

CREDIBLE AUCTIONS: A TRILEMMA

MOHAMMAD AKBARPOUR
Graduate School of Business, Stanford University

SHENGWU LI
Society of Fellows, Harvard University

Consider an extensive-form mechanism, run by an auctioneer who communicates sequentially and privately with bidders. Suppose the auctioneer can deviate from the rules provided that no single bidder detects the deviation. A mechanism is *credible* if it is incentive-compatible for the auctioneer to follow the rules. We study the optimal auctions in which only winners pay, under symmetric independent private values. 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.

1. INTRODUCTION

THE STANDARD MECHANISM DESIGN PARADIGM assumes that the auctioneer has full commitment. She binds herself to follow the rules, and cannot deviate after observing the bids, even when it is profitable *ex post* to renege (McAfee and McMillan (1987)). This contrasts starkly with the way we model participants; incentive compatibility “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” (Hurwicz (1972)).

In this paper, we study incentive compatibility for the auctioneer. We require that the auctioneer, having promised in advance to abide by certain rules, should not find it profitable to “cheat,” where cheating is defined as behavior that can be made to look “legal” to each participant by misrepresenting the preferences of the other participants. For instance, in a second-price auction, the auctioneer can profit by exaggerating the second-highest bid. Thus, as Vickrey (1961) observed, the first-price auction is “automatically self-policing,” while the second-price auction requires special arrangements that tie the auctioneer’s hands.¹

To proceed, we must choose a communication structure for the bigger game played by the bidders and the auctioneer. Clearly, if the bidders simultaneously and publicly announce their bids, then the problem is trivial, and reduces to the case of full commitment. However, such announcements are uncommon in real-world auctions. Most bidders at high-stakes auction houses do not place bids audibly, and instead use secret signals that

Mohammad Akbarpour: mohamwad@stanford.edu

Shengwu Li: shengwu_li@fas.harvard.edu

This paper has been previously distributed with the title “Credible Mechanisms.” We thank Ben Brooks, Jeremy Bulow, Sylvain Chassang, Alex Frankel, Ben Golub, Matthew Jackson, Emir Kamenica, Paul Klemperer, 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. We also thank multiple anonymous referees for their constructive feedback. All errors remain our own.

¹Rothkopf, Teisberg, and Kahn (1990) argued that some real-world auctioneers avoid second-price auctions because bidders fear that the auctioneer may cheat.

other bidders cannot detect. These signals “may be in the form of a wink, a nod, scratching an ear, lifting a pencil, tugging the coat of the auctioneer, or even staring into the auctioneer’s eyes—all of them perfectly legal” (Cassady (1967)). Many bidders are not present in the auction room at all, but instead bid over the Internet or by telephone.² Christie’s and Sotheby’s are legally permitted to call out fake (‘chandelier’) bids to give the impression of higher demand; the New York Times reports that, because of this practice, “bidders have no way of knowing which offers are real.”³ An industry newsletter for online advertising auctions reports:⁴

In a second-price auction, raising the price floors after the bids come in allows [online auctioneers] to make extra cash off unsuspecting buyers [...] This practice persists because neither the publisher nor the ad buyer has complete access to all the data involved in the transaction, so unless they get together and compare their data, publishers and buyers won’t know for sure who their vendor is ripping off.

To formalize these opportunities for rule-breaking, we assume that the auctioneer communicates privately with each bidder. This allows the auctioneer to misrepresent any bidder’s preferences to any other bidder. Of course, many real-world mechanisms have some public communication. For instance, in typical auctions for art or wine, the auctioneer reveals the clearing price but hides the identity of the winner (Ashenfelter (1989)).⁵ As another example, the U.S. National Resident Matching Program publishes aggregate statistics about the match, but does not publish information that identifies individual doctors or hospitals.⁶ In general, who shares information, and what information they share, depends on context-specific features that are outside our framework. Fully private communication is a tractable benchmark, so it is a natural place to start.⁷

Consider any *protocol*: a pair consisting of an extensive-form mechanism and a strategy profile for the bidders. The auctioneer runs the mechanism as follows: Starting from the initial history, she picks up the telephone and conveys a message to the bidder who is called to play (an information set), along with a set of acceptable replies (actions). The bidder chooses a reply. The auctioneer keeps making telephone calls, sending messages 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 bidder intrinsically observes certain features of the outcome. For instance, each bidder observes whether or not he wins the object, and how much he pays, but not how much other bidders pay.

By participating in the protocol, each bidder 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, bidder i ’s observation could still have an *innocent explanation*. That is, when the auctioneer plays by the rules, there exist types for the other bidders that result in that same observation for i .

²The Wall Street Journal reports, “Many auction rooms are sparsely attended these days despite widespread interest in the items being sold, with most bids coming in online or, even more commonly, by phone.” *Why auction rooms seem empty these days*, The Wall Street Journal, June 15, 2014.

³*Genteel auction houses turning aggressive*, The New York Times, April 24, 2000.

⁴*How SSPs use deceptive price floors to squeeze ad buyers*, Digiday, Sep 13, 2017

⁵One justification for keeping the winner’s identity secret is that it gives bidders incentives to defect from collusive arrangements. However, publishing the clearing price does not rule out cheating, since each losing bidder may believe that someone else placed the second-highest bid.

⁶The match rules also limit the information that participants can share. 2019 Main Residency Match Participation Agreement for Applicants and Programs, Sections 4.4 and 4.6.

⁷A fuller analysis would model the auctioneer’s reasons for privacy and disclosure, such as post-auction strategic interaction (Dworczak (2020)).

Given a protocol, some deviations may be *safe*, in the sense that for every type profile, each bidder's observation has an innocent explanation. That is, every observation that a bidder 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.⁸

Instead of just choosing a different outcome, the auctioneer may also alter the way she communicates with bidders. 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.

Having defined the framework, we now consider how credibility interacts with other design features. 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), Edelman, Ostrovsky, and Schwarz (2007)).⁹ The first-price auction is static (“sealed-bid”)—each bidder 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.¹⁰ 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)).

We study the implications of credibility in the independent private values (IPV) model (Myerson (1981)). For now, we assume that the value distributions are regular and symmetric, and restrict attention to auctions in which only winning bidders make (or receive) transfers. Under these assumptions, the second-price auction (with reserve) is the unique strategy-proof static optimal auction (Green and Laffont (1977), Holmström (1979), Milgrom and Segal (2002)). The second-price auction is not credible, so no optimal auction is strategy-proof, static, and credible. This raises two natural questions: Is any auction static and credible? Is any auction strategy-proof and credible?

Our first result is as follows: The first-price auction (with reserve) is the unique static credible optimal auction. This implies that, in the class of static mechanisms, we must

⁸An 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))

⁹The Dutch (descending) auction, in which the price falls until one bidder claims the object, is less prevalent (Krishna (2010, p. 2)).

¹⁰Using data from U.S. Forest Service timber auctions, Athey, Levin, and Seira (2011) found that “sealed bid auctions attract more small bidders, shift the allocation toward these bidders, and can also generate higher revenue.”

choose between incentive-compatibility for the auctioneer and dominant strategies for the bidders.

Static mechanisms include the direct revelation mechanisms, in which each bidder simply reports his type. Thus, when designing credible protocols, restricting attention to revelation mechanisms loses generality. The problem is that revelation mechanisms reveal too much. 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. What happens when we instead allow arbitrary communication protocols—when we use the full richness of extensive forms to regulate who knows what, and when?

For the next result, we discretize the type space, so that optimal clock auctions can be represented as extensive-form mechanisms.

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

Notably, this result does not use open outcry bidding to ensure good behavior by the auctioneer. Given an ascending auction with an optimal reserve, the auctioneer prefers to follow the rules even though she communicates with each bidder individually by telephone. If the auctioneer places chandelier bids, then she runs the risk that bidders will quit. In equilibrium, this deters her from placing chandelier bids at any price above the reserve.

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.

Next, we generalize these results by relaxing the assumption that only winners make transfers and that the distributions are symmetric. The credible static auctions are now *twin-bid* auctions. This is a larger class that includes all-pay auctions and first-price auctions with entry fees. In a twin-bid auction, each bidder chooses from a set of feasible bids, where a bid is a pair of numbers specifying what he pays if he wins and what he pays if he loses. After taking all bids, the auctioneer chooses a winner that maximizes revenue. Under mild assumptions, twin-bid auctions are not strategy-proof.

Under asymmetry, the static strategy-proof optimal auctions are *virtual* second-price auctions: each bid is scored as its corresponding virtual value, and the winner pays the




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.

least bid he could have reported while still having the highest score. Correspondingly, the credible strategy-proof optimal auctions are *virtual* ascending auctions: bids are scored according to their virtual values, so one bidder's price may rise faster than another's. Thus, general extensive forms enable the auctioneer to credibly reject higher bids in favor of lower bids, when it is optimal to do so.

For practical purposes, should an auction be static, strategy-proof, or credible? It depends. Some Internet advertising auctions must be conducted in milliseconds, so latency precludes the use of multi-round protocols. Strategy-proofness matters when bidders are inexperienced or have opportunities for rent-seeking espionage. Credibility matters especially when bidders are anonymous to each other or require that their bids be kept private. These real-world concerns are outside the model. Our purpose is not to elevate some criterion as essential, but to investigate which combinations are possible.

1.1. *Related Work*

We are far from the first to conceive of games of imperfect information as being conducted by a central mediator under private communication. Von Neumann and Morgenstern explicated such games as being run by “an umpire who supervises the course of play,” conveying to each player only such information as is required by the rules (Von Neumann and Morgenstern (1953, pp. 69–84)). Similarly, Myerson (1986) considered multi-stage games in which “all players communicate confidentially with the mediator, so that no player directly observes the reports or recommendations of the other players.”

The papers closest to ours are Dequiedt and Martimort (2015) and Li (2017). In Dequiedt and Martimort (2015), two agents simultaneously and privately report their types to the principal, who can misrepresent each agent's report to the other agent. If we restrict attention to revelation mechanisms, then our definition of credibility is equivalent to their requirement that the principal report truthfully. However, this restriction loses some generality, so our model instead permits the auctioneer to communicate sequentially with bidders by adopting extensive-form mechanisms. Li (2017) proposed a definition of bilateral commitment power, and also introduced the messaging game that we use here. The definition in Li (2017) is restricted to dominant-strategy mechanisms, whereas credibility allows for Bayes–Nash mechanisms. Also, Li (2017) did not model the incentives faced by the auctioneer, which is the entire subject of the present study.

Our paper is related to the literature on mechanisms with imperfect commitment, in which some parts of the outcome are chosen freely by the designer after observing the agents' reports (Baliga, Corchon, and Sjöström (1997), Bester and Strausz (2000, 2001)). Our paper also relates to the literature that studies multi-period auction design with limited commitment (Milgrom (1987), McAfee and Vincent (1997), Skreta (2006, 2015), Liu, Mierendorff, Shi, and Zhong (2019)). In this paradigm, the auctioneer chooses a mechanism in each period, but cannot commit today to the mechanisms that she will choose in the future. In particular, if the object remains unsold, then the auctioneer may attempt to sell the object again. Essentially, these papers have a post-auction game, and require that the auctioneer is sequentially rational. Our machinery instead permits the auctioneer to misrepresent bidders' preferences during the auction.

Some papers model auctions as bargaining games in which the auctioneer cannot commit to close a sale (McAdams and Schwarz (2007a), Vartiainen (2013)). These papers fix a particular stage game, in which players can solicit, make, or accept offers, and study equilibria of the repeated game. The auctioneer does not promise to obey any rules—she is constrained only by the structure of the repeated game. In our model, the auctioneer instead promises in advance to abide by certain rules, and can only deviate from those rules

in ways that have innocent explanations. Thus, if the auctioneer promises to run a first-price auction, then she must conclude the auction after collecting the bids. By contrast, [McAdams and Schwarz \(2007a\)](#) and [Vartiainen \(2013\)](#) permit the auctioneer to restart play in the next period, exploiting the new information that she has learned.

Several papers study auctioneer cheating in specific auction formats, such as shill-bidding in second-price auctions ([McAdams and Schwarz \(2007b\)](#), [Rothkopf and Harstad \(1995\)](#), [Porter and Shoham \(2005\)](#)) and in ascending auctions with common values ([Chakraborty and Kosmopoulou \(2004\)](#), [Lamy \(2009\)](#)). [Loertscher and Marx \(2017\)](#) allowed the auctioneer to choose when to stop the clocks in a two-sided clock auction. We contribute to this literature by providing a definition of auctioneer incentive-compatibility that is not tied to a particular format, and can thus be used as a design criterion.

Our paper contributes to the line of research that studies standard auction formats by relaxing various assumptions of the benchmark model ([Milgrom and Weber \(1982\)](#), [Maskin and Riley \(1984\)](#), [Bulow, Huang, and Klemperer \(1999\)](#), [Fang and Morris \(2006\)](#), [Hafalir and Krishna \(2008\)](#), [Bergemann, Brooks, and Morris \(2017, 2019\)](#)). While the usual approach is to compare the standard formats in terms of expected revenue, we instead characterize the standard formats with a few simple desiderata. Of course, the desiderata of [Figure 1](#) do not exhaust the considerations of real-world auctioneers; factors such as interdependent values, risk aversion, and informational robustness importantly affect the choice of format.

2. MODEL

2.1. Definitions

We now define the model. Proofs omitted from the main text are in [Appendix B](#). The environment consists of:

1. A finite set of agents, N .
2. A set of outcomes, X .
3. A type space, $\Theta_N = \times_{i \in N} \Theta_i$, endowed with σ -algebra \mathcal{F} .
4. A probability measure $D : \mathcal{F} \rightarrow [0, 1]$.
5. Agent utilities $u_i : X \times \Theta_i \rightarrow \mathbb{R}$.
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.¹¹

We represent the rules of the mechanism as an extensive game form with imperfect information. This specifies the information that will be provided to each agent, the choices each agent will make, and the outcomes that will result, *assuming that the auctioneer follows the rules*. Crucially, we are not yet modeling the ways that the auctioneer can deviate.

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 I](#). The full definition of extensive forms is familiar to most readers, so we relegate further detail to [Appendix A](#). We restrict attention to mechanisms with perfect

¹¹In 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.

