¹⁵
5. The auction ends if one of three conditions obtains:
 - (a) If there are no active bidders. In that case, the object is not sold.
 - (b) If only the high bidder is active. In that case, the object is sold to the high bidder at his last bid.
 - (c) If the high bidder has bid θ_i^K , and no active bidder has higher tie-breaking priority. In that case, the object is sold to the high bidder at his last bid.

We pause to note a mild indeterminacy: When an active bidder is called to play, it could be that the available bid is not yet enough to become the high bidder. For instance, bidder i might choose whether to place a bid of 50 or quit, even though the current high bid is 100. In that case, types of i between 50 and 100 could place the bid or could quit. However, S_i is measurable with respect to i 's information sets, so it must be that bidder i never quits when he might still win.

¹⁵Notice that, since i 's strategy must be measurable with respect to i 's information sets, this implies that if i 's type is above the *least possible high bid associated with that information set*, he places a bid.

Observation 3. *If F is regular, then there exists an optimal ascending auction. In any ascending auction, IR-0 binds and IC binds locally downward. Given an optimal reserve ρ^* , ascending auctions maximize the virtual value of the winning bidder. Thus, by Propositions 3 and 4, such an auction is optimal.*

The definition of extensive-form mechanisms permits the auctioneer to communicate with agents in any order, to convey any information (or no information) to the agent called to play, and to ask that agent to report any partition of his type space. Thus, there are many optimal auctions. However, the optimal auctions that are credible and strategy-proof are exactly the ascending auctions. To be precise:

Theorem 2. *Assume F is regular and (G, S_N) is orderly and optimal. (G, S_N) is credible and strategy-proof if and only if (G, S_N) is an ascending auction.*

Proof overview. Suppose (G, S_N) is credible and strategy-proof. To prove that (G, S_N) is an ascending auction, we must show that for *any* extensive form that is not an ascending auction, there exists a profitable safe deviation for the auctioneer. A key feature of ascending auctions is that whenever the winner is not yet decided, all types of i that might still win pool on the same action. This is stated precisely in Lemma 3, and is closely related to unconditional winner privacy as defined by Milgrom and Segal (2017). If at some history winning types do not pool, then the auctioneer can exploit one type by deviating to charge him a higher price. In the case of a second-price auction, the auctioneer simply exaggerates the value of the second-highest bid. In general, however, the deviation must be more subtle in order to be safe - instead of just choosing a different outcome, the auctioneer may systematically misrepresent agents' actions midway through the extensive form. We construct an algorithm that produces a profitable safe deviation for any such extensive form.

Suppose (G, S_N) is an ascending auction. By inspection, it is strategy-proof. What remains is to show that it is credible. Suppose that the auctioneer has a profitable safe deviation. For every agent i , S_i remains a best response to any safe deviation by the auctioneer. Thus, since the auctioneer has a profitable safe deviation, she can openly commit to that deviation without altering the agents' incentives - we can define a new protocol (G', S'_N) that is BIC and yields strictly more expected revenue than (G, S_N) . But (G, S_N) is optimal, so this is a contradiction. (The full proof is in the Appendix.) \square

By Theorem 1, restricting attention to revelation mechanisms forces a sharp choice between incentives for the auctioneer and strategy-proofness for the agents. Theorem 2 shows that allowing other extensive forms relaxes this trade-off.

While first-price auctions and ascending auctions seem to be disparate formats, they share a common feature. In both formats, if an agent might win the auction without being called to play again, then that agent knows exactly how much he will pay for the

object. Thus, we can regard each agent as placing bids in the course of the auction, with the assurance that if he wins without further intervention, he will pay his bid. This ‘pay-as-bid’ feature is shared by all credible auctions:

Proposition 6 (Pay-as-bid). *Assume (G, S_N) is credible. Suppose i is called to play at information set I_i , takes some action a , and might win without being called to play again. Then there is a unique price $t_i(I_i, a)$ that i will pay if he wins without being called to play again.*

Proof. Suppose i might win without being called to play again, and there are two distinct prices $t_i < t'_i$ that i might pay in that circumstance. The auctioneer has a profitable deviation; when i is meant to pay t_i , she can deviate to charge t'_i . (G, S_N) is pruned, so this has an innocent explanation, and (G, S_N) is not credible. \square

Proposition 6 provides a consideration in favor of multi-stage auctions. Suppose we wish to have bidder i ’s payment depend on bidder j ’s private information. In order for the auction to be credible, bidder i must place a bid that incorporates that information, which requires i to learn that information during the auction.

3.3 An Auction Trilemma

Clearly, the quasi-first-price auction is not strategy-proof, except in the degenerate case that the optimal reserve is almost equal to the highest type. Thus, Theorems 1 and 2 yield the following corollary.

Corollary 1 (Auction Trilemma). *Assume F is regular and (G, S_N) is orderly. Assume that for any optimal reserve $\rho^* < \theta_i^K - 3\epsilon$. No ϵ -optimal (G, S_N) is static, strategy-proof, and credible. However, there exist ϵ -optimal (G, S_N) that are:*

1. *static and strategy-proof (the second-price auction),*
2. *static and credible (the first-price auction),*
3. *strategy-proof and credible (the ascending auction).*

Credibility, **staticity**, and strategy-proofness are all desirable properties. No optimal auction has all three, and picking any two of three characterizes a standard auction format (Figure 1).

4 Asymmetric Distributions

So far we have assumed that bidders’ values are independently and identically distributed. Suppose instead that type spaces and probability distributions could be asymmetric. That

is, it may be that for agents i, j and some $k \in \mathbb{Z}$, $\theta_i^k \neq \theta_j^k$, and each agent has a probability mass function $p_i : \Theta_i \rightarrow [0, 1]$. Define f_i and F_i similarly, and

$$\eta_i(\theta_i) \equiv \theta_i - \frac{1 - F_i(\theta_i)}{f_i(\theta_i)} \quad (9)$$

Under asymmetry, first-price auctions remain credible but may no longer be ϵ -optimal. To see why this is so, consider the following simple example.

Example 2. *There are two bidders. Bidder 1's value is uniformly distributed with support $\Theta_1 = \{\epsilon, 2\epsilon, \dots, 1\}$ and bidder 2's value is uniformly distributed with support $\Theta_2 = \{1 + \epsilon, 1 + 2\epsilon, \dots, 2\}$, where $\epsilon = 2^{-k}$ for positive integer k . Thus, the virtual values are $\eta_1(\theta_1) = 2\theta_1 - 1$ for bidder 1 and $\eta_2(\theta_2) = 2\theta_2 - 2$ for bidder 2.*

Proposition 7. *There exist asymmetric value distributions such that no (G, S_N) is credible, static, and ϵ -optimal.*

Proof overview. In Example 2, when ϵ is small enough, no (G, S_N) is credible, static, and ϵ -optimal. When θ_1 is slightly below 1 and θ_2 is slightly above 1, $\eta_1(\theta_1) > \eta_2(\theta_2)$. Thus, an ϵ -optimal mechanism must sell to bidder 1 at some such type profile. Suppose (G, S_N) is credible and static, so if an agent wins he pays his bid $b_i(\theta_i)$. By interim individual rationality, $b_1(\theta_1) \leq \theta_1$. By ϵ -optimality, $\theta_1 < b_2(\theta_2)$. At type profile (θ_1, θ_2) , the auctioneer can safely deviate to sell the object to θ_2 , which yields strictly more revenue, a contradiction. (The full proof is in the Appendix.) \square

Example 2 demonstrates that, under asymmetry, the best credible static mechanism may deliver materially less revenue than the optimum.¹⁶ The next natural question is: When value distributions are regular but asymmetric, can any credible auction attain the optimum? It turns out that not only is there a credible optimal auction, there is even a strategy-proof credible optimal auction.

We modify the ascending auction of Definition 13, instead scoring bids according to their virtual value.

Definition 14. *(G, S_N) is a **virtual ascending auction** if:*

1. All bidders start as **active**, with initial bids $(b_i)_{i \in N} := (\theta_i^0)_{i \in N}$.
2. The **high bidder** is the active bidder with the highest **virtual value** $\eta_i(b_i)$ that is non-negative (breaking ties according to \triangleright).
3. At each non-terminal history, some active bidder i (other than the high bidder) is called to play, and he chooses either to place a **bid** in Θ_i or to **quit**.

¹⁶This example can be modified so that both bidders' type distributions have the same support, but disjoint supports simplify the argument.

- (a) *At each information set, there is a unique action that places a bid.*
 - (b) *Each bid is no less than the last bid that i placed.*
 - (c) *Each bid is no more than is necessary for i to become the high bidder.*
 - (d) *If i quits, then he is no longer active.*
4. *i 's strategy specifies that he places a bid if and only if his type is weakly more than that bid.*
5. *The auction ends if one of three conditions obtains:*
- (a) *If there are no active bidders. In that case, the object is not sold.*
 - (b) *If only the high bidder is active. In that case, the object is sold to the high bidder at his last bid.*
 - (c) *If the high bidder i has bid θ_i^K , and no active bidder can make a bid with a strictly higher virtual value. In that case, the object is sold to the high bidder at his last bid.*

To illustrate, consider the type distributions of Example 2. The virtual ascending auction proceeds as follows: At the beginning, the virtual value of bidder 2 is at least 2ϵ so he is the current high bidder, with bid $1 + \epsilon$. We first ask bidder 1 to bid $\frac{1}{2} + \epsilon$. If he quits, bidder 2 wins and pays $1 + \epsilon$. If bidder 1 places the bid, then we know his virtual value is at least as high as bidder 2's virtual value, so he becomes the new high bidder. We then alternate between the bidders, asking each to increase his bid by ϵ . This implements the optimal allocation.

It is easy to verify that the virtual ascending auction is optimal and strategy-proof. To verify that it is credible as well, note that our proof strategy for the symmetric case applies *mutatis mutandis* to this case as well. The virtual ascending auction is optimal, and the agents' strategies are a best response to any safe deviation. Thus, if the auctioneer has a profitable safe deviation, she could openly commit to that deviation, producing a new BIC protocol that yields strictly more expected revenue than the optimum, which is a contradiction.