TABLE I
NOTATION FOR EXTENSIVE GAME FORMS

Name	Notation	Representative Element
histories	H	h
precedence relation over histories	$<$	
reflexive precedence relation	\leq	
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	$\text{succ}(h)$	
actions available at I_i	$A(I_i)$	

recall and finite depth (i.e., there exists some $K \in \mathbb{N}$ such that no history has more than K predecessors).

An **interim strategy** is a function from information sets to available actions, $\sigma_i : \mathcal{I}_i \rightarrow A$, satisfying $\sigma_i(I_i) \in A(I_i)$. Let Σ_i denote the set of i 's interim strategies, and denote an interim strategy profile by $\sigma_N = (\sigma_i)_{i \in N}$. An **ex ante strategy** is a function from types to interim strategies, $S_i : \Theta_i \rightarrow \Sigma_i$. An ex ante strategy profile is $S_N = (S_i)_{i \in N}$, which implies an interim strategy profile for each type profile, $S_N(\theta_N) = (S_i(\theta_i))_{i \in N}$. We use $S_i(I_i, \theta_i)$ to denote the action played under ex ante strategy S_i at information set I_i by type θ_i .

By convention, many papers make statements about mechanisms that implicitly refer to a particular equilibrium of the mechanism, such as the claim “second-price auctions are efficient.” To reduce ambiguity, we will state our results explicitly for pairs (G, S_N) consisting of a mechanism and a strategy profile, which we refer to as a *protocol*.

Let $x^G(\sigma_N)$ denote the outcome in G , when agents play according to σ_N . Let $u_i^G(\sigma_i, \sigma_{-i}, \theta_i) \equiv u_i(x^G(\sigma_i, \sigma_{-i}), \theta_i)$.

Several definitions that follow take the expectation of a utility function.¹² We implicitly restrict attention to protocols and strategies such that the relevant utility function is measurable.

DEFINITION 1: (G, S_N) is **Bayes incentive-compatible (BIC)** if, for all $i \in N$, for all $\theta_i \in \Theta_i$,

$$S_i(\theta_i) \in \arg \max_{\sigma_i} \mathbb{E}_{\theta_{-i}} [u_i^G(\sigma_i, S_{-i}(\theta_{-i}), \theta_i)].$$

2.2. Pruning

At first glance, when constructing extensive-form mechanisms, it may seem important to keep track of off-path beliefs. However, if certain histories are off-path at every type profile, then we can delete those histories without altering the mechanism’s incentive properties. Similarly, if an agent is called to play, but reveals no outcome-relevant information about his type, we can skip that step without undermining incentives. Thus, we

¹²Definitions 1, 4, and 11.

restrict attention to the class of *pruned* protocols.¹³ 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 $z(\sigma_N)$ denote the terminal history that results from interim strategy profile σ_N . Formally, we have the following:

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 $|\text{succ}(h)| \geq 2$.
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) $x^G(S_N(\theta_i, \theta_{-i})) \neq x^G(S_N(\theta'_i, \theta_{-i}))$.

By the next proposition, when our concern is to construct a BIC protocol, 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) . If the type space Θ_N is finite and the probability measure D has full support, then every information set in a pruned protocol is reached with positive probability, which implies that any Bayes-Nash equilibrium survives equilibrium refinements that restrict off-path beliefs.

2.3. A Messaging Game

We now explicitly model the auctioneer¹⁴ as a player (denoted 0). The auctioneer has utility $u_0 : X \times \Theta_N \rightarrow \mathbb{R}$.

The auctioneer promises in advance to run some protocol (G, S_N) . We now describe a messaging game G^* that includes the auctioneer as a player. In G^* , the auctioneer contacts players privately and sequentially. At each step, she contacts some agent i , sending a message that corresponds to one of i 's information sets in the mechanism G . Agent i replies with one of the actions available at that information set. At any step, the auctioneer can choose an outcome x and end the game. Thus, the auctioneer can deviate from G by altering the sequence of players or information sets, or by choosing different outcomes.

Formally, the messaging game generated by protocol (G, S_N) is defined as follows: Let the auctioneer's message space be $M = \bigcup_i \mathcal{I}_i$.

1. The auctioneer chooses to:
 - (a) Either: Select outcome $x \in X$ and end the game.
 - (b) Or: Go to step 2.

¹³This is stronger than the definition of pruning in Li (2017), which includes only the first requirement.

¹⁴We 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.

2. The auctioneer chooses some agent $i \in N$ and sends a message $m = I_i \in \mathcal{I}_i$.
3. Agent i privately observes message I_i and chooses reply $r \in \mathcal{A}(I_i)$.
4. The auctioneer privately observes r .
5. Go to step 1.

There exists an auctioneer strategy in the messaging game that ‘follows the rules’ of the mechanism G . These rules prescribe which agents to contact, in what order, what messages to send, when to end the game, and what outcome to choose.

We use S_0^G to denote the rule-following auctioneer strategy. Formally, S_0^G is defined by the following algorithm: Initialize $\hat{h} := h_\emptyset$. At each step, if \hat{h} is a terminal history in G , end the game and choose outcome $g(\hat{h})$. Else, contact agent $P(\hat{h})$ and send message $m = I_{P(\hat{h})}$ such that $\hat{h} \in I_{P(\hat{h})}$. Upon receiving reply r , update $\hat{h} := h' \in \text{succ}(\hat{h}) \mid \mathcal{A}(h') = r$, and iterate.¹⁵

We now make a substantive restriction: The auctioneer can only deviate in ways that no agent can detect. Formally, in the messaging game, agent i observes the sequence of communication between himself and the auctioneer $(m_i^k, r_i^k)_{k=1}^{T_i}$, and directly observes some details of the outcome, as specified by the partition Ω_i . An observation for i is a tuple $((m_i^k, r_i^k)_{k=1}^{T_i}, \omega_i)$, where ω_i is the cell of Ω_i that contains the outcome.¹⁶ Let $o_i(S_0, S_N, \theta_N)$ be i ’s observation when the auctioneer plays S_0 , the agents play S_N , and the type profile is θ_N .

DEFINITION 3: Given some promised strategy profile (S_0, S_N) , auctioneer strategy \hat{S}_0 is **safe** if, for all agents $i \in N$ and all type profiles θ_N , there exists $\hat{\theta}_{-i}$ such that $o_i(\hat{S}_0, S_N, \theta_N) = o_i(S_0, S_N, (\theta_i, \hat{\theta}_{-i}))$. $S_0^*(S_0, S_N)$ denotes the set of safe strategies.

G^* is the messaging game restricted to $S_0^*(S_0^G, S_N)$; this constrains the auctioneer to only play safe deviations from the rule-following strategy.¹⁷

Definition 3 permits the auctioneer to deviate only if every agent’s observation has an **innocent explanation**; there must exist $\hat{\theta}_{-i}$ such that i ’s observation is consistent with the auctioneer playing S_0^G , the agents playing S_N , and the other agents’ types being $\hat{\theta}_{-i}$.

DEFINITION 4: (G, S_N) is **credible** if

$$S_0^G \in \arg \max_{S_0 \in S_0^*(S_0^G, S_N)} \mathbb{E}_{\theta_N} [u_0(S_0, S_N, \theta_N)],$$

where $u_0(S_0, S_N, \theta_N)$ is the utility to the auctioneer from the outcome that results from (S_0, S_N) when the type profile is θ_N .

This parallels the definition of agent incentive compatibility in **Hurwicz (1972)**:

¹⁵We have not defined S_0^G at information sets in the messaging game that are ruled out by S_0^G . Since we are not considering trembles by the auctioneer, all such strategies are outcome-equivalent, and this omission is harmless.

¹⁶Note 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.

¹⁷In describing the messaging game G^* , we have not specified payoffs for infinite sequences of communication. However, since (G, S_N) has finite depth, any safe deviation has finite length, and the resulting payoffs are well-defined.

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.

In our definition, the auctioneer is allowed to behave in ways that can be made to look “legal” by misrepresenting the preferences of the other agents, with the proviso that the fictitious preferences should be within certain “plausible” limits. These limits are defined by the type space.

Instead of just choosing different outcomes, Definition 4 permits the auctioneer to modify G by altering the sequence of information sets. This may materially expand the auctioneer’s strategic opportunities, as 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. Agent 1 is assumed to observe 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 follows the rules. 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, following the rules 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 4 takes the expectation of θ_N with respect to the ex ante distribution D . However, when Θ_N is finite and D has full support, Definition 4 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




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

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 rules otherwise. Thus, if S_0 is an ex ante best response, then its corresponding continuation strategies are also best responses along the equilibrium path-of-play.

When our concern is to construct a credible protocol, it is also without loss of generality to consider only pruned protocols.

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 G^* .*

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 $S'_0 \in S_0^*(S_0, S_N)$, then $S_0^*(S'_0, S_N) \subseteq S_0^*(S_0, S_N)$. Thus, if (S'_0, S_N) is a Bayes–Nash equilibrium of the messaging game restricted to $S_0^*(S_0, S_N)$, then it is also a Bayes–Nash equilibrium of the messaging game restricted to $S_0^*(S'_0, S_N)$.*

Definition 4 is stated for pure strategies, but can be generalized to allow the auctioneer to mix. To do so, we simply extend the definition of extensive game forms so that G includes chance moves. We then specify that \hat{S}_0 is **safe** if, for all agents $i \in N$ and all type profiles θ_N , for any observation of agent i that occurs for some realization of the auctioneer’s randomization under $(\hat{S}_0, S_N, \theta_N)$, there exists $\hat{\theta}_{-i}$ so that the same observation occurs for some realization of the auctioneer’s randomization under $(S_0^G, S_N, \theta_i, \hat{\theta}_{-i})$.

In some settings, auctioneer randomization is needed to deliver the right incentives for the agents. However, randomization does not improve auctioneer incentives: We cannot construct a credible protocol (G, S_N) by randomizing over deterministic non-credible protocols. Given randomized (G, S_N) , let (G', S_N) be a deterministic protocol in which we fix a particular realization of the auctioneer’s randomization. Suppose (G, S_N) is credible, so the auctioneer is indifferent between S_0^G and $S_0^{G'}$. Switching from G to G' shrinks the set of innocent explanations, and therefore the set of safe deviations. The auctioneer preferred S_0^G to any safe deviation in the larger set, and therefore prefers $S_0^{G'}$ to any safe deviation in the smaller set, so (G', S_N) is credible.

In the settings we are about to consider, randomization is not helpful for agent incentives.¹⁸ Thus, we will restrict attention to deterministic protocols.

3. CREDIBLE OPTIMAL AUCTIONS

We now study credible auctions in the independent private values (IPV) model (Myerson (1981)). We make this choice for two reasons: First, this is a benchmark model in auction theory, so using it shows that the results are driven by credibility, and not by

¹⁸In auctions with independent private values, there always exists a deterministic mechanism that maximizes expected revenue. For instance, we can run a second-price auction that scores bids according to their ironed virtual value, breaking ties deterministically (Myerson (1981)).

some hidden feature of an unusual model.¹⁹ Second, in the symmetric IPV model, revenue equivalence implies that the standard auctions start on an equal footing—the value distribution does not tip the scales in favor of a particular format, unlike the model with affiliated signals (Milgrom and Weber (1982)) or the model with risk aversion (Maskin and Riley (1984)).

Assume there are at least two agents, henceforth referred to as **bidders**. An outcome $x = (y, t_N)$ consists of a winner $y \in N \cup \{0\}$ and a profile of payments (one for each bidder) $t_N \in \mathbb{R}^{|N|}$, so $X = (N \cup \{0\}) \times \mathbb{R}^{|N|}$.

Bidders have private values, that is,

$$u_i((y, t_N), \theta_i) = 1_{i=y} v_i(\theta_i) - t_i,$$

where $v_i : \Theta_i \rightarrow \mathbb{R}$. We will abuse notation slightly, and use θ_i to refer both to i 's type, and to the real number associated with that type.

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

1. Either: $y \neq i, y' \neq i$, and $t_i = t'_i$,
2. Or: $y = y' = i$ and $t_i = t'_i$.

The auctioneer desires revenue, and her value for the object is normalized to zero:²⁰

$$u_0((y, t_N), \theta_N) = \sum_{i \in N} t_i.$$

An allocation rule is a function $\tilde{y} : \Theta_N \rightarrow N \cup \{0\}$, and a transfer rule is a function $\tilde{t}_N : \Theta_N \rightarrow \mathbb{R}^{|N|}$. For any protocol (G, S_N) , we can consider its induced allocation rule and transfer rule $(\tilde{y}^{G, S_N}(\cdot), \tilde{t}_N^{G, S_N}(\cdot))$. Where it is clear, we suppress the dependence on (G, S_N) to ease notation.

Let $\pi(G, S_N) = \mathbf{E}_{\theta_N}[\sum_{i \in N} \tilde{t}_i^{G, S_N}(\theta_N)]$ denote the expected revenue of (G, S_N) . We will specify the relevant distribution shortly.

DEFINITION 5: (G, S_N) is **optimal** if it maximizes $\pi(G, S_N)$ subject to the constraints:

1. Incentive compatibility: (G, S_N) is BIC.
2. Voluntary participation: For all i , there exists σ'_i that ensures that i does not win and has a zero net transfer, regardless of σ'_{-i} .²¹

3.1. Credible Static Optimal Auctions

We now characterize credible static optimal auctions. Assume that $\Theta_i = [0, 1]$ and that θ_i is independently distributed according to continuous full-support density $f_i : [0, 1] \rightarrow \mathbb{R}$.

We restrict attention to protocols such that:

¹⁹As Brooks and Du (2018) observed, “The IPV model has been broadly accepted as a useful benchmark when values are private, but there is no comparably canonical model when values are common.”

²⁰The results that follow would require only small modifications if the auctioneer’s payoff was a weighted average of revenue and social welfare.

²¹There are several standard ways of defining participation constraints, not entirely equivalent for our purposes. This definition appears in Maskin and Riley (1984). The existence of this non-participating strategy is used in the proof of Proposition 6, to establish that strategy-proof auctions are winner-paying. If we merely required that the bidder’s interim expected payoff was non-negative, then the conclusion would not follow.

1. For all θ_N , $\tilde{y}(\cdot)$ and $\tilde{t}_i(\cdot)$ are measurable functions (with respect to the Borel σ -algebra on Θ_N).²²
2. For all θ_i , $\tilde{y}(\theta_i, \cdot)$ and $\tilde{t}_i(\theta_i, \cdot)$ are measurable functions (with respect to the Borel σ -algebra on Θ_{-i}).
3. For all θ_{-i} , $\tilde{y}(\cdot, \theta_{-i})$ and $\tilde{t}_i(\cdot, \theta_{-i})$ are measurable functions (with respect to the Borel σ -algebra on Θ_i).

These conditions ensure that expected transfers and allocations are well-defined, both ex ante and interim. These are implicit in almost all papers with continuum type spaces and transferable utility. We make these restrictions explicit because the proof of Theorem 1 runs into some measure-theoretic subtleties.²³

DEFINITION 6: (G, S_N) is **static** if, for each bidder i , i has exactly one information set, and for every terminal history z , there exists $h < z$ such that $P(h) = i$.

Next, we prove that, in a credible static auction, the winner makes a payment that essentially depends only on his own type. Thus, we can regard each bidder as placing bids, with the assurance that if he wins the object, he pays exactly his bid.

THEOREM 1—Pay-as-Bid: *If (G, S_N) is credible and static, then, for each bidder i , there exists a function $\tilde{b}_i : \Theta_i \rightarrow \mathbb{R}$ such that almost everywhere in Θ_N , if $\tilde{y}(\theta_i, \theta_{-i}) = i$, then $\tilde{t}_i(\theta_i, \theta_{-i}) = \tilde{b}_i(\theta_i)$.*

PROOF: If the pay-as-bid property does not hold, then we can construct a safe deviation that raises payments on a positive-measure set. However, we cannot simply charge the ‘highest safe payment’ point-by-point, because there may be uncountably many opponent type profiles consistent with i winning the object, and the pointwise supremum of an uncountable family of measurable functions may not be measurable.

LEMMA 1—Hajlasz and Malý (2002)²⁴: *Let Φ be a family of measurable functions defined on a set $E \subseteq \mathbb{R}^n$. There exists a countable subfamily $\hat{\Phi} \subseteq \Phi$ such that, for all $\phi \in \Phi$, $\sup \hat{\Phi} \geq \phi$ almost everywhere.*

Consequently, let $(\theta_{-i}^k)_{k=1}^\infty$ be a countable subset of opponent type profiles, such that for all θ_{-i} , $\sup_k \tilde{t}_i(\cdot, \theta_{-i}^k) \geq \tilde{t}_i(\cdot, \theta_{-i})$ almost everywhere in Θ_i . $\sup_k \tilde{t}_i(\cdot, \theta_{-i}^k)$ is measurable.

We assert that $\tilde{b}_i(\cdot) = \sup_k \tilde{t}_i(\cdot, \theta_{-i}^k)$. Suppose the set

$$\left\{ \theta_N \mid \tilde{y}(\theta_N) = i \text{ and } \tilde{t}_i(\theta_i, \theta_{-i}) \neq \sup_k \tilde{t}_i(\theta_i, \theta_{-i}^k) \right\}$$

has positive measure. Then the set

$$Q = \left\{ \theta_N \mid \tilde{y}(\theta_N) = i \text{ and } \tilde{t}_i(\theta_i, \theta_{-i}) < \sup_k \tilde{t}_i(\theta_i, \theta_{-i}^k) \right\}$$

²²That is, for any $J \subseteq N \cup \{0\}$, its preimage $\{\theta_N \mid \tilde{y}(\theta_N) \in J\}$ is a Borel set, and for any Borel set $J \subseteq \mathbb{R}$, its preimage $\{\theta_N \mid \tilde{t}_i(\theta_N) \in J\}$ is a Borel set.

²³The restrictions rule out, for instance, that we can fix a Vitali set V , and specify that bidder 1 has transfer 1 if $\theta_1 \in V$ and $\theta_2 = 0.5$, and transfer 0 otherwise, in which case 1’s expected transfer conditional on $\theta_2 = 0.5$ is not defined since V is not measurable.

²⁴Lemma 2.6 in Hajlasz and Malý (2002), which is a special case of Lemma 2.6.1 in Meyer-Nieberg (1991).

has positive measure. Since transfers and allocations can only change when the action profile changes, Q is measurable with respect to the equilibrium action profiles.

We now construct a safe deviation: Fix some finite K . If the bidder's chosen actions are consistent with any type profile $(\theta_i, \theta_{-i}) \in Q$, then instead charge bidder i $\max_{k \leq K} \tilde{t}_i(\theta_i, \theta_{-i}^k)$. Let k^* denote the arg max. If $\tilde{y}(\theta_i, \theta_{-i}^{k^*}) = i$, then allocate the object to i ; else keep the object. Otherwise, play according to S_0^G . This deviation takes the maximum of finitely many measurable functions, so the resulting transfer $\tilde{t}_i^K : \Theta_N \rightarrow \mathbb{R}$ is measurable.

For K large enough, this deviation is profitable. In particular, for any $\theta_N \in Q$, $\tilde{t}_i^K(\theta_N)$ is non-decreasing in K . Thus, by the monotone convergence theorem,

$$\begin{aligned} \lim_{K \rightarrow \infty} \mathbb{E}_{\theta_N} [\tilde{t}_i^K(\theta_N) \mid \theta_N \in Q] &= \mathbb{E}_{\theta_N} \left[\lim_{K \rightarrow \infty} \tilde{t}_i^K(\theta_N) \mid \theta_N \in Q \right] \\ &= \mathbb{E}_{\theta_N} \left[\sup_k \tilde{t}_i(\theta_i, \theta_{-i}^k) \mid \theta_N \in Q \right] > \mathbb{E}_{\theta_N} [\tilde{t}_i(\theta_N) \mid \theta_N \in Q], \end{aligned}$$

which completes the proof. Q.E.D.

DEFINITION 7: (G, S_N) is a **first-price auction** if (G, S_N) is static, and each bidder i either chooses a bid b_i from a set $B_i \subset \mathbb{R}_0^+$ or declines to bid, such that:

1. Each bidder i pays b_i if he wins and 0 if he loses.
2. If any bidder places a bid, then some maximal bidder wins the object. Otherwise, no bidder wins.

If Clauses 1 and 2 hold almost everywhere in Θ_N , then (G, S_N) is a **first-price auction almost everywhere**.

We represent a reserve price by restricting the set B_i .

For the next theorem, we assume that the distributions are symmetric, that is, $f_i(\cdot) = f_j(\cdot)$ for all i, j , and regular, that is, $\theta_i - \frac{1 - F_i(\theta_i)}{f_i(\theta_i)}$ is strictly increasing. We also restrict attention to winner-paying protocols.

DEFINITION 8: (G, S_N) is **winner-paying** if, for all θ_N , if $\tilde{t}_i(\theta_N) \neq 0$, then $\tilde{y}(\theta_N) = i$.

THEOREM 2: Assume the distributions are symmetric and regular. Assume (G, S_N) is winner-paying and optimal. If (G, S_N) is a first-price auction, then (G, S_N) is credible and static. If (G, S_N) is credible and static, then (G, S_N) is a first-price auction almost everywhere.

PROOF: Suppose (G, S_N) is a first-price auction. (G, S_N) is static by definition. Every safe deviation that sells the object involves charging some bidder his bid, so no safe deviation yields more revenue than following the rules. Thus, (G, S_N) is credible.

Suppose (G, S_N) is credible and static. By Theorem 1, there exists a function $\tilde{b}_i : \Theta_i \rightarrow \mathbb{R}$ such that, almost everywhere in Θ_N , if type θ_i wins, then i pays $\tilde{b}_i(\theta_i)$. (G, S_N) is optimal, so the participation constraint of the lowest type binds, and we can pick a non-negative function $\tilde{b}_i : \Theta_i \rightarrow \mathbb{R}_0^+$. We now partition i 's actions into bidding actions $B_i = \{\tilde{b}_i(\theta_i) \mid \theta_i \in \Theta_i \text{ and } \exists \theta_{-i} : \tilde{y}(\theta_i, \theta_{-i}) = i\}$, and actions that decline. (G, S_N) is winner-paying, so Clause 1 of Definition 7 holds almost everywhere.

(G, S_N) is optimal, which determines the allocation rule and i 's interim expected transfer almost everywhere (Myerson (1981)). By BIC, i 's interim expected transfer is increasing in i 's type, and the distributions are symmetric and regular, so almost everywhere the

winner has a maximal type, and thus a maximal bid. Thus, Clause 2 of Definition 7 holds almost everywhere. Thus (G, S_N) is a first-price auction almost everywhere. *Q.E.D.*

We now relax the assumption that the distributions are symmetric and regular, and that the protocol is winner-paying and optimal. In particular, rather than requiring that the protocol be optimal, we will require that, with probability 1, no bidder knows at the interim stage that he will win for sure.

DEFINITION 9: (G, S_N) is **contestable** if, almost everywhere in Θ_N , if $\tilde{y}(\theta_i, \theta_{-i}) = i$, then there exists θ'_{-i} such that $\tilde{y}(\theta_i, \theta'_{-i}) \neq i$.

Since $\Theta_i = \Theta_j = [0, 1]$, optimal auctions are contestable.

The first-price auctions of Theorem 2 generalize to a larger class that permits the auctioneer to extract transfers from losing bidders, though each losing bidder's transfer must depend only on his own bid.

Whether this class is of more than technical interest will vary from case to case. Most economically important auctions, such as those for art, for mineral rights, for spectrum, or for online advertising, do not extract payments from losing bidders. Some real-world auctions may need to respect ex post individual rationality, since otherwise one party will try to annul the contract afterwards. The resulting transaction costs may constrain the auctioneer to use winner-paying protocols.

We now state the definition that generalizes first-price auctions.

DEFINITION 10: (G, S_N) is a **twin-bid auction** if (G, S_N) is static, and each bidder chooses a two-dimensional bid (b_i^W, b_i^L) from a set $B_i \subset \mathbb{R}^2$ such that:

1. Each bidder i pays b_i^W if he wins and b_i^L if he loses.
2. If any bidder places a bid such that $b_i^W - b_i^L > 0$, then some bidder wins the object.
3. If i wins the object, then $b_i^W - b_i^L \geq \max\{0, \max_{j \neq i} b_j^W - b_j^L\}$.

If Clauses 1, 2, and 3 hold almost everywhere in Θ_N , then (G, S_N) is a **twin-bid auction almost everywhere**.

Twin-bid auctions include first-price auctions and all-pay auctions, though the credibility of all-pay auctions is sensitive to the assumption that the object is costless to provide. (More generally, $b_i^W - b_i^L$ must be no less than the auctioneer's cost of provision, which rules out standard all-pay auctions.²⁵) Twin-bid auctions also encompass first-price auctions with entry fees (b_i^L is the entry fee), and first-price auctions in which losing bidders are paid fixed compensation ($b_i^L < 0$). Bidders who place higher bids may also receive more compensation if they lose; under the assumptions of Maskin and Riley (1984), this is the form of the optimal auction for symmetric bidders with constant absolute risk aversion.²⁶

THEOREM 3: Assume (G, S_N) is contestable. If (G, S_N) is a twin-bid auction, then (G, S_N) is credible and static. If (G, S_N) is credible and static, then (G, S_N) is a twin-bid auction almost everywhere.

²⁵The case when $b_i^W - b_i^L$ is exactly equal to the cost of provision was studied in Dequiedt and Martimort (2007), an early draft of Dequiedt and Martimort (2015).

²⁶Theorem 14 (Maskin and Riley (1984, pp. 1506–1507)). This claim follows from their Equations (75) and (77), since μ is non-decreasing.

The proof of Theorem 3 does not rely on independence, so the characterization holds even with correlated types. Twin-bid auctions are not strategy-proof, except in degenerate cases.

DEFINITION 11: (G, S_N) is **strategy-proof** if, for all $i \in N$, for all S'_{-i} , for all $\theta_i \in \Theta_i$.

$$S_i(\theta_i) \in \arg \max_{\sigma_i} \mathbb{E}_{\theta_{-i}} [u_i^G(\sigma_i, S'_{-i}(\theta_{-i}), \theta_i)].$$

The definition above requires that S_i is a best response to all S'_{-i} , taking the expectation with respect to θ_{-i} . It is natural to consider a stronger definition that requires S_i to be a best response to all S'_{-i} and all θ'_{-i} . Under private values, these definitions are equivalent.

PROPOSITION 3: *Let (G, S_N) be such that there exist $\theta_i < \theta'_i < \theta''_i < \theta'''_i$, θ_{-i} , and θ'_{-i} such that $\tilde{y}(\theta_i, \theta_{-i}) \neq i = \tilde{y}(\theta'_i, \theta_{-i})$ and $\tilde{y}(\theta''_i, \theta'_{-i}) \neq i = \tilde{y}(\theta'''_i, \theta'_{-i})$. If (G, S_N) is a twin-bid auction, then (G, S_N) is not strategy-proof.*

What happens to Theorem 3 if we remove the assumption that the protocol is contestable? In that case, then some bidder i could have actions that win the object for sure, even when the difference $b_i^W - b_i^L$ is not high enough to satisfy Clause 3 of Definition 10. Since there is only one object for sale, at most one bidder can have incontestable actions. The characterization of credible static mechanisms is otherwise unchanged. We omit the proof, since it is an easy modification of the proof of Theorem 3.

3.2. Credible and Strategy-Proof Optimal Auctions

We now characterize credible strategy-proof optimal auctions. In particular, we will show that certain ascending auctions are credible and strategy-proof.

We must make a modeling choice, because ascending auctions with discrete steps are not optimal for continuum type spaces. We could proceed by introducing a model for continuous-time auctions, as in Milgrom and Weber (1982). However, we wish to argue that credibility and strategy-proofness select ascending auctions out of a general class, and there is not yet any theory of continuous-time games that rivals the generality of extensive-form games.²⁷

Consequently, our approach is to discretize the type space, so that clock auctions (and many other dynamic protocols) can be optimal. Let $\Theta_i = \{\theta_i^1, \dots, \theta_i^k\}$. Assume $\theta_i^1 = 0$ and that $\theta_i^{k+1} - \theta_i^k > 0$. (We are continuing to abuse notation, and use θ_i^k to refer both to i 's k th type, and to $v_i(\theta_i^k)$, the real number associated with that type.)