In summary, asymmetry interacts with credibility in subtle ways. Clearly, when type distributions are asymmetric, standard first-price auctions and ascending auctions may not be optimal. First-price auctions remain credible¹⁷, but Example 2 shows that there may be no way to modify them to restore optimality (while retaining credibility and sealed bids). Ascending auctions may not even be credible - the auctioneer may profitably deviate by running agents' price clocks out-of-sync. However, modifying the scoring rule restores both credibility and optimality, while retaining strategy-proofness.

¹⁷First-price auctions are 'robustly' credible, in the sense that they are credible regardless of the joint distribution of types.

The foregoing arguments suggest another reason to look outside the class of static mechanisms: General extensive forms allow the auctioneer to credibly commit to treat bidders asymmetrically.¹⁸ Under asymmetric distributions, the best credible auction can yield strictly more revenue than the best static credible auction. However, this gap depends on the restriction that losing bidders do not make payments. The following modified all-pay auction is static, credible, and optimal: Each type θ_i makes a bid equal to the expected payment of θ_i in the virtual ascending auction. Every agent pays his bid, and the auctioneer awards the object to the agent with the highest non-negative virtual value.¹⁹

5 A ‘prior-free’ definition

The definition of credibility depends on the joint distribution of agent types (Definition 6). It may be useful to have a definition that is ‘prior-free’, for settings such as matching or maxmin mechanism design.

Definition 15. *Given (G, S_N) , $S_0 \in \mathcal{S}_0^*(S_0^G, S_N)$ is **always-profitable** if, for all θ_N :*

$$u_0(S_0, S_N, \theta_N) \geq u_0(S_0^G, S_N, \theta_N) \quad (10)$$

with strict inequality for some θ_N .

*(G, S_N) is **weakly credible** if no safe deviation is always-profitable.*

For comparison, (G, S_N) is credible if no safe deviation is profitable in expectation. Weak credibility allows one to dispense with strong assumptions about the auctioneer’s beliefs.

What happens if we replace “credible” with “weakly credible” in the statement of Theorems 1 and 2? Since any credible protocol is weakly credible, this weakens one direction of implication but strengthens the other. Weak credibility still suffices to characterize both auctions.

Theorem 1*. *Assume F is regular and (G, S_N) is orderly and BIC. (G, S_N) is weakly credible and static if and only if (G, S_N) is a quasi-first-price auction. There exists a quasi-first-price auction that is ϵ -optimal.*

Theorem 2*. *Assume F is regular and (G, S_N) is orderly and optimal. (G, S_N) is weakly credible and strategy-proof if and only if (G, S_N) is an ascending auction.*

¹⁸However, when one considers worst-case equilibria, standard first-price auctions have good revenue guarantees (Roughgarden, 2009; Feldman et al., 2016; Bergemann et al., 2017).

¹⁹Of course, all-pay auctions are rarely used in practice. This may be because losing bidders can find ways to renege on their payments, or because bidders dislike the prospect of making large payments and getting nothing in return. The robustness of all-pay auctions to auctioneer cheating is studied in an early draft of Dequiedt and Martimort (2015), circulated under the title “Mechanism Design with Private Communication”.

Proof. Since credible protocols are weakly credible, the “if” direction is implied by Theorems 1 and 2.

In the proof of Theorem 1, we show that if (G, S_N) is static but not a quasi-first-price auction, then there exists a safe deviation that is always-profitable. This proves the “only if” direction of Theorem 1*.

In the proof of Theorem 2, we show that if (G, S_N) is strategy-proof but not an ascending auction, then there exists a safe deviation that is always-profitable. This proves the “only if” direction of Theorem 2*. \square

6 Discussion

To raise the old question once more: What accounts for the popularity of common real-world auction forms? “What determines which form will (or should) be used in any particular circumstance?” (Milgrom and Weber, 1982)

Our analysis indicates that the first-price auction and the ascending auction are not just historical accidents, but are game-theoretic solutions to a well-defined commitment problem. Furthermore, each auction form is the unique solution in its class.

It is worth considering why real-world auctioneers might lack full commitment power. Vickrey (1961) suggests that the seller could delegate the task of running the auction to a third-party who has no stake in the outcome. However, auction houses such as Sotheby’s, Christie’s, and eBay charge commissions that are piecewise-linear functions of the sale price.²⁰ Running an auction takes effort, and many dimensions of effort are not contractible. Robust contracts reward the auctioneer linearly with revenue (Carroll, 2015), so it is difficult to employ a third-party who is both neutral and well-motivated.²¹

When an auctioneer makes repeated sales, reputation could help enforce the full-commitment outcome. However, the force of reputation depends on the discount rate and the detection rate of deviations. Safe deviations are precisely those that a bidder could not detect immediately. Online advertising auctions are repeated frequently, so it is plausible that bidders could examine the statistics to detect foul play.²² However,

²⁰<http://www.sothebys.com/en/news-video/blogs/all-blogs/sotheby-s-at-large/2016/10/important-update-regarding-sothebys-buyers-premium.html>, <http://www.christies.com/buying-services/buying-guide/financial-information/>, <http://pages.ebay.com/help/sell/fees.html>, accessed 11/5/2017.

²¹As Myerson (2009) observes, “The problems of motivating hidden actions can explain why efficient institutions give individuals property rights, as owners of property are better motivated to maintain it. But property rights give people different vested interests, which can make it more difficult to motivate them to share their private information with each other.”

²²However, bidders in online advertising auctions have expressed concerns that supply-side platforms (SSPs) are deviating from the rules of the second-price auction. The industry news website Digiday alleged, “Rather than setting price floors as a flat fee upfront, some SSPs are setting high price floors after their bids come in as a way to squeeze out more money from ad buyers who believe they are bidding into a second-price auction”. <https://digiday.com/marketing/ssps-use-deceptive-price-floors-squeeze-ad-buyers/>, accessed 11/30/2017.

some economically important auctions are infrequent or not repeated at all - for instance, auctions for wireless spectrum or for the privatization of state-owned industries. Even established auction houses such as Christie's and Sotheby's have faced regulatory scrutiny, based in part on concerns that certain deviations are difficult for individual bidders to detect.

Modern auctioneers could use cryptography to prove that the rules of the auction have been followed, without disclosing additional information to bidders. Cryptographic verification relies on digital infrastructure: Participants typically need access to a public bulletin board, a sound method of creating and sharing public keys, and a time-lapse encryption service that provides public keys and commits to release the corresponding decryption keys only at pre-defined times (Parkes et al., 2015).²³ It can be costly to construct this infrastructure, and to persuade bidders that it works as the auctioneer claims. By using credible mechanisms, auctioneers may increase the resources and attention available for substantive purposes.

Not all auctioneers have full commitment power, just as not all firms are Stackelberg leaders. When the auctioneer lacks full commitment, it can be hazardous for bidders to reveal all their information at once. In a first-price auction, a bidder 'reveals' his value in return for a guarantee that his report completely determines the price he might pay.²⁴ In an ascending auction, a bidder reports whether his value is above b only when the auctioneer (correctly) asserts that bids below b are not enough to win. Credibility is a shared foundation for these seemingly disparate design features. How this principle extends to other environments is an open question.

References

- Athey, S., Levin, J., and Seira, E. (2011). Comparing open and sealed bid auctions: Evidence from timber auctions. *The Quarterly Journal of Economics*, 126(1):207–257.
- Baliga, S., Corchon, L. C., and Sjöström, T. (1997). The theory of implementation when the planner is a player. *Journal of Economic Theory*, 77(1):15–33.
- Bergemann, D., Brooks, B., and Morris, S. (2017). First-price auctions with general information structures: Implications for bidding and revenue. *Econometrica*, 85(1):107–143.
- Bergemann, D. and Morris, S. (2005). Robust mechanism design. *Econometrica*, 73(6):1771–1813.

²³Bidders may even need special training or software assistance to play their part in a cryptographic protocol.

²⁴This property is generalized in a natural way by the 'first-price' menu auction (Bernheim and Whinston, 1986).

- Bernheim, B. D. and Whinston, M. D. (1986). Menu auctions, resource allocation, and economic influence. *The Quarterly Journal of Economics*, 101(1):1–31.
- Bester, H. and Strausz, R. (2001). Contracting with imperfect commitment and the revelation principle: the single agent case. *Econometrica*, 69(4):1077–1098.
- Burguet, R. and Sakovics, J. (1996). Reserve prices without commitment. *Games and Economic Behavior*, 15(2):149–164.
- Caillaud, B. and Mezzetti, C. (2004). Equilibrium reserve prices in sequential ascending auctions. *Journal of Economic Theory*, 117(1):78–95.
- Carroll, G. (2015). Robustness and linear contracts. *American Economic Review*, 105(2):536–563.
- Cassady, R. (1967). *Auctions and auctioneering*. Univ of California Press.
- Chakraborty, I. and Kosmopoulou, G. (2004). Auctions with skill bidding. *Economic Theory*, 24(2):271–287.
- Chung, K.-S. and Ely, J. C. (2007). Foundations of dominant-strategy mechanisms. *The Review of Economic Studies*, 74(2):447–476.
- Dequiedt, V. and Martimort, D. (2015). Vertical contracting with informational opportunism. *American Economic Review*, 105(7):2141–2182.
- Elkind, E. (2007). Designing and learning optimal finite support auctions. In *Proceedings of the eighteenth annual ACM-SIAM symposium on Discrete algorithms*, pages 736–745. Society for Industrial and Applied Mathematics.
- Feldman, M., Lucier, B., and Nisan, N. (2016). Correlated and coarse equilibria of single-item auctions. In *International Conference on Web and Internet Economics*, pages 131–144. Springer.
- Grant, D. (2014). Why auction rooms seem empty these days. *The Wall Street Journal*.
- Green, J. and Laffont, J.-J. (1977). Characterization of satisfactory mechanisms for the revelation of preferences for public goods. *Econometrica*, pages 427–438.
- Holmström, B. (1979). Groves’ scheme on restricted domains. *Econometrica*, pages 1137–1144.
- Hurwicz, L. (1972). On informationally decentralized systems. In McGuire, C. B. and Radner, R., editors, *Decision and Organization*, chapter 14, pages 297–336. North-Holland, Amsterdam.