Types are independently distributed, with probability mass function $f_i : \Theta_i \rightarrow (0, 1]$ and corresponding $F_i(\theta_i^k) = \sum_{l=1}^k f_i(\theta_i^l)$. F_N is **symmetric** if, for all i, j , $\Theta_i = \Theta_j$ and $F_i = F_j$.

The virtual values machinery in Myerson (1981) applies mutatis mutandis to the discrete setting.

²⁷For an explanation of some difficulties involved in continuous-time game theory, see Simon and Stinchcombe (1989).

DEFINITION 12: For each k , we define the **virtual value** of θ_i^k to be²⁸

$$\eta_i(\theta_i^k) \equiv \theta_i^k - \frac{1 - F_i(\theta_i^k)}{f_i(\theta_i^k)} (\theta_i^{k+1} - \theta_i^k). \tag{1}$$

$F_N = (F_i)_{i \in N}$ is **regular** if, for all i , $\eta_i(\cdot)$ is strictly increasing.

Optimal auctions have a characterization in terms of virtual values when certain constraints bind. $\check{u}_i^{G, S_N}(k, k')$ denotes the expected utility of bidder i when his type is θ_i^k and he plays as though his type is $\theta_i^{k'}$.

PROPOSITION 4—**Elkind (2007)**: Assume F_N is regular and (G, S_N) satisfies the constraints in Definition 5. (G, S_N) is optimal if and only if:

1. Participation constraints bind for the lowest types. $\forall i : \check{u}_i^{G, S_N}(1, 1) = 0$.
2. Incentive constraints bind locally downward. $\forall i : \forall k \geq 2 : \check{u}_i^{G, S_N}(k, k) = \check{u}_i^{G, S_N}(k, k - 1)$.
3. The allocation maximizes virtual value. $\forall \theta_N$:
 - (a) If $\max_i \eta_i(\theta_i) > 0$, then $\tilde{y}(\theta_N) \in \arg \max_i \eta_i(\theta_i)$.
 - (b) If $\eta_i(\theta_i) < 0$, then $i \neq \tilde{y}(\theta_N)$.

Ties occur with positive probability under discrete type spaces, although the probability goes to 0 as we make the discretization finer. For convenience, we will assume that the protocol breaks ties deterministically according to a fixed priority order.

DEFINITION 13: Let ρ denote a reserve price. (G, S_N) is **orderly** if there exists a strict total order \triangleright on $\rho \cup \bigcup_{i \in N} \Theta_i$, such that

1. i wins the object if and only if $\theta_i \triangleright \rho$ and $\theta_i \triangleright \theta_j$ for all $j \neq i$.
2. $\theta_i \triangleright \theta_j$ if $\theta_i > \theta_j$.
3. $\theta_i \triangleright \rho$ if and only if $\theta_i \geq \rho$.
4. If $\theta_i = \theta_j$, $\theta'_i = \theta'_j$, and $\theta_i \triangleright \theta_j$, then $\theta'_i \triangleright \theta'_j$.

We use \min^\triangleright to denote the minimum of a set with respect to \triangleright , and \max^\triangleright similarly.

For two bidders, an ascending auction can proceed as follows: Choose an optimal reserve ρ^* . Ask each bidder to either bid ρ^* or quit. If both quit, keep the object. Otherwise, alternate between the bidders, asking each bidder to either raise his bid by one increment or quit. When there is only one bidder left, allocate the object to that bidder and charge him his current bid.

The following definition generalizes the two-bidder auction, and will shortly be used to characterize credible strategy-proof optimal auctions.

DEFINITION 14: (G, S_N) is an **ascending auction** if:

1. The induced allocation rule $\tilde{y}(\cdot)$ is orderly.
2. The induced payment rule satisfies **threshold pricing**, that is,

$$\tilde{z}_i(\theta_N) = \begin{cases} \min\{\theta'_i \in \Theta_i \mid \tilde{y}(\theta'_i, \theta_{-i}) = i\} & \text{if } \tilde{y}(\theta_N) = i, \\ 0 & \text{otherwise.} \end{cases}$$

²⁸Since $1 - F_i(\theta_i)$ is equal to 0 at the upper bound, we can define θ_i^{k+1} arbitrarily for the purposes of Equation (1).

3. All bidders start with a profile of initial bids $b_N := (\theta_i^1)_{i \in N}$.
4. At each non-terminal history h , some bidder i is called to play, and is offered some price $p \in \Theta_i$. Each available action either accepts the price or quits.
 - (a) The offered price p is no less than i 's current bid.
 - (b) If i accepts the price, then his bid is updated $b_i := p$.
 - (c) If i quits, then he is not called to play again, does not win the object, and pays 0.
 - (d) When i is called to play at h , he knows the offered price, and knows, for each action, whether that action accepts or quits.
5. At each information set I_i :
 - (a) Either: There is a unique action that accepts.
 - (b) Or: If i plays any action that accepts, then he is not called to play again, wins the object, and pays the offered price p .
6. If i wins the object, then he pays his current bid.

By definition, ascending auctions have threshold pricing and bid-or-quit decisions. This implicitly limits the price offers that each bidder receives. In particular, no bidder is ever offered a price higher than is needed to win, given the profile of current bids and the tie-breaking order.

In stating Definition 14, we have deliberately omitted what each bidder is told about the other bidders. The protocol could require that each bidder is informed about the number of active bidders or the identities of the active bidders. The protocol could specify that each bidder places an increasing sequence of bids, receiving no other information until he quits or is the last bidder left. These all count as ascending auctions for the purposes of the definition.

OBSERVATION 3: *If F_N is regular and symmetric, then there exists an optimal ascending auction. In any ascending auction, participation constraints bind for the lowest types and incentive constraints bind locally downward. Given an optimal reserve $\rho^* = \min_k \theta_i^k \mid \eta_i(\theta_i^k) > 0$, the ascending auction maximizes the virtual value of the winning bidder. By Proposition 4, such an auction is optimal.*

Credibility and strategy-proofness pin down most of the game tree, except for certain corner cases. In particular, consider optimal auctions for a single bidder; the bidder will win and pay the reserve price if and only if his type is above the reserve. Thus, many extensive-form mechanisms are credible. For instance, the auctioneer could first ask the bidder to report whether his type is even or odd, and then to report whether his type is above the reserve. There exist many protocols that are essentially ‘long-winded’ ways to make a take-it-or-leave-it offer. These situations sometimes occur in multi-bidder auctions, for instance, if bidder 1 quits before the reserve is met, and only bidder 2 is left.

For these corner cases, ruling out long-winded protocols yields a simpler and more transparent characterization. Rather than describing all the drawn-out ways that the auctioneer could make a take-it-or-leave-it offer, we will instead assume that, once the auctioneer is in a position to make such an offer, she does so directly and in a single step. In particular, we will restrict attention to the class of *concise* protocols, defined as follows.

DEFINITION 15: Under (G, S_N) , i **faces a posted price** at h if $P(h) = i$ and there exists a price τ_h such that, at all terminal histories that follow h , if i wins, then i pays τ_h . (G, S_N) is **concise** if, for any history h , if i faces a posted price at h , then:

1. The information set containing h is singleton.
2. For all $h' > h$, $P(h') \neq i$.

We go on to characterize credible strategy-proof optimal auctions. Notably, optimality and strategy-proofness together imply that the protocol is winner-paying. Thus, we do not need to make that assumption separately in the results that follow.

The definition of extensive-form mechanisms permits the auctioneer to communicate with bidders in any order, to convey information to the bidder called to play, and to ask that bidder 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, we have the following:

THEOREM 4: *Assume that F_N is regular and symmetric and that (G, S_N) is optimal. If (G, S_N) is an ascending auction, then (G, S_N) is credible and strategy-proof.*

Assume additionally that (G, S_N) is concise and orderly. If (G, S_N) is credible and strategy-proof, then (G, S_N) is an ascending auction.

PROOF OVERVIEW: 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 bidder 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 bidders' incentives—we can define a new protocol (G', S'_N) that is BIC and has voluntary participation, but yields strictly more expected revenue than (G, S_N) . But (G, S_N) is optimal, a contradiction.

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. Fix a protocol and a history h where bidder i is called to play. Consider the types θ_i consistent with h , such that there exists θ_{-i} consistent with h , such that i wins at (θ_i, θ_{-i}) . A key feature of ascending auctions is that, at each history, either these types pool on the same action, or i faces a posted price. This 'winner-pooling' property is stated precisely in Proposition 11, and is closely related to unconditional winner privacy as defined by Milgrom and Segal (2020). If these types do not pool and i does not face a posted price, then the auctioneer can sometimes deviate to charge i a higher price. In 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 bidders' actions midway through the extensive form. We construct an algorithm that produces a profitable safe deviation for any such extensive form. This establishes that auctions that satisfy the desiderata are winner-pooling, which pins down the extensive form of the game. (The proof is in the Appendix.) *Q.E.D.*

By Theorem 2, restricting attention to revelation mechanisms forces a sharp choice between incentives for the auctioneer and strategy-proofness for the bidders. Theorem 4 shows that allowing other extensive forms relaxes this trade-off.

The characterization in Theorem 4 assumes optimality. This is not just a feature of our proof technique: the ascending auction is credible because it is optimal. If the reserve price is below-optimal, then the auctioneer could profitably deviate by chandelier bidding up to the optimal reserve. If the type distributions are asymmetric, then the auctioneer may profitably deviate by enforcing bidder-specific reserve prices.²⁹ We characterize the asymmetric case in Theorem 6.

²⁹Symmetric beliefs may seem like a knife-edge case. However, in some real-world auctions, strong bidders can mask their identities and bid through proxies so as to avoid discriminatory pricing. When faced with anonymous bidders, it is quite reasonable for auctioneers to hold symmetric beliefs.

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

THEOREM 5—Extensive Pay-as-Bid: *Assume (G, S_N) is credible. Suppose that, with positive probability, i is called to play at information set I_i , takes some action a , and wins without being called to play again. Conditional on that event, there is a price $t_i(I_i, a)$ that i will pay with probability 1.*

PROOF: Suppose that the event obtains, and there are two distinct prices $t_i < t'_i$, such that i pays each with positive conditional probability. The auctioneer has a profitable safe deviation: when i is meant to pay t_i , she can deviate to charge t'_i , so (G, S_N) is not credible. *Q.E.D.*

Theorem 5 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. The converse of Theorem 5 is not true. For a counterexample, consider a ‘pay-as-bid’ static auction that allocates the object to the bidder who placed the second-highest bid.

Theorem 4 assumed that the distribution was symmetric; we now state a version that allows asymmetry. To proceed, we define a technical condition on the distribution. Clauses 1 and 2 of the following definition require that the distribution is generic, which removes distractions from tie-breaking. Clause 3 states that for any $\eta_i(\theta'_i)$ in the interior of the convex hull of $\eta_j(\Theta_j)$, we can find θ_j with virtual value ‘just below’ $\eta_i(\theta'_i)$. This is implied by continuum type spaces and continuous densities, but must be assumed separately for finite type spaces.

DEFINITION 16: F_N is **interleaved** if, $\forall i \neq j$:

1. $\forall \theta_i, \theta_j : \eta_i(\theta_i) \neq \eta_j(\theta_j)$,
2. $\forall \theta_i : \eta_i(\theta_i) \neq 0$,
3. $\forall \theta_i, \theta'_i$: if $\eta_i(\theta_i) < \eta_i(\theta'_i)$ and $\eta_j(\theta_j^1) < \eta_i(\theta'_i) < \eta_j(\theta_j^{K_j})$, then $\exists \theta_j : \eta_i(\theta_i) < \eta_j(\theta_j) < \eta_i(\theta'_i)$.

Under asymmetry, we can construct an optimal auction by modifying the ascending auction to score bids according to their corresponding virtual values, and to sell only when the high bidder’s virtual value is positive.

DEFINITION 17: (G, S_N) is a **virtual ascending auction** if

1. The induced allocation rule is

$$\tilde{y}(\theta_N) \in \begin{cases} \arg \max_i \eta_i(\theta_i) & \text{if } \max_i \eta_i(\theta_i) > 0, \\ \{0\} & \text{otherwise.} \end{cases}$$

2. The protocol satisfies Clauses 2 to 6 of Definition 14.

THEOREM 6: *Assume that F_N is regular and that (G, S_N) is optimal. If (G, S_N) is a virtual ascending auction, then (G, S_N) is credible and strategy-proof.*

Assume additionally that F_N is interleaved and that (G, S_N) is concise. If (G, S_N) is credible and strategy-proof, then (G, S_N) is a virtual ascending auction.

Virtual ascending auctions score bids asymmetrically: Bidder i may be asked to bid \$100 in order to beat j 's bid of \$50, and then to bid \$101 to beat j 's bid of \$51. Since the auctioneer is communicating privately, she could safely deviate to equalize the prices that bidders face (provided Θ_i and Θ_j overlap enough). Nonetheless, it is incentive-compatible for the auctioneer to follow the rules. For each bidder, truthful bidding is a best response to any safe deviation. Thus, if the auctioneer has a profitable safe deviation, then she could openly promise to deviate without undermining bidders' incentives. In that case, the original protocol was not optimal, a contradiction. It may seem intuitive that the auctioneer cannot credibly reject higher bids in favor of lower bids, but multi-round communication permits her to do so.

The virtual ascending auction can be modified to deal with irregular distributions: we simply alter Definition 17 to use ironed virtual values instead of virtual values, following the construction in Elkind (2007). In effect, if we iron virtual values in the interval θ_i^k to $\theta_i^{k'}$, the auctioneer promises ahead of time to jump i 's price directly from θ_i^k to $\theta_i^{k'+1}$. The proof that this is credible is the same as in the regular case.

Finally, the virtual ascending auction can be used to construct a static credible optimal auction. Consider a modified all-pay auction; each type θ_i makes a bid equal to the expected payment of θ_i in the virtual ascending auction, to be paid regardless of whether he wins. The winner is the bidder with the highest virtual value. This twin-bid auction is BIC and optimal, but neither strategy-proof nor ex post individually rational.³⁰

3.3. A Note on the Dutch Auction

The Dutch (descending) auction is neither strategy-proof nor static, but it is credible. In a Dutch auction, the price falls until one bidder claims the object. Thus, each bidder only sees a sequence of descending prices $(p_i^1, p_i^2, p_i^3, \dots)$; once he claims the object, he wins at that price. Consequently, once one bidder makes a claim, it is not safe to deviate—the auctioneer must sell to that bidder at his current price. Fixing S_N , each bidder has a claim-price $p_i(\theta_i)$ at which he will agree. For a given θ_N , the rule-following auctioneer strategy yields revenue $\max_{i \in N} p_i(\theta_i)$. No safe deviation results in bidder i paying more than $p_i(\theta_i)$, so the revenue from following the rules first-order stochastically dominates the revenue from any safe deviation.

4. EXTENSIONS

In the Supplemental Material (Akbarpour and Li (2020)), we study a number of extensions to the benchmark model.

First, we relax the assumption that bidders' types are independent, so that the optimal auction extracts full surplus (Cremer and McLean (1988)). Static Cremer–McLean mechanisms are not credible, since two type profiles with the same winning bidder may have

³⁰This format is closely related to the 'all-pay' procurement auctions studied in Dequiedt and Martimort (2015).

different profiles of transfers. Even using extensive-form mechanisms does not in general allow credible full-surplus extraction.

Next, we assume symmetric and affiliated type spaces, and constrain the auctioneer to use ex post incentive-compatible and ex post individually rational mechanisms. In this setting, a modified ascending auction is optimal (Roughgarden and Talgam-Cohen (2013)), and is also credible.

Finally, we assume independent private values, and relax the assumption that there is a single object for sale. Instead, the feasible sets of winning bidders are a matroid. We prove that there exists a credible strategy-proof optimal auction.

5. ALTERNATIVE DEFINITIONS

5.1. Group-Credible Mechanisms

Our main purpose in this paper is to study auctioneer incentives under private communication. Nonetheless, it is natural to consider what happens under other communication structures. Here, we develop an extension that permits agents to share information in groups, and show that increasing information-sharing makes it harder for the auctioneer to deviate.

Essentially, we partition agents into groups in advance, and permit each group of agents to share information after the auction, so that the auctioneer can only hide deviations by misrepresenting the behavior of other groups. Let Λ be a partition on N , and let λ denote a cell of Λ .

DEFINITION 18: Given some promised strategy profile (S_0, S_N) , auctioneer strategy \hat{S}_0 is Λ -safe if, for all groups $\lambda \in \Lambda$ and all type profiles θ_N , there exists $\hat{\theta}_{-\lambda}$ such that for all $i \in \lambda$, $o_i(\hat{S}_0, S_N, \theta_N) = o_i(S_0, S_N, (\theta_\lambda, \hat{\theta}_{-\lambda}))$. $\mathcal{S}_0^\Lambda(S_0, S_N)$ denotes the set of Λ -safe strategies.

Definition 18 permits the auctioneer to deviate only if every group’s observations have an **innocent explanation**; there must exist $\hat{\theta}_{-\lambda}$ such that all observations by agents in λ are consistent with the auctioneer playing S_0^G , the agents playing S_N , and the other groups’ types being $\hat{\theta}_{-\lambda}$. Notably, the order of quantifiers in Definition 18 requires a single explanation to be offered to the entire group, which is more demanding than if we permit each observation in the group to have a different explanation.

Coarser partitions imply more information-sharing between agents.

DEFINITION 19: (G, S_N) is Λ -credible if

$$S_0^G \in \arg \max_{S_0 \in \mathcal{S}_0^\Lambda(S_0^G, S_N)} \mathbb{E}_{\theta_N} [u_0(S_0, S_N, \theta_N)].$$

PROPOSITION 5: If Λ is coarser than Λ' and (G, S_N) is Λ' -credible, then (G, S_N) is Λ -credible.

PROOF: We will prove that, if Λ is coarser than Λ' , then $\mathcal{S}_0^\Lambda(S_0^G, S_N) \subseteq \mathcal{S}_0^{\Lambda'}(S_0^G, S_N)$. From that, Proposition 5 follows immediately.

Take any $\hat{S}_0 \in \mathcal{S}_0^\Lambda(S_0^G, S_N)$, any group $\lambda' \in \Lambda'$, and any θ_N . Since Λ is coarser than Λ' , we can find a group $\lambda \in \Lambda$ such that $\lambda \supseteq \lambda'$. Let $\hat{\theta}_{-\lambda}$ be such that, for all $i \in \lambda$, $o_i(\hat{S}_0, S_N, \theta_N) = o_i(S_0^G, S_N, (\theta_\lambda, \hat{\theta}_{-\lambda}))$. Observe that $(\hat{\theta}_{-\lambda}, \theta_{\lambda \setminus \lambda'})$ is an innocent explanation for group λ' at type profile θ_N . Thus, $\hat{S}_0 \in \mathcal{S}_0^{\Lambda'}(S_0^G, S_N)$. Q.E.D.

One interpretation of Proposition 5 is that starting with a Λ -credible mechanism and increasing information-sharing does not undermine auctioneer incentives. Equivalently, starting with a mechanism that is not Λ -credible and reducing information-sharing does not restore auctioneer incentives. When Λ is the finest partition, then Definition 19 is equivalent to Definition 4.

The second-price auction is not Λ -credible, unless Λ is the coarsest partition. If even a single bidder is unwilling to share information about his bids, then the auctioneer can profitably deviate by misrepresenting that bidder's behavior.

5.2. A 'Prior-Free' Definition

The definition of credibility depends on the joint distribution of agent types (Definition 4). It may be useful to have a definition that is 'prior-free,' for settings such as matching or maxmin mechanism design.

DEFINITION 20: Given (G, S_N) , $S_0 \in \mathcal{S}_0^*(S_0^G, S_N)$ is **always-profitable** if, for all θ_N ,

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

with strict inequality for some θ_N .

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

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

With continuum type spaces, credibility neither implies nor is implied by prior-free credibility. This is because some always-profitable deviations are strictly profitable only on a zero-measure set.

Replacing credibility with prior-free credibility does not essentially change any of our characterizations. Indeed, for the continuum type spaces, requiring prior-free credibility sharpens the results, since it pins down the payment rule even on measure-zero sets:

THEOREM 7: *Assume the continuum type-space model of Section 3.1.*

Assume that (G, S_N) is winner-paying and optimal. (G, S_N) is prior-free credible and static if and only if (G, S_N) is a first-price auction.

With finite type spaces, prior-free credibility is a weaker requirement than credibility. Nonetheless, prior-free credibility is enough to pin down the extensive form of the ascending auction:

THEOREM 8: *Assume the finite type-space model of Section 3.2.*

Assume that F_N is regular and symmetric and that (G, S_N) is optimal. If (G, S_N) is an ascending auction, then (G, S_N) is prior-free credible and strategy-proof.

Assume additionally that (G, S_N) is concise and orderly. If (G, S_N) is prior-free credible and strategy-proof, then (G, S_N) is an ascending auction.

6. DISCUSSION

It is worth considering why real-world auctioneers might lack full commitment power. Vickrey (1961) suggested 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.³¹ 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.³²

Many real-world auctions allow for private communication. There are several reasons for this practice. First, bidders frequently desire privacy for reasons both intrinsic and strategic. A mobile operator may be unwilling to publicize its value for a band of spectrum, because its rivals will take advantage of this information. In recent spectrum auctions in Ireland, the Netherlands, Austria, and Switzerland, the auctioneer did not disclose the losing bids, even after the auction (Dworczak (2020)). Second, auctioneers want to prevent collusion. Thus, in many procurement auctions, bidders are forbidden from conferring—they must submit their bids only to the auctioneer. Third, in auctions that take place over the Internet, bidders are often anonymous to each other, which prevents them from sharing information.

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.³³ However, 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, Thorpe, and Li (2015)).³⁴ 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

³¹As of January 2020, Christie's and Sotheby's use three-piece functions, with the auctioneer's share starting at 25%, then falling to 20%, and then 13.5%. For most categories, eBay charges 10% of the sale price, up to a maximum of US\$750.

³²As Myerson (2009) observed, "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."

³³However, 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.

³⁴Bidders may even need special training or software assistance to play their part in a cryptographic protocol.

in return for a guarantee that his report completely determines the price he might pay.³⁵ 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 designs. How this principle extends to other environments is an open question.

APPENDIX A: DEFINITION OF EXTENSIVE GAME FORMS WITH CONSEQUENCES IN X

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

1. H is a set of histories, along with a binary relation $<$ on H that represents precedence:

- (a) $<$ is a partial order, and $(H, <)$ form an arborescence,³⁶ with root denoted h_\emptyset .
- (b) We use $h \preceq h'$ if $h = h'$ or $h < h'$.
- (c) $Z \equiv \{h \in H : \neg \exists h' : h < h'\}$.
- (d) $\text{succ}(h)$ 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 $\text{succ}(h)$.
- (b) $A(h)$ denotes the actions available at h :

$$A(h) \equiv \bigcup_{h' \in \text{succ}(h)} A(h').$$

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. $g : Z \rightarrow X$.

APPENDIX B: PROOFS OMITTED FROM THE MAIN TEXT

B.1. *Proposition 1*

Suppose that (G, S_N) does not satisfy Clause 1 of Definition 2. We can modify (G, S_N) so that it satisfies Clause 1, remains BIC, and results in the same outcomes for each type profile.

In particular, suppose there exists h such that there is no θ_N such that $h \preceq z(S_N(\theta_N))$. Since the game tree has finite depth, 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 $(<', \mathcal{A}', P', (\mathcal{I}'_i)_{i \in N}, g')$ and S'_N so that they are as in G , but restricted to the new smaller

³⁵This property is generalized in a natural way by the ‘first-price’ menu auction (Bernheim and Whinston (1986)).

³⁶That is, a directed rooted tree such that every edge points away from the root.

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 each type θ_i preferred $S_i(\theta_i)$ to any interim strategy that played a' at I'_i , his new interim strategy $S'_i(\theta_i)$ remains incentive-compatible. Thus, the transformed (G', S'_N) is BIC. We do this for all such histories simultaneously, to produce a protocol that satisfies Clause 1.

Suppose that (G, S_N) satisfies Clause 1 but not Clause 2. We now modify (G, S_N) so that it satisfies Clause 1 and Clause 2, remains BIC, and results in the same outcomes for each type profile.

Suppose there exists $h \notin Z$ such that $|\text{succ}(h)| = 1$. 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 $H \setminus I_i$. We do this for all singleton-action histories simultaneously, to produce a protocol that satisfies Clause 1 and Clause 2.

We now take (G, S_N) that satisfies Clauses 1 and 2, and transform it to satisfy Clause 3. Informally, our argument proceeds as follows: Suppose there is some h at which Clause 3 is not satisfied, where we denote $i = P(h)$. Upon reaching h , i 's continuation strategy no longer affects the outcome. Consider a modified protocol (G', S'_N) : Play proceeds exactly as in (G, S_N) , except after history h is reached. Whenever, under (G, S_N) , i would be called to play at h' where $h \leq h'$, we instead skip i 's turn and continue play as though i chose the action that would be selected by some type θ_i .

Formally, suppose Clauses 1 and 2 hold for (G, S_N) , but 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. $x^G(S_N(\theta_i, \theta_{-i})) \neq x^G(S_N(\theta'_i, \theta_{-i}))$.

Since Clause 1 holds, 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 . First, we delete any history h' such that $h \leq h'$ and $P(h') = i$; this ensures that i is no longer called to play after h . Next, we delete any history h' such that $h \leq h'$ and there does not exist θ''_{-i} such that $h' \leq z(S_N(\theta_i, \theta''_{-i}))$; this has the effect of 'automating' play as though i has type θ_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' . We perform this deletion simultaneously for all histories that violate Clause 3.

By construction, for all θ'_i , if i is playing as though his type is θ'_i and we would have reached some deleted 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 a deleted history under (G, S_N) , then the outcomes under (G, S_N) and (G', S'_N) are identical. Thus, (G', S'_N) is BIC, satisfies Clauses 1, 2, and 3, and results in the same outcomes for each type profile.

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 are not reached at any type profile). Suppose the auctioneer had a profitable safe deviation S'_0 from $S_0^{G'}$. The auctioneer could make that same deviation in the messaging game generated by (G, S_N) (with her play specified arbitrarily after actions that correspond to deleted histories). At every type profile, the agents never reply with actions corresponding to the deleted histories, so the auctioneer's deviation is in $S_0^*(S_0^G, S_N)$.

Consider the second transformation (deleting all histories with singleton action sets). 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 any deleted information set, the auctioneer delays sending the corresponding message until the last possible moment. That is, consider S_0 that is the same as S'_0 , except that:

1. If agent i last received message I_i , and S'_0 specifies that the auctioneer sends I'_i to i , let $(I_i^1, I_i^2, \dots, I_i^K)$ denote the sequence of deleted information sets that i would have encountered between I_i and I'_i under G (this sequence is possibly empty, and is unique by perfect recall). S_0 specifies that the auctioneer first sends $(I_i^1, I_i^2, \dots, I_i^K)$ and then (immediately thereafter) sends I'_i .

2. If agent i last received message I_i , and S'_0 specifies that the auctioneer chooses outcome x , let $(I_i^1, I_i^2, \dots, I_i^K)$ denote the (possibly empty) sequence of deleted information sets that i would have encountered (under G) between I_i and some terminal history $z \succ I_i$ such that $\exists \omega_i \in \Omega_i : \{g(z)\} \cup \{x\} \in \omega_i$. At least one such history exists because S'_0 is a safe deviation. S_0 specifies that the auctioneer sends $(I_i^1, I_i^2, \dots, I_i^K)$ before choosing x . S_0 is a profitable safe deviation from S_0^G .

Consider the third transformation (deleting histories where i is called to play, following any 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 $S_0^{G'}$. The auctioneer can make that same deviation from S_0^G , except that she delays any 'outcome-irrelevant' queries to i until just before she selects the outcome.