- Krishna, V. (2010). *Auction Theory*. Academic Press, San Diego, USA.
- Lamy, L. (2009). The shill bidding effect versus the linkage principle. *Journal of Economic Theory*, 144(1):390–413.
- Li, S. (2017). Obviously strategy-proof mechanisms. *American Economic Review*, 107(11):3257–87.
- Liu, Q., Mierendorff, K., Shi, X., et al. (2014). Auctions with limited commitment. Technical report, Working Paper.
- Lobel, I. and Paes Leme, R. (2017). Dynamic mechanism design under positive commitment. *working paper*.
- Loertscher, S. and Marx, L. M. (2017). Optimal clock auctions. *working paper*.
- Lucking-Reiley, D. (2000). Vickrey auctions in practice: From nineteenth-century philately to twenty-first-century e-commerce. *The Journal of Economic Perspectives*, 14(3):183–192.
- Maskin, E. and Tirole, J. (1990). The principal-agent relationship with an informed principal: The case of private values. *Econometrica*, pages 379–409.
- Maskin, E. and Tirole, J. (1992). The principal-agent relationship with an informed principal, ii: Common values. *Econometrica*, pages 1–42.
- McAdams, D. and Schwarz, M. (2007). Credible sales mechanisms and intermediaries. *American Economic Review*, 97(1):260–276.
- McAfee, R. P. and McMillan, J. (1987). Auctions and bidding. *Journal of economic literature*, 25(2):699–738.
- McAfee, R. P. and Vincent, D. (1997). Sequentially optimal auctions. *Games and Economic Behavior*, 18(2):246–276.
- Milgrom, P. (1987). Auction theory. In Bewley, T., editor, *Advances in Economic Theory: Fifth World Congress*, pages 1–32. Cambridge University Press, Cambridge.
- Milgrom, P. and Segal, I. (2017). Deferred-acceptance auctions and radio spectrum reallocation. Technical report, working paper.
- Milgrom, P. R. and Weber, R. J. (1982). A theory of auctions and competitive bidding. *Econometrica: Journal of the Econometric Society*, pages 1089–1122.
- Myerson, R. B. (1981). Optimal auction design. *Mathematics of Operations Research*, 6(1):58–73.

- Myerson, R. B. (1983). Mechanism design by an informed principal. *Econometrica*, pages 1767–1797.
- Myerson, R. B. (2009). Fundamental theory of institutions: a lecture in honor of leo hurwicz. *Review of Economic Design*, 13(1-2):59.
- Myerson, R. B. and Reny, P. J. (2016). Broad sequential equilibria of multi-stage games with infinite sets of signals and actions. *Working paper*.
- Parkes, D. C., Thorpe, C., and Li, W. (2015). Achieving trust without disclosure: Dark pools and a role for secrecy-preserving verification. In *Third Conference on Auctions, Market Mechanisms and Their Applications*.
- Porter, R. and Shoham, Y. (2005). On cheating in sealed-bid auctions. *Decision Support Systems*, 39(1):41–54.
- Rothkopf, M. H. and Harstad, R. M. (1995). Two models of bid-taker cheating in vickrey auctions. *Journal of Business*, pages 257–267.
- Rothkopf, M. H., Teisberg, T. J., and Kahn, E. P. (1990). Why are Vickrey auctions rare? *Journal of Political Economy*, pages 94–109.
- Roughgarden, T. (2009). Intrinsic robustness of the price of anarchy. In *Proceedings of the forty-first annual ACM symposium on Theory of computing*, pages 513–522. ACM.
- Shubik, M. (2004). *The Theory of Money and Financial Institutions: Volume 1*. MIT Press, Massachusetts, USA.
- Simon, L. K. and Stinchcombe, M. B. (1989). Extensive form games in continuous time: Pure strategies. *Econometrica*, pages 1171–1214.
- Skreta, V. (2015). Optimal auction design under non-commitment. *Journal of Economic Theory*, 159:854–890.
- Vartiainen, H. (2013). Auction design without commitment. *Journal of the European Economic Association*, 11(2):316–342.
- Vickrey, W. (1961). Counterspeculation, auctions, and competitive sealed tenders. *The Journal of Finance*, 16(1):8–37.
- Wilson, R. (1987). Game theoretic analysis of trading processes. In Bewley, T., editor, *Adv. Econ. Theory*. Cambridge University Press.

A Definition of Extensive Game Forms with Consequences in X

An *extensive game form with consequences in X* is a tuple $(H, \prec, P, A, \mathcal{A}, (\mathcal{I}_i)_{i \in N}, g)$, where:

1. H is a set of histories, along with a binary relation \prec on H that represents precedence.
 - (a) \prec is a partial order, and (H, \prec) form an arborescence²⁵.
 - (b) h_\emptyset denotes $h \in H : \neg \exists h' : h' \prec h$
 - (c) $Z \equiv \{h \in H : \neg \exists h' : h \prec h'\}$
 - (d) $\sigma(h)$ denotes the set of immediate successors of h .
2. P is a player function. $P : H \setminus Z \rightarrow N$.
3. A is a set of actions.
4. $\mathcal{A} : H \setminus h_\emptyset \rightarrow A$ labels each non-initial history with the last action taken to reach it.
 - (a) For all h , \mathcal{A} is one-to-one on $\sigma(h)$.
 - (b) $A(h)$ denotes the actions available at h .

$$A(h) \equiv \bigcup_{h' \in \sigma(h)} \mathcal{A}(h') \quad (11)$$

5. \mathcal{I}_i is a partition of $\{h : P(h) = i\}$ such that:
 - (a) $A(h) = A(h')$ whenever h and h' are in the same cell of the partition.
 - (b) For any $I_i \in \mathcal{I}_i$, we denote: $P(I_i) \equiv P(h)$ for any $h \in I_i$. $A(I_i) \equiv A(h)$ for any $h \in I_i$.
 - (c) Each action is available at only one information set: If $a \in A(I_i)$, $a' \in A(I'_j)$, $I_i \neq I'_j$ then $a \neq a'$.
6. g is an outcome function. It associates each terminal history with an outcome.
$$g : Z \rightarrow X$$

²⁵That is, a directed rooted tree such that every edge points away from the root.

B Proofs omitted from the main text

B.1 Proposition 1

For each of the three clauses in Definition 2, we show that if (G, S_N) does not satisfy this clause, we can transform (G, S_N) to have strictly fewer histories, such that the transformed protocol is BIC and results in the same outcomes for each type profile. Since the set of histories in (G, S_N) is finite, it follows that there exists a BIC (G', S'_N) that cannot be reduced further, but results in the same outcomes as (G, S_N) .

Suppose there exists h such that there is no θ_N such that $h \preceq z(S_N, \theta_N)$. Since the game tree is finite, we can locate an earliest possible h ; that is, an h such that no predecessor satisfies this property. Consider h' that immediately precedes h , and the information set I'_i such that $h \in I'_i$. There is some action a' at I'_i that is not played by any type of i that reaches I'_i . We can delete all histories that follow i playing a' at I'_i (and define $(\prec', \mathcal{A}', P', (\mathcal{I}'_i)_{i \in N}, g')$ and S'_N so that they are as in G , but restricted to the new smaller set of histories H'). Since these histories were off the path of play, their deletion does not affect the incentives of agents in $N \setminus i$. Since i preferred his original S_i to any strategy that played a' at I'_i , his new strategy S'_i remains incentive-compatible. Thus, the transformed (G', S'_N) is BIC.

Suppose there exists $h \notin Z$ such that $|\sigma(h)| = 1$. We simply rewrite the transformed game (G', S'_N) that deletes h (and all the other histories in that same information set) and ‘automates’ i ’s singleton action at h . That is, for all $h' \in I_i$ for I_i such that $h \in I_i$, we remove h' from the set of histories, and define $(\prec', \mathcal{A}', P', (\mathcal{I}'_i)_{i \in N}, g')$ and S'_N so that they are as in G , but restricted to $H \setminus I_i$. (G', S'_N) is BIC.

Suppose there exists $h \notin Z$, such that for $i = P(h)$, there does not exist $\theta_i, \theta'_i, \theta_{-i}$ such that

1. $h \prec z(S_N, (\theta_i, \theta_{-i}))$
2. $h \prec z(S_N, (\theta'_i, \theta_{-i}))$
3. $x^G(S_N, (\theta_i, \theta_{-i})) \neq x^G(S_N, (\theta'_i, \theta_{-i}))$

If there does not even exist (θ_i, θ_{-i}) such that $h \prec z(S_N, (\theta_i, \theta_{-i}))$, then our argument about the first clause applies. Suppose there exists (θ_i, θ_{-i}) such that $h \prec z(S_N, (\theta_i, \theta_{-i}))$. Upon reaching h , we can henceforth ‘automate’ play as though i had type θ_i . That is, we can delete any history h' such that $h \preceq h'$ and $P(h') = i$, or $h \preceq h'$ and there does not exist θ''_{-i} such that $h' \prec z(S_N, (\theta_i, \theta''_{-i}))$. Given the new smaller set of histories H' , we again define $(\prec', \mathcal{A}', P', (\mathcal{I}'_i)_{i \in N}, g')$ and S'_N so that they are as in G , but restricted to H' .

By construction, for all θ'_i , if i is playing as though his type is θ'_i and we would have reached history h under (G, S_N) , then the outcome is the same under (G', S'_N) as when i is playing as though his type is θ_i under (G, S_N) (which by hypothesis is the same as when

i is playing as though his type is θ'_i under (G, S_N)). Plainly, if we would not have reached history h under (G, S_N) , then the outcomes under (G, S_N) and (G', S'_N) are identical. Thus, (G', S'_N) is BIC.

This completes the proof of Proposition 1.

B.2 Proposition 2

To prove Proposition 2, we show that each of the three transformations we used in the proof of Proposition 1 also preserve credibility. That is, for each (G', S'_N) that is produced from (G, S_N) by one of the three transformations, if the auctioneer has a profitable safe deviation from $S_0^{G'}$, then she also has a profitable safe deviation from S_0^G .

Consider the first transformation (deleting all histories that follow action a at I_i , when a is never chosen on the path of play). Suppose the auctioneer had a profitable safe deviation S'_0 from $S_0^{G'}$. The auctioneer could make that same deviation, but additionally offer the response $\lambda(a)$ whenever sending the message $\lambda(I_i)$. By hypothesis, agent i never selects $\lambda(a)$ as a reply, so for any θ_N and any j , j 's resulting observation has an innocent explanation. Thus, the auctioneer also has a profitable safe deviation from S_0^G .