Formally, take any θ_N , i , and $\hat{\theta}_{-i}$ such that $o_i(S'_0, S_N, \theta_N) = o_i(S_0^{G'}, S_N, (\theta_i, \hat{\theta}_{-i}))$. If $o_i(S_0^{G'}, S_N, (\theta_i, \hat{\theta}_{-i})) \neq o_i(S_0^G, S_N, (\theta_i, \hat{\theta}_{-i}))$, then this can only be because $o_i(S_0^{G'}, S_N, (\theta_i, \hat{\theta}_{-i}))$ contains additional communication at the end of the sequence that corresponds to deleted histories at which i is called to play. Let h be the earliest such deleted history that would be encountered under (G, S_N) at type profile $(\theta_i, \hat{\theta}_{-i})$. We can 'fill in' the missing communication for agent i , as follows. Initialize $\hat{h} := h$:

1. If $\hat{h} \in Z$, then terminate.
2. Else if $P(\hat{h}) \neq i$, then for $I_{P(\hat{h})}$ such that $\hat{h} \in I_{P(\hat{h})}$:
 - (a) $\hat{h} := h' \mid h' \in \text{succ}(\hat{h})$ and $S_{P(\hat{h})}(I_{P(\hat{h})}, \hat{\theta}_{P(\hat{h})}) = \mathcal{A}(h')$.
 - (b) Go to step 1.
3. Else:
 - (a) Send (to agent i) message I_i such that $\hat{h} \in I_i$.
 - (b) Upon receiving reply a , choose $\hat{h} := h' \mid \mathcal{A}(h') = a$ and $h' \in \text{succ}(\hat{h})$.
 - (c) Go to step 1.

Since (under S_N) i 's play in the deleted histories makes no difference to the outcome, delaying communication with i until the outcome is about to be selected results in a safe deviation. Thus, whenever S'_0 would select an outcome, we can run the above algorithm for every agent whose resulting observation would not have an innocent explanation, and then select the same outcome, thus producing a profitable safe deviation from S_0^G . This completes the proof of Proposition 2.

B.3. Theorem 3

Suppose (G, S_N) is a twin-bid auction. (G, S_N) is static by definition. Given any profile of bids $(b_i^W, b_i^L)_{i \in N}$, every safe deviation charges b_i^W if bidder i wins and b_i^L if he loses, so the auctioneer prefers S_0^G to any safe deviation. Thus, (G, S_N) is credible.

Suppose (G, S_N) is credible and static. By Theorem 1, there exists a function $\tilde{b}_i^W : \Theta_i \rightarrow \mathbb{R}$ such that, almost everywhere in Θ_N , if type θ_i wins, then i pays $\tilde{b}_i^W(\theta_i)$.

By Lemma 1, let $(\underline{\theta}_{-i}^k)_{k=1}^\infty$ be a countable subset such that, for all θ_{-i} , $\inf_k \tilde{t}_i(\cdot, \underline{\theta}_{-i}^k) \leq \tilde{t}_i(\cdot, \theta_{-i})$ almost everywhere in Θ_i .

Define

$$\tilde{t}_i^L(\theta_i, \theta_{-i}) = \begin{cases} \tilde{t}_i(\theta_i, \theta_{-i}) & \text{if } \tilde{y}(\theta_i, \theta_{-i}) \neq i, \\ \inf_k \tilde{t}_i(\theta_i, \underline{\theta}_{-i}^k) - 1 & \text{otherwise.} \end{cases}$$

Intuitively, the function constructed above ‘penalizes’ the auctioneer’s revenue from i unless the type profile is consistent with i losing.

Since $\tilde{y}(\cdot, \theta_{-i})$, $\tilde{t}_i(\cdot, \theta_{-i})$, and $\inf_k \tilde{t}_i(\cdot, \underline{\theta}_{-i}^k)$ are measurable, it follows that $\tilde{t}_i^L(\cdot, \theta_{-i})$ is measurable. Again applying Lemma 1, let $(\bar{\theta}_{-i}^k)_{k=1}^\infty$ be a countable subset of opponent type profiles such that, for all θ_{-i} , $\sup_k \tilde{t}_i^L(\cdot, \bar{\theta}_{-i}^k) \geq \tilde{t}_i^L(\cdot, \theta_{-i})$ almost everywhere in Θ_i .

We now assert that, almost everywhere in Θ_N , if type θ_i does not win, then that type is charged $\tilde{b}_i^L(\theta_i) = \sup_k \tilde{t}_i^L(\cdot, \bar{\theta}_{-i}^k)$. Suppose the set

$$\left\{ \theta_N \mid \tilde{y}(\theta_N) \neq i \text{ and } \tilde{t}_i(\theta_i, \theta_{-i}) \neq \sup_k \tilde{t}_i^L(\theta_i, \bar{\theta}_{-i}^k) \right\}$$

has positive measure. Observe that for (θ_i, θ_{-i}) in the above set, $\tilde{t}_i(\theta_i, \theta_{-i}) = \tilde{t}_i^L(\theta_i, \theta_{-i})$. Consequently, the set

$$Q = \left\{ \theta_N \mid \tilde{y}(\theta_N) \neq i \text{ and } \inf_k \tilde{t}_i(\theta_i, \underline{\theta}_{-i}^k) \leq \tilde{t}_i(\theta_i, \theta_{-i}) < \sup_k \tilde{t}_i^L(\theta_i, \bar{\theta}_{-i}^k) \right\}$$

has positive measure. Q is measurable with respect to the equilibrium action profiles.

We now construct a profitable safe deviation. Fix some finite K . If the bidders’ chosen actions are consistent with any type profile $(\theta_i, \theta_{-i}) \in Q$, charge $\max\{\tilde{t}_i(\theta_i, \theta_{-i}), \max_{k \leq K} \tilde{t}_i^L(\theta_i, \bar{\theta}_{-i}^k)\}$, without changing the allocation or the other bidders’ transfers. Otherwise, play according to S_0^G . The resulting transfer $\tilde{t}_i^K : \Theta_N \rightarrow \mathbb{R}$ is measurable. Notice that our construction of Q and $\tilde{t}_i^L(\cdot)$ means that we charge more than $\tilde{t}_i(\theta_i, \theta_{-i})$ only if $\max_{k \leq K} \tilde{t}_i^L(\theta_i, \bar{\theta}_{-i}^k)$ is consistent with i losing.

For K large enough, this deviation is profitable. In particular, for all $(\theta_i, \theta_{-i}) \in Q$, $\tilde{t}_i^K(\theta_i, \theta_{-i})$ is non-decreasing in K and converges as $K \rightarrow \infty$ to $\sup_k \tilde{t}_i^L(\theta_i, \bar{\theta}_{-i}^k)$. Thus, by the monotone convergence theorem,

$$\begin{aligned} \lim_{K \rightarrow \infty} \mathbb{E}_{\theta_N} [\tilde{t}_i^K(\theta_N) \mid \theta_N \in Q] &= \mathbb{E}_{\theta_N} \left[\sup_k \tilde{t}_i^L(\theta_i, \bar{\theta}_{-i}^k) \mid \theta_N \in Q \right] \\ &> \mathbb{E}_{\theta_N} [\tilde{t}_i(\theta_N) \mid \theta_N \in Q], \end{aligned}$$

which establishes that the deviation is profitable.

We have shown that there exist $\tilde{b}_i^W : \Theta_i \rightarrow \mathbb{R}$ and $\tilde{b}_i^L : \Theta_i \rightarrow \mathbb{R}$ such that, almost everywhere in θ_N , i pays $\tilde{b}_i^W(\theta_i)$ if $\tilde{y}(\theta_i, \theta_{-i}) = i$ and $\tilde{b}_i^L(\theta_i)$ if $\tilde{y}(\theta_i, \theta_{-i}) \neq i$. If, for all θ_{-i} , $\tilde{y}(\theta_i, \theta_{-i}) \neq i$, then we set $\tilde{b}_i^W(\theta_i)$ to be equal to $\tilde{b}_i^L(\theta_i) - 1$. We then define $B_i = \{(\tilde{b}_i^W(\theta_i), \tilde{b}_i^L(\theta_i)) \mid \theta_i \in \Theta_i\}$, which implies that Clause 1 of Definition 10 holds almost everywhere. Let Y denote the subset of Θ_N on which Clause 1 holds.

Suppose then that Clause 2 does not hold on a positive measure set. Then, for some bidder i , the set

$$\{\theta_N \mid \tilde{y}(\theta_N) = 0 \text{ and } \tilde{b}_i^W(\theta_i) - \tilde{b}_i^L(\theta_i) > 0\} \cap Y$$

has positive measure. The auctioneer can raise expected revenue by deviating at all type profiles in this set, allocating the object to i and charging $\tilde{b}_i^W(\theta_i)$. Thus Clause 2 holds almost everywhere.

Suppose then that Clause 3 does not hold on a positive measure set. (G, S_N) is contestable, so for some bidder i , the set

$$Q' = \left\{ \theta_N \mid \tilde{y}(\theta_N) = i \text{ and } \tilde{b}_i^W(\theta_i) - \tilde{b}_i^L(\theta_i) < \max_{j \neq i} \left\{ 0, \tilde{b}_j^W(\theta_j) - \tilde{b}_j^L(\theta_j) \right\} \right. \\ \left. \text{and } \exists \theta'_{-i} : \tilde{y}(\theta_i, \theta'_{-i}) \neq i \right\} \cap Y$$

has positive measure. The auctioneer can raise expected revenue by deviating at all type profiles in this set. Take any type profile in $\theta_N \in Q'$:

1. If $\tilde{b}_i^W(\theta_i) - \tilde{b}_i^L(\theta_i) < 0$, then keep the object and changes i 's payment to $\tilde{b}_i^L(\theta_i)$.
2. Else, if $\tilde{b}_i^W(\theta_i) - \tilde{b}_i^L(\theta_i) < \max_{j \neq i} \tilde{b}_j^W(\theta_j) - \tilde{b}_j^L(\theta_j)$, then award the object to the bidder who maximizes the right-hand side, changes i 's payment to $\tilde{b}_i^L(\theta_i)$ and the other bidder's payment to $\tilde{b}_j^W(\theta_j)$.

Hence, Clause 3 holds almost everywhere, which completes the proof.

B.4. Proposition 3

Suppose (G, S_N) is a twin-bid auction and strategy-proof. Strategy-proofness requires

$$\tilde{t}_i(\theta'_i, \theta_{-i}) - \tilde{t}_i(\theta_i, \theta_{-i}) \leq \theta'_i, \tag{2}$$

$$\theta''_i \leq \tilde{t}_i(\theta'''_i, \theta_{-i}) - \tilde{t}_i(\theta'_i, \theta_{-i}). \tag{3}$$

$\tilde{y}(\cdot)$ is non-decreasing in θ_i , so $\tilde{y}(\theta_i, \theta_{-i}) \neq i$ and $\tilde{y}(\theta'''_i, \theta_{-i}) = i$. It follows that

$$\tilde{t}_i(\theta_i, \theta_{-i}) = \tilde{t}_i(\theta_i, \theta'_{-i}) = \tilde{t}_i(\theta'_i, \theta'_{-i}),$$

$$\tilde{t}_i(\theta'''_i, \theta'_{-i}) = \tilde{t}_i(\theta'''_i, \theta_{-i}) = \tilde{t}_i(\theta'_i, \theta_{-i}),$$

where the first equality in each line follows from the definition of a twin-bid auction and the second equality follows from strategy-proofness. Substituting into Equation (3) yields

$$\theta''_i \leq \tilde{t}_i(\theta'_i, \theta_{-i}) - \tilde{t}_i(\theta_i, \theta_{-i}),$$

which contradicts Equation (2).

B.5. Theorem 4

B.5.1. Ascending \rightarrow Credible, Strategy-Roof

Here we show that if (G, S_N) is optimal and an ascending auction, then it is strategy-proof and credible. Essentially, we will exploit the fact that ascending auctions are obviously strategy-proof (Li (2017)); that is, starting from any information set at which bidder i deviates from S_i , every possible outcome from the deviation is no better than every possible outcome from S_i . This implies that (G, S_N) is strategy-proof, and also implies that S_i is a best response to any safe deviation by the auctioneer.

Since we are holding fixed (G, S_N) , we will drop the superscripts on \tilde{y}^{G, S_N} and \tilde{t}_i^{G, S_N} to reduce clutter.

LEMMA 2: *If (G, S_N) is an ascending auction, then (G, S_N) is strategy-proof.*

We will prove that, given any ascending auction (G, S_N) , for all interim strategies σ_i and all S'_{-i} , θ_i , and θ_{-i} ,

$$u_i^G(S_i(\theta_i), S'_{-i}(\theta_{-i}), \theta_i) \geq u_i^G(\sigma_i, S'_{-i}(\theta_{-i}), \theta_i), \tag{4}$$

which implies that (G, S_N) is strategy-proof.

Consider the paths of play induced by $(S_i(\theta_i), S'_{-i}(\theta_{-i}))$ and $(\sigma_i, S'_{-i}(\theta_{-i}))$. If these are identical, then Equation (4) holds trivially. Otherwise, let h be any history at which these paths diverge, that is, the history at which $S_i(\theta_i)$ and σ_i choose different actions for the first time. There are three cases to consider.

Case 1: Suppose that at h , σ_i plays a quitting action. After the quitting action, i does not win and pays 0, so $u_i^G(\sigma_i, S'_{-i}(\theta_{-i}), \theta_i) = 0$. Since (G, S_N) is pruned, we can find θ'_{-i} such that $(S_i(\theta_i), S'_i(\theta_{-i}))$ and $(S_i(\theta_i), S_i(\theta'_{-i}))$ result in the same path-of-play. $u_i^G(S_i(\theta_i), S'_{-i}(\theta_{-i}), \theta_i) = u_i^G(S_i(\theta_i), S_{-i}(\theta'_{-i}), \theta_i) \geq 0$, where the inequality follows since (G, S_N) is orderly and has threshold pricing.

Case 2: Suppose that at h , σ_i plays an accepting action, and $S_i(\theta_i)$ quits. Then $u_i^G(S_i(\theta_i), S'_{-i}(\theta_{-i}), \theta_i) = 0$. Since (G, S_N) is pruned, we can find $(\theta'_i, \theta'_{-i})$ such that $(\sigma_i, S'_{-i}(\theta_{-i}))$ and $(S_i(\theta'_i), S_{-i}(\theta'_{-i}))$ result in the same path-of-play, so

$$u_i^G(\sigma_i, S'_{-i}(\theta_{-i}), \theta_i) = u_i^G(S_i(\theta'_i), S_{-i}(\theta'_{-i}), \theta_i). \tag{5}$$

(G, S_N) is orderly and $S_i(\theta_i)$ specifies that i quits at h , so $\tilde{y}(\theta_i, \theta'_{-i}) \neq i$. (G, S_N) has threshold pricing, so if $\tilde{y}(\theta'_i, \theta'_{-i}) \neq i$, then $u_i^G(S_i(\theta'_i), S_{-i}(\theta'_{-i}), \theta_i) = 0$. If $\tilde{y}(\theta'_i, \theta'_{-i}) = i$, then by threshold pricing, $\tilde{t}_i(\theta'_i, \theta'_{-i}) > \theta_i$, so $u_i^G(S_i(\theta'_i), S_{-i}(\theta'_{-i}), \theta_i) < 0$.

Case 3: Suppose that at h , σ_i plays an accepting action, and $S_i(\theta_i)$ plays a distinct accepting action. Since there is more than one accepting action, we must be in the case governed by Clause 5(b) of Definition 14. Thus, both actions result in i winning at the current offered price p , and yield the same utility.

This completes the proof of Lemma 2.

LEMMA 3: *Let (G, S_N) be an ascending auction. For every bidder i , if S'_0 is a safe deviation, then S_i is a best response to (S'_0, S_{-i}) .*

Let S'_0 be an arbitrary safe deviation. Take any type θ_i . Suppose that $S_i(\theta_i)$ and deviating strategy $S'_i(\theta_i)$ choose different actions for the first time after receiving message I_i . There

are three cases to consider; we will show that, in each case, $S'_i(\theta_i)$ is not a profitable deviation.

Case 1: Upon receiving message I_i , $S'_i(\theta_i)$ quits. S'_0 is safe, so $S'_i(\theta_i)$ results in zero utility after receiving I_i . S'_0 is safe and (G, S_N) is orderly and has threshold pricing, so $S_i(\theta_i)$ results in at least zero utility.

Case 2: Upon receiving message I_i , $S_i(\theta_i)$ quits and $S'_i(\theta_i)$ accepts. This results in some observation $o_i(S'_0, (S'_i, S_{N \setminus i}), \theta_N)$, which pins down whether i wins and how much i pays. (G, S_N) is pruned, so we can find θ'_i such that $o_i(S'_0, (S'_i, S_{N \setminus i}), \theta_N) = o_i(S'_0, S_N, (\theta'_i, \theta_{-i}))$. S'_0 is safe, so we can find θ'_{-i} such that $o_i(S'_0, S_N, (\theta'_i, \theta_{-i})) = o_i(S'_0, S_N, (\theta'_i, \theta'_{-i}))$. If $\tilde{y}(\theta'_i, \theta'_{-i}) \neq i$, then since (G, S_N) has threshold pricing, i does not profit by deviating. Suppose $\tilde{y}(\theta'_i, \theta'_{-i}) = i$. (G, S_N) is orderly and S_i specifies that θ_i quits at I_i , so since θ'_{-i} is consistent with reaching I_i , $\tilde{y}(\theta_i, \theta'_{-i}) \neq i$. (G, S_N) has threshold pricing, so $\tilde{t}_i(\theta'_i, \theta'_{-i}) > \theta_i$. Thus, i does not profit by deviating in this case.

Case 3: Upon receiving message I_i , $S_i(\theta_i)$ accepts a price, and $S'_i(\theta_i)$ accepts a price, with a distinct action. Then I_i satisfies Clause 5(b) of Definition 14. S'_0 is safe, so both actions result in i winning at the offered price p .

This completes the proof of Lemma 3.

Suppose now that (G, S_N) is an ascending auction but 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' . By Lemma 3, for all i , S_i is a best response to (S'_0, S_{-i}) , so (G', S_N) is also BIC. (G', S_N) has voluntary participation. (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 $\pi(G', S_N) > \pi(G, S_N)$, so (G, S_N) is not optimal. Thus, if (G, S_N) is optimal and an ascending auction, then (G, S_N) is credible. This completes the proof of the first claim in Theorem 4.

B.5.2. Credible, Strategy-Proof \rightarrow Ascending

We start by deriving several properties of credible strategy-proof optimal (G, S_N) , without assuming that F_N is regular or symmetric.

PROPOSITION 6: *If (G, S_N) is optimal and strategy-proof, then (G, S_N) is winner-paying.*

PROOF: For all (θ_i, θ_{-i}) , if $\tilde{y}(\theta_i, \theta_{-i}) \neq i$, then $\tilde{t}_i(\theta_i, \theta_{-i}) \leq 0$. Suppose not. (G, S_N) satisfies voluntary participation. When i 's opponents imitate θ_{-i} ,³⁷ type θ_i can profitably deviate to non-participation if $\tilde{t}_i(\theta_i, \theta_{-i}) > 0$, contradicting strategy-proofness.

$\theta_i^1 \leq 0$, so $\eta_i(\theta_i^1) < 0$. (G, S_N) is optimal, so θ_i^1 never wins (by Proposition 4). θ_i^1 's participation constraint binds, so for all θ_{-i} , $\tilde{t}_i(\theta_i^1, \theta_{-i}) = 0$.

Take any (θ_i, θ_{-i}) . If $\tilde{y}(\theta_i, \theta_{-i}) \neq i$ and $\tilde{t}_i(\theta_i, \theta_{-i}) > 0$, then when i 's opponents imitate θ_{-i} , θ_i^1 can profitably imitate θ_i , contradicting strategy-proofness. Thus, (G, S_N) is winner-paying. Q.E.D.

PROPOSITION 7: *If (G, S_N) is strategy-proof, then the allocation rule is monotone. That is, if $\theta_i < \theta'_i$ and $\tilde{y}(\theta_i, \theta_{-i}) = i$, then $\tilde{y}(\theta'_i, \theta_{-i}) = i$.*

PROOF: Suppose not, so $\tilde{y}(\theta'_i, \theta_{-i}) \neq i$. By strategy-proofness, $-\tilde{t}_i(\theta'_i, \theta_{-i}) \geq \theta'_i - \tilde{t}_i(\theta_i, \theta_{-i})$, which implies $-\tilde{t}_i(\theta'_i, \theta_{-i}) > \theta_i - \tilde{t}_i(\theta_i, \theta_{-i})$, so θ_i can profitably imitate θ'_i , a contradiction. Q.E.D.

³⁷Formally, define S'_j such that for all $j \neq i$, I_j , and θ'_j , $S'_j(I_j, \theta'_j) = S_j(I_j, \theta_j)$

DEFINITION 21: (G, S_N) has **threshold pricing** if

$$\tilde{t}_i(\theta_N) = \begin{cases} \min\{\theta'_i \in \Theta_i \mid \tilde{y}(\theta'_i, \theta_{-i}) = i\} & \text{if } \tilde{y}(\theta_N) = i, \\ 0 & \text{otherwise.} \end{cases} \tag{6}$$

PROPOSITION 8: *If (G, S_N) is optimal and strategy-proof, then (G, S_N) has threshold pricing.*

PROOF: Proposition 6 pins down the payments whenever $\tilde{y}(\theta_N) \neq i$.

We prove the rest by induction. (G, S_N) is optimal, so θ_i^1 's participation constraint binds. Thus, Equation (6) holds when for θ_i^1 . Suppose that Equation (6) holds for all $\theta_i^{k'}$ such that $k' \leq k$. We prove it holds for θ_i^{k+1} .

Take any θ_{-i} . There are three cases to consider.

If $\tilde{y}(\theta_i^k, \theta_{-i}) = i$, then strategy-proofness implies that $\tilde{y}(\theta_i^{k+1}, \theta_{-i}) = i$ and $\tilde{t}_i(\theta_i^{k+1}, \theta_{-i}) = \tilde{t}_i(\theta_i^k, \theta_{-i}) = \min_{\theta'_i \in \Theta_i} \theta'_i \mid \tilde{y}(\theta'_i, \theta_{-i}) = i$.

If $\tilde{y}(\theta_i^{k+1}, \theta_{-i}) \neq i$, then $\tilde{t}_i(\theta_i, \theta_{-i}) = 0$.

Notice that, in the previous two cases, θ_i^{k+1} is exactly indifferent between S_i and deviating to imitate type θ_i^k . Finally, suppose $\tilde{y}(\theta_i^k, \theta_{-i}) \neq i$ and $\tilde{y}(\theta_i^{k+1}, \theta_{-i}) = i$. $\tilde{t}_i(\theta_i^{k+1}, \theta_{-i}) \leq \theta_i^{k+1}$, since (G, S_N) is strategy-proof. If $\tilde{t}_i(\theta_i^{k+1}, \theta_{-i}) < \theta_i^{k+1}$, then (G, S_N) is not optimal, since the incentive constraints do not bind locally downward (Proposition 4). Thus, $\tilde{t}_i(\theta_i^{k+1}, \theta_{-i}) = \theta_i^{k+1}$, and the inductive step is proved. Q.E.D.

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,

$$\Theta_i^h = \{\theta_i \mid \forall h', h'' \leq h : [h' \in I_i, h'' \in \text{succ}(h')] \rightarrow [S_i(I_i, \theta_i) = \mathcal{A}(h'')]\}. \tag{7}$$

For $\hat{N} \subseteq N$, let $\Theta_{\hat{N}}^h = \times_{i \in \hat{N}} \Theta_i^h$.

PROPOSITION 9: *If $P(h) \neq i$ and $h' \in \text{succ}(h)$, then $\Theta_i^h = \Theta_i^{h'}$. If $h < h'$, 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 two claims are 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

$$\underline{\theta}_i^h = \min\{\theta_i \in \Theta_i^h\}, \tag{8}$$

$$\bar{\theta}_i^h = \max\{\theta_i \in \Theta_i^h\}. \tag{9}$$

The next proposition states that strategy-proofness constrains what bidders can learn about each others' play midway through the protocol. In essence, it says that if, at some history h where i is called to play, i can affect whether or not θ_j wins, then i cannot (at this information set) rule out the possibility that j 's type is instead some $\theta'_j > \theta_j$.

PROPOSITION 10: *Assume (G, S_N) is optimal and strategy-proof. Take any information set I_i and history $h \in I_i$. Take any $\theta_i, \theta'_i \in \Theta_i^h, \theta_j \in \Theta_j^h$, and $\theta_{N \setminus \{i,j\}} \in \Theta_{N \setminus \{i,j\}}^h$.*

If $\tilde{y}(\theta_i, \theta_j, \theta_{N \setminus \{i,j\}}) = j$ and $\tilde{y}(\theta'_i, \theta_j, \theta_{N \setminus \{i,j\}}) \neq j$, then $\forall \theta'_j > \theta_j : \exists h' \in I_i : \theta'_j \in \Theta_j^{h'}$ and $\theta_{N \setminus \{i,j\}} \in \Theta_{N \setminus \{i,j\}}^{h'}$.

PROOF: Suppose not. We construct a strategy profile S'_{-j} such that θ'_j has a profitable deviation. For $l \in N \setminus \{i, j\}$, let l imitate θ_i ; that is, $\forall I_l : \forall \theta'_l : S'_l(I_l, \theta'_l) = S_l(I_l, \theta_l)$. Let i imitate θ'_i unless he encounters I_i , and let him imitate type θ_i if he has encountered I_i . Formally:

$$\forall I'_i : \forall \theta''_i : S'_i(I'_i, \theta''_i) = \begin{cases} S_i(I'_i, \theta_i) & \text{if } \exists h'' \in I'_i : \exists h''' \in I_i : h''' \preceq h'', \\ S_i(I'_i, \theta'_i) & \text{otherwise.} \end{cases} \tag{10}$$

By Proposition 8, (G, S_N) has threshold pricing. If type θ'_j deviates to imitate θ_j , then (when facing S'_j), the path-of-play passes through I_i , so j wins at price $\min_{\theta'_j \in \Theta_j} \theta'_j \mid \tilde{y}(\theta_i, \theta'_j, \theta_{N \setminus \{i, j\}}) = j$, for a positive surplus since $\theta'_j > \theta_j$. On the other hand, if type θ'_j plays according to S_j , then the path-of-play does not pass through I_i , so j either wins at a strictly higher price $\min_{\theta'_j \in \Theta_j} \theta'_j \mid \tilde{y}(\theta'_i, \theta'_j, \theta_{N \setminus \{i, j\}}) = j$, or does not win and has zero surplus. Thus, j has a profitable deviation, and (G, S_N) is not strategy-proof, a contradiction. *Q.E.D.*

Let W_i^h denote the subset of i 's types that might reach h and then win. Similarly, let L_i^h denote the subset of i 's types that might reach h and then lose:

$$W_i^h = \{ \theta_i \in \Theta_i^h \mid \exists \theta_{-i} \in \Theta_{-i}^h : \tilde{y}(\theta_i, \theta_{-i}) = i \}, \tag{11}$$

$$L_i^h = \{ \theta_i \in \Theta_i^h \mid \exists \theta_{-i} \in \Theta_{-i}^h : \tilde{y}(\theta_i, \theta_{-i}) \neq i \}. \tag{12}$$

DEFINITION 22: (G, S_N) is **winner-pooling** if, for all $I_i, h \in I_i$:

1. Either: $\forall \theta_i, \theta'_i \in W_i^h : S_i(I_i, \theta_i) = S_i(I_i, \theta'_i)$,
2. Or: $W_i^h \cap L_i^h = \emptyset$.

PROPOSITION 11: Assume F_N is symmetric and regular, and (G, S_N) is optimal, orderly, credible, and strategy-proof. (G, S_N) is winner-pooling.

Before starting the proof of Proposition 11, we highlight that this is the reason that we have assumed regularity and orderliness in the statement of Theorem 4. Together, regularity and orderliness imply that, if there are two distinct types $\theta_i < \theta'_i$ in W_i^h that do not pool on the same action, then there exists θ_{-i} such that θ_i loses when facing θ_{-i} , but θ'_i wins. This enables us to construct profitable safe deviations for the auctioneer.³⁸

PROOF: Under the assumptions of Proposition 11, we will show that if (G, S_N) is not winner-pooling, then the auctioneer has a profitable safe deviation, so (G, S_N) is not credible.