Consider the second transformation (deleting all histories in some information set with a singleton action set). Suppose the auctioneer had a profitable safe deviation S'_0 from $S_0^{G'}$. The auctioneer could make that same deviation from S_0^G , except that for the deleted information set I_i , the auctioneer delays sending $\lambda(I_i, A(I_i))$ until the last possible moment. That is consider S_0 that is the same as S'_0 , except that, if the auctioneer has not yet sent $\lambda(I_i, A(I_i))$ to agent i , then:

1. If S'_0 specifies that the auctioneer sends $\lambda(I'_i, A(I'_i))$ for $I'_i \succ I_i$, then S_0 specifies that she first sends $\lambda(I_i, A(I_i))$ and then (immediately thereafter) sends $\lambda(I'_i, A(I'_i))$
2. If S'_0 specifies that the auctioneer chooses an outcome such that the resulting observation for i does not have an innocent explanation under S_0^G , then S_0 specifies that she first sends $\lambda(I_i, A(I_i))$ before choosing that outcome.

S_0 is a profitable safe deviation from S_0^G .

Consider the third transformation (deleting histories where i is called to play, following some history h such that, for any two types of i that reach h , both types of i result in the same outcome). Suppose S'_0 was a profitable safe deviation from $S_0^{G'}$.

If the observation for i that results from S'_0 does not have an innocent explanation under S_0^G , it must be that (given on all the communication i has seen so far), the outcome S'_0 is about to select can only occur under G at terminal histories that follow h . But by hypothesis, for any θ_i and θ'_i that are consistent with reaching h , and any θ_{-i} consistent with reaching h , the resulting outcome is the same. Thus, let S_0 be exactly as in S'_0 , except that if S'_0 specifies that the auctioneer chooses an outcome such that the resulting

observation for i does not have an innocent explanation under S_0^G , then S_0 specifies that the auctioneer communicates with i as though play started from h and the opponent type profiles were θ_{-i} , for some θ_{-i} consistent with reaching h .

Formally, if S_0' would choose an outcome such that i 's observation has no innocent explanation, then fix some θ_N such that $h \prec z(S_N, \theta_N)$. Initialize $\hat{h} := h$.

1. If $\hat{h} \in Z$, then terminate and choose $x = g(\hat{h})$.
2. Else if $P(\hat{h}) \neq i$, then for $I_{P(\hat{h})}$ such that $\hat{h} \in I_{P(\hat{h})}$:
 - (a) $\hat{h} := h' \mid h' \in \sigma(\hat{h})$ and $S_{P(\hat{h})}(I_{P(\hat{h})}, \theta_{P(\hat{h})}) = \mathcal{A}(h')$.
 - (b) Go to step 1.
3. Else:
 - (a) Send $(m, R) = \lambda(I_i, A(I_i))$ for I_i such that $\hat{h} \in I_i$.
 - (b) Upon receiving $r \in R$, choose $\hat{h} := h' \mid \mathcal{A}(h') = \lambda^{-1}(r)$ and $h' \in \sigma(\hat{h})$.
 - (c) Go to step 1.

Since, under S_i , i 's play in this final stage makes no difference to the outcome, delaying communication with i until the outcome is about to be selected results in a safe deviation. This completes the proof of Proposition 2.

B.3 Proposition 4

Let $y_i^{G,SN}(\theta_i, \theta_{-i})$ denote the probability i gets the object when his type is θ_i and his opponent types are θ_{-i} .

Thus,

$$\begin{aligned} \check{u}_i^{G,SN}(k, k) - \check{u}_i^{G,SN}(k-1, k-1) &= \check{u}_i^{G,SN}(k, k-1) - \check{u}_i^{G,SN}(k-1, k-1) \\ &= \epsilon \mathbf{E}_{\theta_{-i}}[y_i^{G,SN}(\theta^{k-1}, \theta_{-i})] \end{aligned} \quad (12)$$

$$\check{u}_i^{G,SN}(k, k) = \sum_{l=1}^k \epsilon \mathbf{E}_{\theta_{-i}}[y_i^{G,SN}(\theta^{l-1}, \theta_{-i})] = \mathbf{E}_{\theta_{-i}}\left[\sum_{l=1}^k \epsilon y_i^{G,SN}(\theta^{l-1}, \theta_{-i})\right] \quad (13)$$

$$\mathbf{E}_{\theta_i}(\mathbf{E}_{\theta_{-i}}(\text{revenue from } i)) = \mathbf{E}_{\theta_{-i}}(\mathbf{E}_{\theta_i}(\text{welfare from } i)) - \mathbf{E}_{\theta_i}(\mathbf{E}_{\theta_{-i}}(\text{utility for } i)) \quad (14)$$

$$\mathbf{E}_{\theta_{-i}}(\mathbf{E}_{\theta_i}(\text{welfare from } i)) = \mathbf{E}_{\theta_{-i}}\left[\sum_{k=0}^K \epsilon f(\theta^k) \theta^k y_i^{G,SN}(\theta^k, \theta_{-i})\right] \quad (15)$$

$$\begin{aligned}
\mathbf{E}_{\theta_i}(\mathbf{E}_{\theta_{-i}}(\text{utility for } i)) &= \mathbf{E}_{\theta_{-i}}\left[\sum_{k=1}^K \epsilon f(\theta^k) \sum_{l=1}^k \epsilon y_i^{G,S_N}(\theta^{l-1}, \theta_{-i})\right] \\
&= \mathbf{E}_{\theta_{-i}}\left[\sum_{k=0}^K (1 - F(\theta^k)) \epsilon y_i^{G,S_N}(\theta^k, \theta_{-i})\right] \\
&= \mathbf{E}_{\theta_{-i}}\left[\sum_{k=0}^K \epsilon f(\theta^k) \frac{1 - F(\theta^k)}{f(\theta^k)} y_i^{G,S_N}(\theta^k, \theta_{-i})\right] \quad (16)
\end{aligned}$$

Combining Equations 15 and 16 yields:

$$\mathbf{E}_{\theta_i}(\mathbf{E}_{\theta_{-i}}(\text{revenue from } i)) = \mathbf{E}_{\theta_{-i}}\left[\sum_{k=0}^K \epsilon f(\theta^k) y_i^{G,S_N}(\theta^k, \theta_{-i}) \underbrace{\left[\theta^k - \frac{1 - F(\theta^k)}{f(\theta^k)}\right]}_{\text{virtual value}}\right] \quad (17)$$

Summing across agents yields Proposition 4.

B.4 Proposition 5

Since (G, S_N) is BIC, every type of i prefers playing according to S_i , instead of playing as though his type is ϵ lower. That is,

$$\forall i : \forall k > 1 : \check{u}_i^{G,S_N}(k, k) \geq \check{u}_i^{G,S_N}(k, k-1) \quad (18)$$

Thus,

$$\begin{aligned}
\check{u}_i^{G,S_N}(k, k) - \check{u}_i^{G,S_N}(0, 0) &= \sum_{l=1}^k \check{u}_i^{G,S_N}(l, l) - \check{u}_i^{G,S_N}(l-1, l-1) \\
&\geq \sum_{l=1}^k \check{u}_i^{G,S_N}(l, l-1) - \check{u}_i^{G,S_N}(l-1, l-1) = \mathbf{E}_{\theta_{-i}}\left[\sum_{l=1}^k \epsilon y_i^{G,S_N}(\theta^{l-1}, \theta_{-i})\right] \quad (19)
\end{aligned}$$

Using the same steps as in Proposition 4 (but with inequalities) yields:

$$0 \leq \mathbf{E}_{\theta_N}\left[\sum_{i \in N} y_i^{G,S_N}(\theta_N) \eta(\theta_i)\right] - \pi(G, S_N) - \sum_{i \in N} \check{u}_i^{G,S_N}(0, 0) \quad (20)$$

Since (G, S_N) is BIC, every type of i prefers playing according to S_i , instead of playing as though his type is ϵ higher. That is,

$$\forall i : \forall k > 1 : \check{u}_i^{G,S_N}(k-1, k-1) \geq \check{u}_i^{G,S_N}(k-1, k) \quad (21)$$

$$\begin{aligned}
\check{u}_i^{G,S_N}(k,k) - \check{u}_i^{G,S_N}(0,0) &= \sum_{l=1}^k \check{u}_i^{G,S_N}(l,l) - \check{u}_i^{G,S_N}(l-1,l-1) \\
&\leq \sum_{l=1}^k \check{u}_i^{G,S_N}(l,l) - \check{u}_i^{G,S_N}(l-1,l) = \sum_{l=1}^k \epsilon \mathbf{E}_{\theta_{-i}}[y_i^{G,S_N}(\theta^l, \theta_{-i})]
\end{aligned} \tag{22}$$

$$\begin{aligned}
\mathbf{E}_{\theta_i}(\mathbf{E}_{\theta_{-i}}(\text{utility for } i)) - \check{u}_i^{G,S_N}(0,0) &\leq \mathbf{E}_{\theta_{-i}}\left[\sum_{k=1}^K \epsilon f(\theta^k) \sum_{j=1}^k \epsilon y_i^{G,S_N}(\theta^j, \theta_{-i})\right] \\
&= \mathbf{E}_{\theta_{-i}}\left[\sum_{k=0}^K (1 - F(\theta^k)) \epsilon y_i^{G,S_N}(\theta^k, \theta_{-i})\right] + \mathbf{E}_{\theta_{-i}}\left[\sum_{k=0}^K \epsilon f(\theta^k) (y_i^{G,S_N}(\theta^k, \theta_{-i}) - y_i^{G,S_N}(\theta^0, \theta_{-i}))\right] \\
&= \mathbf{E}_{\theta_{-i}}\left[\sum_{k=0}^K \epsilon f(\theta^k) \frac{1 - F(\theta^k)}{f(\theta^k)} y_i^{G,S_N}(\theta^k, \theta_{-i})\right] + \mathbf{E}_{\theta_{-i}}\left[\sum_{k=0}^K \epsilon f(\theta^k) (y_i^{G,S_N}(\theta^k, \theta_{-i}) - y_i^{G,S_N}(\theta^0, \theta_{-i}))\right]
\end{aligned} \tag{23}$$

Again using the same steps as in Proposition 4 yields:

$$\begin{aligned}
\mathbf{E}_{\theta_N}\left[\sum_{i \in N} y_i^{G,S_N}(\theta_N) \eta(\theta_i)\right] - \pi(G, S_N) - \sum_{i \in N} \check{u}_i^{G,S_N}(0,0) \\
\leq \epsilon \mathbf{E}_{\theta_N}\left[\underbrace{\sum_{i \in N} y_i^{G,S_N}(\theta_N)}_{\leq 1} - \underbrace{\sum_{i \in N} y_i^{G,S_N}(\theta_i^0, \theta_{-i})}_{\geq 0}\right] \leq \epsilon
\end{aligned} \tag{24}$$

B.5 Theorem 2

Given (G, S_N) , let Θ_i^h denote the types of i that are consistent with i 's actions up to history h , that is:

$$\Theta_i^h = \{\theta_i \mid \forall h', h'' \preceq h : [h' \in I_i, h'' \in \sigma(h')] \rightarrow [S_i(I_i, \theta_i) = \mathcal{A}(h'')]\} \tag{25}$$

Proposition 8. *If $h \prec h'$ then $\Theta_i^h \supseteq \Theta_i^{h'}$. If $h \in I_i$ and $h' \in I_i$, then $\Theta_i^h = \Theta_i^{h'}$.*

The first is clear by inspection. The second follows because the definition of Θ_i^h invokes only i 's past information sets and actions, and G has perfect recall. Thus, we define $\Theta_i^{I_i} = \Theta_i^h \mid h \in I_i$.

Define also:

$$\underline{\theta}_i^h = \min\{\theta_i \in \Theta_i^h\} \tag{26}$$

$$\bar{\theta}_i^h = \max\{\theta_i \in \Theta_i^h\} \tag{27}$$

B.5.1 credible, strategy-proof \rightarrow ascending

Definition 16. (G, S_N) has **threshold pricing** if, for some reserve price $\rho \in \Theta_i$ and tie-breaking order \triangleright , if i wins the object at type profile θ_N , then i pays

$$\min\{\theta'_i \in \Theta_i \mid \theta'_i \triangleright \max_{j \neq i} \theta_j \text{ and } \theta'_i \triangleright \rho\} \quad (28)$$

Lemma 1. If (G, S_N) is optimal, orderly, and strategy-proof, then (G, S_N) has threshold pricing.

Proof. We proceed by induction. The inductive hypothesis is: For all $\theta_i \leq \theta_i^k$, if i wins the object, i pays $\min\{\theta'_i \in \Theta_i \mid \theta'_i \triangleright \max_{j \neq i} \theta_j \text{ and } \theta'_i \triangleright \rho\}$.

(G, S_N) is optimal, so IR-0 binds and IC binds locally downward by Proposition 3. Consequently, if $\theta_i = \rho$ and i wins the object, i must pay ρ . Thus, the inductive hypothesis holds for all k such that $\theta_i^k \leq \rho$.

Suppose the inductive hypothesis holds for some k such that $\theta_i^k \geq \rho$. We now show that it holds for $k + 1$. Suppose i wins the object at type profile $(\theta_i^{k+1}, \theta_{-i})$. There are two cases to consider: Either i also wins at $(\theta_i^k, \theta_{-i})$, or i does not. If i wins at $(\theta_i^k, \theta_{-i})$, then since (G, S_N) is strategy-proof, i pays the same amount at both type profiles, namely $\min\{\theta'_i \in \Theta_i \mid \theta'_i \triangleright \max_{j \neq i} \theta_j \text{ and } \theta'_i \triangleright \rho\}$. (Otherwise, when the opponents play as though their types are θ_{-i} , either θ_i^k can profitably deviate to imitate θ_i^{k+1} , or vice versa.)

Suppose i does not win at $(\theta_i^k, \theta_{-i})$. (G, S_N) is optimal, so IC binds locally downward. Type θ_i^{k+1} is exactly indifferent between playing as though his type is θ_i^{k+1} and playing as though his type is θ_i^k whenever $\theta_i^{k+1} \neq \min\{\theta'_i \in \Theta_i \mid \theta'_i \triangleright \max_{j \neq i} \theta_j \text{ and } \theta'_i \triangleright \rho\}$. Thus, for IC to bind locally downward, when $\theta_i^{k+1} = \min\{\theta'_i \in \Theta_i \mid \theta'_i \triangleright \max_{j \neq i} \theta_j \text{ and } \theta'_i \triangleright \rho\}$, i 's payoff must be 0, which means that he must pay θ_i^{k+1} , so the inductive hypothesis holds for $k + 1$. \square

Since F is regular, and (G, S_N) is optimal and orderly, the reserve price must be an optimal reserve price ρ^* .

We define the types of i that will win when facing types θ_{-i} and reserve ρ^* , that is

$$\Theta_i^{\text{win}}(\theta_{-i}, \rho^*) = \{\theta_i \mid \theta_i \triangleright \rho^* \text{ and } \theta_i \triangleright \max_{j \neq i} \theta_j\} \quad (29)$$

Let $\tau_i(\theta_{-i}, \rho)$ denote the price that i pays if he wins when the opponent types are θ_{-i} and the reserve price is ρ , where we set this equal to the highest possible type if i is unable to win. That is,

$$\tau_i(\theta_{-i}, \rho) = \begin{cases} \min\{\theta'_i \in \Theta_i \mid \theta'_i \triangleright \max_{j \neq i} \theta_j \text{ and } \theta'_i \triangleright \rho\} & \text{if } \Theta_i^{\text{win}}(\theta_{-i}, \rho) \neq \emptyset \\ \theta_i^K & \text{otherwise} \end{cases} \quad (30)$$

Lemma 2 (Mutual secrecy). *Assume (G, S_N) is strategy-proof, has threshold pricing, and is orderly with reserve price ρ . For any distinct $i, j \in N$, any $I_j, h \in I_j, \theta_{-i} \in \Theta_{-i}^h$. If $\Theta_i^h \cap \Theta_i^{\text{win}}(\theta_{-i}, \rho) \neq \emptyset$ and $\bar{\theta}_j^h \triangleright \tau_i(\theta_{-i}, \rho)$, then for all $\theta_i \in \Theta_i^{\text{win}}(\theta_{-i}, \rho)$, there exists $h' \in I_j$ such that $(\theta_i, \theta_{-i}) \in \Theta_N^{h'}$.*

Proof. Suppose the antecedent is true but the consequent is not. Pick $\theta_i \in \Theta_i^h \cap \Theta_i^{\text{win}}(\theta_{-i}, \rho)$. Pick $\theta'_i \in \Theta_i^{\text{win}}(\theta_{-i}, \rho)$ such that there does not exist $h' \in I_j$ such that $(\theta'_i, \theta_{-i}) \in \Theta_N^{h'}$. (Trivially, $\theta_i \neq \theta'_i$.) We will define S'_{-i} such that i has a profitable deviation, so (G, S_N) is not strategy-proof.

Suppose for now that $\theta_i \triangleleft \theta'_i$. For $l \in N \setminus \{i, j\}$, S_l specifies that l plays as though his type is θ_l . S_j specifies that j plays as though his type is $\bar{\theta}_j^h$, unless he encounters I_j . Upon encountering I_j , j plays thereafter as though his type is θ_j . i now has a profitable deviation. When i 's type is θ'_i , playing S_i means that j never encounters I_j , so i has payoff $\max\{0, \theta'_i - \tau_i((\bar{\theta}_j^h, \theta_{N \setminus \{i, j\}}), \rho)\}$. On the other hand, if i deviates to play as though his type is θ_i , then since $(\bar{\theta}_j^h, \theta_{N \setminus \{i, j\}}) \in \Theta_{-i}^h$, j encounters I_j with certainty. Thus i 's payoff is $\theta'_i - \tau_i(\theta_{-i}, \rho) > \max\{0, \theta'_i - \tau_i((\bar{\theta}_j^h, \theta_{N \setminus \{i, j\}}), \rho)\}$, where the strict inequality follows since $\theta'_i \triangleright \tau_i(\theta_{-i}, \rho)$ and $\bar{\theta}_j^h \triangleright \tau_i(\theta_{-i}, \rho)$.

Suppose instead that $\theta_i \triangleright \theta'_i$. For $l \in N \setminus \{i, j\}$, S_l specifies that l plays as though his type is θ_l . S_j specifies that j plays as though his type is θ_j , unless he encounters I_j . Upon encountering I_j , j plays thereafter as though his type is $\bar{\theta}_j^h$. When i 's type is θ_i , playing S_i means that j encounters I_j with certainty, so i has payoff $\max\{0, \theta_i - \tau_i((\bar{\theta}_j^h, \theta_{N \setminus \{i, j\}}), \rho)\}$. On the other hand, if i deviates to play as though his type is θ'_i , then j never encounters I_j . Thus i 's payoff is $\theta_i - \tau_i(\theta_{-i}, \rho) > \max\{0, \theta_i - \tau_i((\bar{\theta}_j^h, \theta_{N \setminus \{i, j\}}), \rho)\}$, where the strict inequality follows since $\theta_i \triangleright \tau_i(\theta_{-i}, \rho)$ and $\bar{\theta}_j^h \triangleright \tau_i(\theta_{-i}, \rho)$. \square

Lemma 3 (Potential winners pool). *Assume (G, S_N) is optimal, strategy-proof and credible. For all $I_i, h \in I_i, \theta_i \in \Theta_i^h$, if $\theta_i \triangleright \tau_i(\underline{\theta}_{-i}^h, \rho^*)$ and there exists $j \neq i$ such that $\bar{\theta}_j^h \triangleright \tau_i(\underline{\theta}_{-i}^h, \rho^*)$, then $S_i(I_i, \theta_i) = S_i(I_i, \tau_i(\underline{\theta}_{-i}^h, \rho^*))$.*

Proof. Suppose not. Take some earliest h and I_i such that Lemma 3 does not hold; i.e. h such that for all $h' \prec h$, Lemma 3 holds at h' . We will exhibit a profitable safe deviation that the auctioneer can undertake, upon encountering h while running G . By Lemma 1, (G, S_N) has threshold pricing.

Define $\theta_i^* \equiv \min^{\triangleright} \{\theta_i | \theta_i \triangleright \tau_i(\underline{\theta}_{-i}^h, \rho^*) \text{ and } S_i(I_i, \theta_i) \neq S_i(I_i, \tau_i(\underline{\theta}_{-i}^h, \rho^*))\}$. Note that since Lemma 3 holds for every $h' \prec h$, $\theta_i^* \in \Theta_i^h$. Define h^* to be the immediate successor of h such that $\theta_i^* \in \Theta_i^{h^*}$.

Consider the agents who might still win the object conditional on reaching h , that is, $N^* \equiv \{j \in N | \Theta_j^{\text{win}}(\underline{\theta}_{-j}^h, \rho^*) \cap \Theta_j^h \neq \emptyset\}$. For each agent j , we pick a ‘nemesis’; this is the highest-priority agent in $N^* \setminus j$. That is, $\psi(j) \equiv \max^{\triangleright} N^* \setminus j$.²⁶

²⁶This choice of nemesis deals with boundary cases; it ensures that if there is *any* offer from an agent in N^* that j cannot beat because of the tie-breaking rule, we can pick a nemesis type that j cannot beat.

Since Lemma 3 holds for every $h' \prec h$, for all $j \in N^* \setminus \{i\}$, $\Theta_j^{\text{win}}(\theta_{-j}^h, \rho^*) \subseteq \Theta_j^h$.

Consider, for any type θ , the least upper bound type of j , that is $\min\{\theta_j \mid \theta_j \succeq \theta\}$. We can then define the set (equivalence class) of types such that have the same least upper bound. Formally:

$$\text{CLASS}_j(\theta) \equiv \begin{cases} \{\theta' \mid \min\{\theta_j \mid \theta_j \succeq \theta'\} = \min\{\theta_j \mid \theta_j \succeq \theta\}\} & \text{if } \{\theta_j \mid \theta_j \succeq \theta\} \neq \emptyset \\ \{\theta' \mid \theta' \succ \theta_j^K\} & \text{otherwise} \end{cases}. \quad (31)$$

Given S_0^G (with corresponding λ), we now exhibit a (partial) behavioral strategy that deviates from S_0^G upon encountering h^* and is strictly profitable. We describe this algorithmically. (The description is lengthy, because it must produce a safe deviation for any extensive game form in a large class.)

The algorithm is divided into four parts; **First-Query**, **Safety**, **Query**, and **Conclude**. Essentially, the auctioneer first pauses communication with $N \setminus i$ and checks whether i 's type is at least θ_i^* (**First-Query**). If i 's type is revealed to be less than θ_i^* , that implies that i would not win under S_0^G . In that case, the auctioneer communicates with $N \setminus i$ in a way that ensures her payoff is at least as high as under S_0^G (**Safety**). If i 's type is revealed to be at least θ_i^* , the auctioneer then treats θ_i^* as her 'best offer', and communicates with agents in turn seeing if any agent can beat the best offer, updating the best offer whenever she finds a better offer (**Query**). She continues this until the agent with the best offer is the only agent left, whereupon she sells to that agent at a price equal to the best offer (**Conclude**).

This deviation always results in revenue at least as high as under S_0^G , and with positive probability results in strictly more revenue - in particular, when the type profile is $(\theta_i^*, \theta_{-i}^{h^*})$. The key is to write this down precisely, and to show that the deviation can be carried out safely.

We use the notation $:=$ for the assignment operator, and we use $:\in$ to assign an (arbitrary) element from the set on the right-hand side.

This algorithm keeps track of:

1. A best offer, initialized $\beta := \theta_i^*$.
2. A set of 'active' agents, initialized $\hat{N} := N$.
3. A current agent called to play, $\hat{j} := i$.
4. A simulated history, $\hat{h} := h^*$.
5. The last information set each agent saw, if it exists, $(\hat{I}_j)_{j \in N}$.

$$\hat{I}_j := I_j \prec h^* \mid \forall I'_j \prec h^* : I'_j \preceq I_j \quad (32)$$

6. The last action each agent took, $(\hat{a}_j)_{j \in N}$, initialized to be the action taken at \hat{I}_j .

We start at Step 1 of **First-Query**.

First-Query

1. If $\underline{\theta}_i^{\hat{h}} \geq \beta$, then:

(a) $\beta := \underline{\theta}_i^{\hat{h}}$

(b) Assign (\hat{I}_i, \hat{a}_i) to be the latest information set and action that i encountered.

(c) $\hat{j} := \hat{N} \setminus i$.

(d) If $\hat{a}_{\hat{j}} \neq \emptyset$, then for $\theta_{\psi(\hat{j})} \in \text{CLASS}_{\hat{j}}(\beta)$,

$$\hat{h} := h \mid h \in \sigma(\hat{I}_{\hat{j}}) \text{ and } \mathcal{A}(h) = \hat{a}_{\hat{j}} \text{ and } (\theta_{\psi(\hat{j})}, \underline{\theta}_{N \setminus \{\hat{j}, \psi(\hat{j})\}}^{h*}) \in \Theta_{-\hat{j}}^h \quad (33)$$

(e) Else $\hat{h} := h_0$.

(f) Go to **Query**.

2. Else if $\hat{h} \in Z$, then $\beta := \bar{\theta}_i^{\hat{h}}$, $\hat{h} := h^*$, and go to **Safety**.²⁷

3. Else if $P(\hat{h}) = i$

(a) Send $\lambda(I_i, A(I_i))$ for $I_i \mid \hat{h} \in I_i$ to i .

(b) Upon receiving r , set $\hat{h} := h' \mid h' \in \sigma(\hat{h})$ and $\mathcal{A}(h') = \lambda^{-1}(r)$.

(c) Go to step 1.

4. Else:

(a) For $\theta_{\psi(i)} \in \text{CLASS}_i(\beta)$

$$\hat{h} := h \mid h \in \sigma(\hat{h}) \text{ and } (\theta_{\psi(i)}, \underline{\theta}_{N \setminus \{i, \psi(i)\}}^{h*}) \in \Theta_{-i}^h \quad (34)$$

(b) Go to step 1.

Query

1. If $|\hat{N}| = 1$, go to **Conclude**.

2. If $\underline{\theta}_{\hat{j}}^{\hat{h}} \geq \beta$, then:

(a) $\beta := \underline{\theta}_{\hat{j}}^{\hat{h}}$

(b) Assign $(\hat{I}_{\hat{j}}, \hat{a}_{\hat{j}})$ to be the latest information set and action that \hat{j} encountered.

²⁷Here we use β to keep track of the highest type consistent with i 's actions, which we have learned is not enough to win.

(c) $\hat{j} := \hat{N} \setminus \hat{j}$.

(d) If $\hat{a}_{\hat{j}} \neq \emptyset$, then for $\theta_{\psi(\hat{j})} \in \text{CLASS}_{\hat{j}}(\beta)$,

$$\hat{h} := h \mid h \in \sigma(\hat{I}_{\hat{j}}) \text{ and } \mathcal{A}(h) = \hat{a}_{\hat{j}} \text{ and } (\theta_{\psi(\hat{j})}, \theta_{\underline{N} \setminus \{\hat{j}, \psi(\hat{j})\}}^{h*}) \in \Theta_{-\hat{j}}^h \quad (35)$$

(e) Else $\hat{h} := h_0$.

(f) Go to step 1.

3. Else if $\hat{h} \in Z$, then:

(a) $\hat{N} := \hat{N} \setminus \hat{j}$.

(b) $\hat{j} := \hat{N}$.

(c) If $\hat{a}_{\hat{j}} \neq \emptyset$, then for $\theta_{\psi(\hat{j})} \in \text{CLASS}_{\hat{j}}(\beta)$,

$$\hat{h} := h \mid h \in \sigma(\hat{I}_{\hat{j}}) \text{ and } \mathcal{A}(h) = \hat{a}_{\hat{j}} \text{ and } (\theta_{\psi(\hat{j})}, \theta_{\underline{N} \setminus \{\hat{j}, \psi(\hat{j})\}}^{h*}) \in \Theta_{-\hat{j}}^h \quad (36)$$

(d) Else $\hat{h} := h_0$.

(e) Go to step 1.

4. Else if $P(\hat{h}) = \hat{j}$, then:

(a) Send $\lambda(I_{\hat{j}}, A(I_{\hat{j}}))$ for $I_{\hat{j}} \mid \hat{h} \in I_{\hat{j}}$ to \hat{j} .

(b) Upon receiving r , set $\hat{h} := h' \mid h' \in \sigma(\hat{h})$ and $\mathcal{A}(h') = \lambda^{-1}(r)$.

(c) Go to step 1.

5. Else:

(a) For $\theta_{\psi(\hat{j})} \in \text{CLASS}_{\hat{j}}(\beta)$

$$\hat{h} := h \mid h \in \sigma(\hat{h}) \text{ and } (\theta_{\psi(\hat{j})}, \theta_{\underline{N} \setminus \{\hat{j}, \psi(\hat{j})\}}^{h*}) \in \Theta_{-\hat{j}}^h \quad (37)$$

(b) Go to step 1.

Safety

1. If $\hat{h} \in Z$, then terminate and choose $x = g(\hat{h})$.

2. Else if $P(\hat{h}) = i$:

(a) $\hat{h} := h \mid h \in \sigma(\hat{h})$ and $\beta \in \Theta_i^h$.

(b) Go to step 1.

3. Else:

- (a) Choose agent $P(\hat{h})$ and send $(m, R) = \lambda(I_j, A(I_j))$ for I_j such that $\hat{h} \in I_j$.
- (b) Upon receiving $r \in R$, choose $\hat{h} := h' \mid \mathcal{A}(h') = \lambda^{-1}(r)$ and $h' \in \sigma(\hat{h})$.
- (c) Go to step 1.

Conclude

1. If $\hat{h} \in Z$, then terminate and choose $x = g(\hat{h})$.
2. Else if $P(\hat{h}) = \hat{j}$, then:
 - (a) Send $\lambda(I_{\hat{j}}, A(I_{\hat{j}}))$ for $I_{\hat{j}} \mid \hat{h} \in I_{\hat{j}}$ to \hat{j} .
 - (b) Upon receiving r , set $\hat{h} := h' \mid h' \in \sigma(\hat{h})$ and $\mathcal{A}(h') = \lambda^{-1}(r)$.
 - (c) Go to step 1.
3. Else:
 - (a) For $\theta_{\psi(\hat{j})} \in \text{CLASS}_{\hat{j}}(\beta)$

$$\hat{h} := h \mid h \in \sigma(\hat{h}) \text{ and } (\theta_{\psi(\hat{j})}, \theta_{N \setminus \{\hat{j}, \psi(\hat{j})\}}^{h*}) \in \Theta_{-\hat{j}}^h \quad (38)$$

- (b) Go to step 1.

The algorithm terminates in two possible ways. Either it goes from **First-Query** to **Safety**, or it goes from **First-Query** to **Query** to **Conclude**.

If the algorithm goes from **First-Query** to **Safety**, then agent i has (by construction) observed a communication sequence consistent with some terminal history z at which he does not win. All the other agents observe a communication sequence consistent with i playing as though his type is $\bar{\theta}_i^z$. Since $h^* \prec z$, Proposition 8 implies that $\bar{\theta}_i^z \in \Theta_i^{h^*}$, so every agent's observation has an innocent explanation. Moreover, i 's true type is no more than $\bar{\theta}_i^z$, which is strictly less than the least type i would need to win at the non-pooling history where we started. Thus, if the algorithm goes from **First-Query** to **Safety**, revenue is at least as high as it would have been under S_0^G .

If the algorithm goes from **First-Query** to **Query** to **Conclude**, notice that whenever we initiate communication with an agent, if he has seen any communication before, then we pick a simulated history that is among the immediate successors of the last information set he saw, and consistent with the last action that he took (Step 1.d of **First-Query** and steps 2.d and 3.c of **Query**). Moreover, whenever an agent is removed from the set \hat{N} during **Query**, he has already seen a communication sequence consistent with a terminal history at which he does not win (Step 3). Finally, **Conclude** ensures that the agent who does win the object sees a communication sequence consistent with that happening.

It remains to show that we can in fact pick simulated histories in this way; *i.e.* that Step 1.d of **First-Query** and steps 2.d and 3.c of **Query** are well-defined. Consider

two cases; the first time the deviating algorithm communicates with an agent, and all subsequent times.

The first time the deviating algorithm communicates with an agent \hat{j} ,²⁸ we need to pick a history consistent with $(\theta_{\psi(\hat{j})}, \underline{\theta}_{N \setminus \{\hat{j}, \psi(\hat{j})\}}^{h^*})$ for $\theta_{\psi(\hat{j})} \in \text{CLASS}_{\hat{j}}(\beta)$. If \hat{j} has not seen any information sets yet, then h_0 will do. Suppose \hat{j} was called to play at some history h' before we started deviating at h . By hypothesis, Lemma 3 holds at h' . Moreover, for all $l \in N^* \setminus \{\hat{j}\}$, $\Theta_l^{\text{win}}(\underline{\theta}_{-l}, \rho^*) \subseteq \Theta_l^h \subseteq \Theta_l^{h'}$, where the second set inclusion is by Proposition 8. In particular, since \hat{j} 's nemesis $\psi(\hat{j})$ was picked from $N^* \setminus \{\hat{j}\}$, and $\beta \triangleright \theta_i^* \triangleright \underline{\theta}_{\psi(\hat{j})}^{h'}$, this implies that for $\theta_{\psi(\hat{j})} \in \text{CLASS}_{\hat{j}}(\beta)$, $(\theta_{\psi(\hat{j})}, \underline{\theta}_{N \setminus \{\hat{j}, \psi(\hat{j})\}}^{h^*}) \in \Theta_{-\hat{j}}^{h'}$. Thus, Step 1.d of **First-Query** and steps 2.d and 3.c of **Query** are well-defined for this case.

Consider any subsequent time that the deviating algorithm initiates communication with agent \hat{j} . By construction of **First-Query** and **Query**, this implies that the agent has previously been queried and beaten the (then) best offer β^{old} . We have to show that we can pick a history that immediately succeeds $\hat{I}_{\hat{j}}$ that is consistent with $(\theta_{\psi(\hat{j})}, \underline{\theta}_{N \setminus \{\hat{j}, \psi(\hat{j})\}}^{h^*})$ for $\theta_{\psi(\hat{j})} \in \text{CLASS}_{\hat{j}}(\beta)$, for $\beta \triangleright \beta^{\text{old}}$.

When the auctioneer last communicated with \hat{j} , she reached a simulated history $h^{\text{old}} \in \hat{I}_{\hat{j}}$, where $(\theta_{\psi(\hat{j})}^{\text{old}}, \underline{\theta}_{N \setminus \{\hat{j}, \psi(\hat{j})\}}^{h^*}) \in \Theta_{-\hat{j}}^{h^{\text{old}}}$ for $\theta_{\psi(\hat{j})}^{\text{old}} \in \text{CLASS}_{\hat{j}}(\beta^{\text{old}})$. This is the simulated history at which \hat{j} was last called to play, and then took an action that revealed that his type beat $\theta_{\psi(\hat{j})}^{\text{old}}$.

Since at h^{old} , \hat{j} has *not yet* beaten $\theta_{\psi(\hat{j})}^{\text{old}}$, it follows that

$$\Theta_{\psi(\hat{j})}^{h^{\text{old}}} \cap \Theta_{\psi(\hat{j})}^{\text{win}}((\underline{\theta}_{\hat{j}}^{h^{\text{old}}}, \underline{\theta}_{N \setminus \{\hat{j}, \psi(\hat{j})\}}^{h^*}), \rho^*) \neq \emptyset \quad (39)$$

Since \hat{j} 's action at $\hat{I}_{\hat{j}}$ revealed that his type beats $\theta_{\psi(\hat{j})}^{\text{old}}$, it follows that

$$\bar{\theta}_{\hat{j}}^{h^{\text{old}}} \triangleright \tau_{\psi(\hat{j})}((\underline{\theta}_{\hat{j}}^{h^{\text{old}}}, \underline{\theta}_{N \setminus \{\hat{j}, \psi(\hat{j})\}}^{h^*}), \rho^*) \quad (40)$$

Since the new simulated type for $\psi(\hat{j})$ is at least as high as the old simulated type, by Lemma 2 it follows from Equations 39 and 40 that we can pick an alternative history in $\hat{I}_{\hat{j}}$ that is consistent with the new simulated type. That is, since $\theta_{\psi(\hat{j})} \triangleright \theta_{\psi(\hat{j})}^{\text{old}}$, it follows that $\theta_{\psi(\hat{j})} \in \Theta_{\psi(\hat{j})}^{\text{win}}((\underline{\theta}_{\hat{j}}^{h^{\text{old}}}, \underline{\theta}_{N \setminus \{\hat{j}, \psi(\hat{j})\}}^{h^*}), \rho^*)$, so by Lemma 2, there exists $h^{\text{new}} \in \hat{I}_{\hat{j}}$ such that $(\theta_{\psi(\hat{j})}, \underline{\theta}_{\hat{j}}^{h^{\text{old}}}, \underline{\theta}_{N \setminus \{\hat{j}, \psi(\hat{j})\}}^{h^*}) \in \Theta_N^{h^{\text{new}}}$. Thus, we can pick \hat{h} appropriately by defining it to be the immediate successor of h^{new} that is consistent with \hat{j} 's last action, i.e. $\hat{h} := h'' \in \sigma(h^{\text{new}}) \mid \mathcal{A}(h'') = \hat{a}_{\hat{j}}$. Thus, Step 1.d of **First-Query** and steps 2.d and 3.c of **Query** are well-defined.

Observe, finally, that if the algorithm goes from **First-Query** to **Query** to **Conclude**, revenue is at least equal to revenue under S_0^G , and is strictly higher whenever θ_i^* is greater

²⁸That is, the first time that we perform Step 1.d of **First-Query** or steps 2.d or 3.c of **Query** for that agent.

than revenue under S_0^G . Thus, the deviation is safe and results in strictly higher expected profit, and (G, S_N) is not credible. This completes the proof of Lemma 3. \square

Moving from Lemma 3 to Theorem 2 is mostly an exercise in labeling. Bidder i is **active** at h if $\Theta_i^h \cap \Theta_i^{\text{win}}(\underline{\theta}_{-i}^h, \rho^*) \neq \emptyset$.

There are three cases to consider:

1. An active bidder is called to play, and there is more than one active bidder.
2. An inactive bidder is called to play.
3. An active bidder is called to play, and there are no other active bidders. (We leave this case till last - it can only happen when every other bidder has a type below the reserve.)

Take any I_i and $h \in I_i$ such that an active bidder i is called to play. If there exists another active bidder, then there exists $j \neq i$ such that $\bar{\theta}_j^h \triangleright \tau_i(\underline{\theta}_{-i}^h, \rho^*)$. There is some action $S_i(I_i, \theta_i^K)$ that is taken by the highest type of i . Lemma 3 implies that all the types of i that might still win at h , *i.e.* $\{\theta_i \mid \theta_i \geq \tau_i(\underline{\theta}_{-i}^h, \rho^*)\}$, must also play $S_i(I_i, \theta_i^K)$. Thus, any agent who does not play that action has **quit**. The **bid** at I_i is the least type of i consistent with playing $S_i(I_i, \theta_i^K)$, that is $\min^{\triangleright} \{\theta_i \in \Theta_i^h \mid S_i(I_i, \theta_i) = S_i(I_i, \theta_i^K)\}$.²⁹ By Proposition 8, each bid is weakly more than the last bid that i placed.

By construction, all types strictly below the bid quit. Since (G, S_N) is optimal and has threshold pricing, if there is no high bidder, then all types weakly above ρ^* place a bid. Similarly, all types above the current high bid place a bid.

If bidder i quits, then he either has a type lower than the reserve, or we have identified another bidder whose type is greater than i 's (according to the order \triangleright). Thus, once i is inactive, further information about his type no longer affects the outcome, so since (G, S_N) is pruned, only active bidders are called to play. Similarly, if i is the current high bidder at history h , and the auction has not ended, then by Lemma 3, all his types who reach h take the same action, and by (G, S_N) pruned, i is not called to play at h . Thus, if i is called to play at h , he is an active bidder who is not the current high bidder.

Suppose an active bidder i is called to play at h and is the unique active bidder. Since (G, S_N) is pruned, i is not the current high bidder, which implies that there is *no* high bidder - all the other bidders have types below the reserve. Let I_i be such that $h \in I_i$.

In this case, we can define an action a as **quitting** if there is no type above the reserve that plays a , that is:

$$\neg \exists \theta_i \in \Theta_i^h \mid \theta_i \triangleright \rho^* \text{ and } S_i(I_i, \theta_i) = a \quad (41)$$

²⁹Note that this may be below $\tau_i(\underline{\theta}_{-i}^h, \rho^*)$.

For any non-quitting action a , the associated **bid** is:

$$\min^{\triangleright}\{\theta_i \mid \theta_i \triangleright \rho^* \text{ or } [\theta_i \in \Theta_i^h \text{ and } S_i(I_i, \theta_i) = a]\} \quad (42)$$

By construction, if i has a type strictly below the bid associated with a , then he does not play a . If i has a type above the reserve, then he places a bid. However, Lemma 3 does not apply (since there are no other active bidders), so there can be multiple actions that place bids. Again, by Proposition 8, each bid is weakly more than the last bid that i placed.

The three conditions that specify what happens when the auction ends are similarly entailed by optimality and threshold pricing. If there are no active bidders at h , then for all i , $\rho^* \triangleright \bar{\theta}_i^h$. Thus, the object is not sold, and since (G, S_N) is pruned, h is a terminal history. If the high bidder i is the unique active bidder at h , then we know that no bidder in $N \setminus i$ has a higher type than i , and that $\tau_i(\theta_{-i}, \rho^*)$ is equal to i 's current bid. Thus, i must win and pay b , and since (G, S_N) is pruned, h is a terminal history. Finally, if the high bidder has bid θ^K and no active bidder has higher tie-breaking priority, then i must win and pay θ^K , and since (G, S_N) is pruned, h is a terminal history.

This completes the proof that if an orderly optimal protocol is strategy-proof and credible, then it is an ascending auction.

B.5.2 ascending \rightarrow credible, strategy-proof

Now we show that if (G, S_N) is orderly, optimal, and an ascending auction, it is credible and strategy-proof.

That (G, S_N) is strategy-proof is straightforward. It remains to show that (G, S_N) is credible. Observe that for all i , S_i is an *obviously dominant* strategy, in the sense of Li (2017). In particular, for any safe deviation $S'_0 \in \mathcal{S}_0^*(S_0^G, S_N)$ and for any S'_{-i} , S_i is a best response to (S'_0, S'_{-i}) in the messaging game. We offer a direct proof here:

First, consider information sets at which there is a unique action that places a bid. Take any i , I_i , and θ_i such that $\theta_i \in \Theta_i^{I_i}$. Recall that S_i requires that i quit if θ_i is strictly below the bid $b(I_i)$ at I_i , and that i places the bid if θ_i is above the least high bid consistent with reaching I_i . The least high bid consistent with reaching I_i is, formally,

$$\min^{\triangleright}\{\rho^*, \min_{h \in I_i, j \neq i}^{\triangleright} \theta_j^h\} \quad (43)$$

And, since (G, S_N) is optimal and has threshold pricing,

$$b(I_i) \leq \min^{\triangleright}\{\theta'_i \mid \theta'_i \triangleright \min^{\triangleright}\{\rho^*, \min_{h \in I_i, j \neq i}^{\triangleright} \theta_j^h\}\} \quad (44)$$

For any safe deviation S'_0 and for any S'_{-i} , it is optimal for i to quit (upon reaching

information set I_i) if $\theta_i \triangleleft \min\{\rho^*, \min_{h \in I_i, j \neq i} \theta_j^h\}$. In particular, note that under (G, S_N) , if i wins after reaching I_i , he pays at least $\min\{\rho^*, \min_{h \in I_i, j \neq i} \theta_j^h\}$. Thus, for any safe deviation, i 's best possible payoff upon placing a bid is no more than zero, so it is optimal to quit (which yields zero payoff).

For any safe deviation S'_0 and for any S'_{-i} , it is optimal for i to place a bid if θ_i is weakly above that bid. This is because i can quit if the bids required ever rise strictly above θ_i . Under any safe deviation, i cannot be charged more than θ_i unless he (at some later point) bids more than θ_i . Thus, the worst possible payoff from placing a bid is zero, and the best possible payoff from quitting is zero.

By the above arguments and Equation 44, there are three possibilities at each I_i and $\theta_i \in \Theta_i^{I_i}$:

1. $\min\{\rho^*, \min_{h \in I_i, j \neq i} \theta_j^h\} \triangleleft \theta_i$, in which case S_i requires that i place a bid, and this is a best response to (S'_0, S'_{-i}) .
2. $\theta_i \triangleleft b(I_i)$, in which case S_i requires i to quit, and this is a best response to (S'_0, S'_{-i}) .
3. $b(I_i) \trianglelefteq \theta_i \triangleleft \min\{\rho^*, \min_{h \in I_i, j \neq i} \theta_j^h\}$, in which case S_i is underdetermined, and both quitting now or placing the bid and quitting later are best responses to (S'_0, S'_{-i}) .

Finally, consider information sets at which there are multiple bid-placing actions. In this case, under any safe deviation, i is sure to win if and only if he eventually bids above the reserve - this implies that S_i remains a best response to any safe deviation.

Suppose now that (G, S_N) is not credible, so the auctioneer has a profitable safe deviation S'_0 . Consider a corresponding G' in which the auctioneer 'commits openly' to that deviation, that is to say, G' such that S'_0 runs G' . For all i , S_i is a best response to (S'_0, S_{-i}) , so (G', S_N) is also BIC. (We abuse notation slightly to use S_N as a strategy profile for G and G' . Every information set in G' has a corresponding information set in G , so it is clear what is meant.) By hypothesis, S'_0 is a profitable deviation, so $\pi(G', S_N) > \pi(G, S_N)$. But (G, S_N) is optimal, so this is a contradiction. Thus, if (G, S_N) is orderly, optimal, and an ascending auction, (G, S_N) is credible.

B.6 Proposition 7

By the same arguments as in Proposition 6, in a credible static auction there is a unique payment associated with each action. Proposition 5 generalizes easily to the asymmetric case:

Proposition 9. *If (G, S_N) is BIC, then:*

$$0 \leq \mathbf{E}_{\theta_N} \left[\sum_{i \in N} y_i^{G, S_N}(\theta_N) \eta_i(\theta_i) \right] - \pi(G, S_N) - \sum_{i \in N} \check{u}_i^{G, S_N}(0, 0) \leq \epsilon \quad (45)$$

Turning to Example 2, suppose $\epsilon \leq \frac{1}{16}$ and consider the event E' that $\theta_1 \in (\frac{14}{16}, \frac{15}{16}]$ and $\theta_2 \in (\frac{18}{16}, \frac{19}{16}]$. In this event, $\eta_1(\theta_1) \geq \frac{3}{4} > \frac{3}{8} \geq \eta_2(\theta_2)$, so any optimal mechanism assigns the object to bidder 1 with certainty. Take any ϵ -optimal credible static (G, S_N) . If bidder 2 never wins when E' occurs, then the virtual value of the winner under (G, S_N) is bounded away from the optimum by $\frac{1}{16^2}(\frac{3}{4} - \frac{3}{8})$. Thus, by Proposition 9, for small enough ϵ , there exist $\theta'_1 \in (\frac{14}{16}, \frac{15}{16}]$ and $\theta'_2 \in (\frac{18}{16}, \frac{19}{16}]$ such that 1 wins the object in (G, S_N) when the type profile is (θ'_1, θ'_2) . We next establish that the auctioneer has a profitable safe deviation; she should sell the object to bidder 2 at this type profile.

Consider the event E'' that $\theta_1 \in (\frac{15}{16}, 1]$ and $\theta_2 \in (\frac{31}{16}, 2]$, so $\eta_1(\theta_1) \leq 1 < \frac{15}{8} \leq \eta_2(\theta_2)$. If bidder 1 never loses when E'' occurs, then the virtual value of the winner under (G, S_N) is bounded away from the optimum by $\frac{1}{16^2}(\frac{15}{8} - 1)$. By Proposition 9, for small enough ϵ , that there exists $\theta''_1 \in (\frac{15}{16}, 1]$ such that 1 loses with positive probability when his type is θ''_1 . Since (G, S_N) is BIC, each agent's probability of winning is non-decreasing in his type, so there exists θ''_2 such that 1 does not win when the type profile is (θ''_1, θ''_2) .

With probability $\frac{1}{2}$, $\theta_1 \leq \frac{1}{2}$, which implies $\eta_1(\theta_1) \leq 0$. With probability $\frac{1}{16}$, $\theta_2 \in (\frac{17}{16}, \frac{18}{16}]$, which implies $\frac{1}{8} \leq \eta_2(\theta_2)$. Again invoking Proposition 9, for small enough ϵ , there exists $\theta'''_2 \in (\frac{17}{16}, \frac{18}{16}]$ that wins with probability at least $\frac{1}{4}$, because otherwise the virtual value of the winner is bounded away from the optimum by $\frac{1}{4} \cdot \frac{1}{16} \cdot \frac{1}{8}$. Thus, by BIC, θ'_2 wins with probability at least $\frac{1}{4}$, which implies that there exists θ'''_1 such that bidder 2 wins when the type profile is (θ'''_1, θ'_2) .

Let $b_i(\theta_i)$ denote the (unique) payment made by bidder i with type θ_i , if he wins under (G, S_N) . By interim individual rationality, $b_1(\theta'_1) \leq \frac{15}{16}$. We assert that $b_2(\theta'_2) > \frac{15}{16}$. Suppose not. (G, S_N) is BIC, so the lowest type of bidder 2 does not want to imitate θ'_2 , so $\check{u}_2^{G, S_N}(0, 0) \geq \frac{1}{4}(1 + \epsilon - b_2(\theta'_2)) \geq \frac{1}{4}(1 - \frac{15}{16})$. By Proposition 9, for small enough ϵ , (G, S_N) is not ϵ -optimal, a contradiction.

We have established that (G, S_N) requires bidder 1 to win at type profile (θ'_1, θ'_2) , but $b_1(\theta'_1) < b_2(\theta'_2)$ and the auctioneer can safely deviate to sell the object to bidder 2 at $b_2(\theta'_2)$. Thus, (G, S_N) is not credible, a contradiction.