Let h^* be some history at which the winner-pooling property does not hold; we pick h^* such that, for all $h < h^*$, h is not a counterexample to winner-pooling. Since (G, S_N) is orderly and the winner-pooling property held at all predecessors to h^* , it follows that for all i , either $W_i^{h^*} = \emptyset$ or $W_i^{h^*} = \{ \theta_i \mid \theta_i \triangleright \max_{j \neq i} \theta_j^{h^*} \text{ and } \theta_i \triangleright \rho \}$.

Let i^* denote $P(h^*)$, and $I_{i^*}^*$ the corresponding information set. Since the winner-pooling property does not hold at h^* , $W_{i^*}^{h^*} \cap L_{i^*}^{h^*} \neq \emptyset$ and there exist two distinct actions taken by types in $W_{i^*}^{h^*}$ at $I_{i^*}^*$.

³⁸If type spaces were continuous, regularity would by itself imply the desired property for every optimal allocation rule. However, for discrete types, we need to pick a particular allocation rule—and the orderly one will do.

Since (G, S_N) is orderly, $\min_{i^*} W_{i^*}^{h^*} \in W_{i^*}^{h^*} \cap L_{i^*}^{h^*}$. Define

$$\theta_{i^*}^* = \min_{i^*} \theta_{i^*} \in W_{i^*}^{h^*} \mid S_{i^*}(I_{i^*}^*, \theta_{i^*}^*) \neq S_{i^*}(I_{i^*}^*, \min_{i^*} W_{i^*}^{h^*}). \tag{13}$$

We are going to squeeze extra revenue out of bidder i^* when his type is $\theta_{i^*}^*$: by his actions at h^* , he hints that his type is *more than high enough to win*. Let h^{**} be the immediate successor of h^* that would be reached by $\theta_{i^*}^*$, that is,

$$h^{**} = h \mid h \in \text{succ}(h^*) \quad \text{and} \quad \theta_{i^*}^* \in \Theta_{i^*}^h. \tag{14}$$

Since $W_{i^*}^{h^*} \cap L_{i^*}^{h^*} \neq \emptyset$ and (G, S_N) is orderly, $\{j \in N \mid W_j^{h^*} \neq \emptyset\}$ includes i^* and at least one other bidder. For each $i \in N$, we assign a nemesis:

$$\psi(i) = \max_{i^*} \{j \in N \setminus \{i\} \mid W_j^{h^*} \neq \emptyset\}. \tag{15}$$

By choosing i 's nemesis in this way, we ensure a useful property; given any θ_j , we can find $\theta_{\psi(i)}$ such that i has the same allocation and transfer when the highest opponent type is θ_j and when it is $\theta_{\psi(i)}$. Similarly, given any θ_i , we can find $\theta_{\psi(i)}$ that forces i to pay exactly θ_i if he wins (by threshold pricing). Formally, we say $\theta_{\psi(i)}$ is ***i-equivalent to θ_j*** if

$$\{\theta_i \mid \theta_i \geq \theta_j\} = \{\theta_i \mid \theta_i \geq \theta_{\psi(i)}\}, \tag{16}$$

where \geq is the reflexive order implied by the strict order \triangleright .

Given S_0^G , 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. We start by defining several subroutines for the algorithm.

The algorithm calls the following subroutine: Given some variable \hat{h} that takes values in the set of histories, we can start at the initial value of \hat{h} and communicate with i as though the opponent types were θ_{-i} , updating \hat{h} as we go along. When we do this, we say that we **simulate** θ_{-i} against i starting from \hat{h} , until certain specified conditions are met. Formally:

1. If **[conditions]**, STOP.
2. Else if $P(\hat{h}) \neq i$, set $\hat{h} := h \in \text{succ}(\hat{h}) \mid \theta_{-i} \in \Theta_{-i}^h$.
3. Else if $P(\hat{h}) = i$:
 - (a) Send message $I_i \mid \hat{h} \in I_i$ to i .
 - (b) Upon receiving $r \in A(I_i)$, set $\hat{h} := h \mid (h \in \text{succ}(\hat{h}) \text{ and } \mathcal{A}(h) = r)$.
 - (c) Go to step 1.

The algorithm also calls the following subroutine: Given some history h and some θ_{-i} , where i was called to play at h 's immediate predecessor, we may **find the cousin** of h consistent with θ_{-i} . This is the history that immediately follows from the same information set, is consistent with the action i just took, but is also consistent with the opponent types being θ_{-i} . Formally, let $\text{cousin}(h, \theta_{-i})$ be equal to h' such that $\exists I_i : \exists h'', h'''$:

1. $h'', h''' \in I_i$.
2. $h \in \text{succ}(h'')$.
3. $h' \in \text{succ}(h''')$.
4. $\mathcal{A}(h) = \mathcal{A}(h')$.
5. $\theta_{-i} \in \Theta_{-i}^{h'}$.

Clearly, it is not always possible to find such a history. But we will be careful to prove that $\text{cousin}(h, \theta_{-i})$ is well-defined when we call it.

Our algorithm keeps track of several variables:

1. A best offer, initialized $\beta := \theta_i^*$.
2. A set of ‘active’ bidders, initialized $\hat{N} := N$.
3. The bidder we are currently communicating with, $\hat{i} := i^*$.
4. A simulated history, for each bidder: $\hat{h}_{i^*} := h^{**}$ and for $i \in N \setminus \{i^*\}$, $\hat{h}_i := h^*$.

The algorithm proceeds in three stages. At h^{**} , i^* ’s type could be at least $\theta_{i^*}^*$, but it could also be too low to exploit (if some types not in $W_i^{h^*}$ took the same action as $\theta_{i^*}^*$ at h^*). In Stage 1, we check whether i^* ’s type is at least $\theta_{i^*}^*$. If it is, we set β to be the least type consistent with i^* ’s responses, and go to Stage 2. Otherwise, we lower β appropriately, and proceed to Stage 2. In Stage 2, we cycle through the bidders, updating β to be equal to the highest type we have confirmed so far, until we have found the bidder with the highest type (breaking ties with \triangleright). Finally, in Stage 3, we sell to the bidder with the highest type (if it is above the reserve), at a price greater than or equal to the price in the original protocol. We use $:=$ for the assignment operator, and \in to assign an arbitrary element in the set on the right-hand side.

Stage 1

1. Pick $\theta_{\psi(i^*)}$ that is i^* -equivalent to β .
2. Simulate $(\theta_{\psi(i^*)}, \underline{\theta}_{N \setminus \{i^*, \psi(i^*)\}}^{h^*})$ against i^* starting from \hat{h}_{i^*} , until either $\underline{\theta}_{i^*}^{\hat{h}_{i^*}} \geq \beta$ or $\hat{h}_{i^*} \in Z$.
3. If $\underline{\theta}_{i^*}^{\hat{h}_{i^*}} \geq \beta$, then set $\beta := \underline{\theta}_{i^*}^{\hat{h}_{i^*}}$ and go to Stage 2.
4. Else, set $\hat{N} := \hat{N} \setminus \{i^*\}$, $\beta := \min_{i \neq i^*, \theta_i} \theta_i$ and go to Stage 2.

Stage 2

1. If $\hat{N} = 1$, go to Stage 3.
2. Set $\hat{i} := \{i \in \hat{N} \mid \underline{\theta}_i^{\hat{h}_i} \triangleleft \beta\}$.
3. Pick $\theta_{\psi(\hat{i})}$ that is \hat{i} -equivalent to β .
4. If $(\theta_{\psi(\hat{i})}, \underline{\theta}_{N \setminus \{\hat{i}, \psi(\hat{i})\}}^{h^*}) \notin \Theta_{-\hat{i}}^{\hat{h}_i}$, set $\hat{h}_i := \text{cousin}(\hat{h}_i, (\theta_{\psi(\hat{i})}, \underline{\theta}_{N \setminus \{\hat{i}, \psi(\hat{i})\}}^{h^*}))$.
5. Simulate $(\theta_{\psi(\hat{i})}, \underline{\theta}_{N \setminus \{\hat{i}, \psi(\hat{i})\}}^{h^*})$ against \hat{i} starting from \hat{h}_i , until either $\underline{\theta}_{\hat{i}}^{\hat{h}_i} \geq \beta$ or $\hat{h}_i \in Z$.
6. If $\underline{\theta}_{\hat{i}}^{\hat{h}_i} \geq \beta$, set $\beta := \underline{\theta}_{\hat{i}}^{\hat{h}_i}$ and go to Step 1 of Stage 2.
7. Else, set $\hat{N} := \hat{N} \setminus \{\hat{i}\}$ and go to Step 1 of Stage 2.

Stage 3

1. Set $\hat{i} := i \mid i \in \hat{N}$.
2. Pick $\theta_{\psi(\hat{i})}$ that is \hat{i} -equivalent to β .
3. If $(\theta_{\psi(\hat{i})}, \underline{\theta}_{N \setminus \{\hat{i}, \psi(\hat{i})\}}^{h^*}) \notin \Theta_{-\hat{i}}^{\hat{h}_i}$, set $\hat{h}_i := \text{cousin}(\hat{h}_i, (\theta_{\psi(\hat{i})}, \underline{\theta}_{N \setminus \{\hat{i}, \psi(\hat{i})\}}^{h^*}))$.
4. Simulate $(\theta_{\psi(\hat{i})}, \underline{\theta}_{N \setminus \{\hat{i}, \psi(\hat{i})\}}^{h^*})$ against \hat{i} starting from \hat{h}_i , until $\hat{h}_i \in Z$.
5. Choose the outcome that corresponds to that terminal history, $x = g(\hat{h}_i)$, and terminate.

Since (G, S_N) is orderly, the deviation does not change the allocation. In particular, some bidder \hat{i} is removed from \hat{N} only when we know that $\theta_{\psi(\hat{i})} \triangleright \theta_i$; since $\theta_{\psi(\hat{i})}$ is \hat{i} -

equivalent to β , the latter implies that $\beta \triangleright \theta_i$.³⁹ Moreover, since (G, S_N) is orderly and has threshold pricing (by Proposition 8), the resulting algorithm results in transfers that are always at least as high as the transfers under (G, S_N) . The transfers are strictly higher for at least one type profile, namely $(\theta_{i^*}^*, \underline{\theta}_{-i^*}^{h^*})$. Under (G, S_N) , $\tilde{t}_{i^*}(\theta_{i^*}^*, \underline{\theta}_{-i^*}^{h^*}) = \min W_{i^*}^{h^*}$, whereas under the deviation, i^* 's transfer is $\theta_{i^*}^*$. Thus, the deviation is profitable.

It remains to prove that the deviation is safe. When we first start communicating with any bidder \hat{i} under the deviation, we are simulating opponent types that are consistent with h^* , because the winner-pooling property holds at all histories prior to h^* , and we have chosen the simulated nemesis type $\theta_{\psi(\hat{i})}$ to be in $W_{\psi(\hat{i})}^{h^*}$. (Thus, Step 4 of Stage 2 and Step 3 of Stage 3 are not triggered if this is the first time the deviating algorithm is communicating with that bidder.)

Whenever the deviation communicates with some bidder \hat{i} for a second time, we have to prove that we can find cousins (in Step 4 of Stage 2 and Step 3 of Stage 3) in the way the algorithm requires. Let $\theta_{\psi(\hat{i})}^{\text{old}}$ and β^{old} denote the simulated nemesis type and the best offer from the last time the algorithm communicated with \hat{i} . Let $\theta_{\psi(\hat{i})}^{\text{new}}$ and β^{new} denote the current simulated nemesis type and best offer. Observe that we always revise the nemesis type upwards: $\beta^{\text{old}} \trianglelefteq \beta^{\text{new}}$, so $\theta_{\psi(\hat{i})}^{\text{old}} \leq \theta_{\psi(\hat{i})}^{\text{new}}$. If $\theta_{\psi(\hat{i})}^{\text{old}} = \theta_{\psi(\hat{i})}^{\text{new}}$, we are done, since $(\theta_{\psi(\hat{i})}^{\text{old}}, \underline{\theta}_{N \setminus \{\hat{i}, \psi(\hat{i})\}}^{h^*}) \in \Theta_{-\hat{i}}^{\hat{h}_i}$. Otherwise, consider h' , the immediate predecessor of \hat{h}_i . At h' , \hat{i} is called to play, and it is not yet clear whether $\psi(\hat{i})$ wins. In particular, $\theta_{\psi(\hat{i})}^{\text{old}}$ would win against $\underline{\theta}_i^{h'}$, but would lose against $\underline{\theta}_i^{\hat{h}_i}$, that is, $\tilde{y}(\underline{\theta}_i^{h'}, \theta_{\psi(\hat{i})}^{\text{old}}, \underline{\theta}_{N \setminus \{\hat{i}, \psi(\hat{i})\}}^{h^*}) = \psi(\hat{i}) \neq \tilde{y}(\underline{\theta}_i^{\hat{h}_i}, \theta_{\psi(\hat{i})}^{\text{old}}, \underline{\theta}_{N \setminus \{\hat{i}, \psi(\hat{i})\}}^{h^*})$. By Proposition 10, there exists another history h'' in the same information set as h' , such that $(\theta_{\psi(\hat{i})}^{\text{new}}, \underline{\theta}_{N \setminus \{\hat{i}, \psi(\hat{i})\}}^{h^*}) \in \Theta_{-\hat{i}}^{h''}$. Thus, we can find cousins in the way that the algorithm requires.

Observe that, whenever \hat{i} is removed from \hat{N} , he has seen a communication sequence that is consistent with his reaching a terminal history with an opponent type profile such that \hat{i} does not win and has a zero transfer, and the Stage 3 outcome respects that. At Stage 3, the final bidder \hat{i} 's observation is consistent with $(\theta_{\psi(\hat{i})}, \underline{\theta}_{N \setminus \{\hat{i}, \psi(\hat{i})\}}^{h^*})$. Thus, the algorithm produces a profitable safe deviation. Q.E.D.

PROPOSITION 12: *Assume F_N is symmetric and regular, and (G, S_N) is optimal, orderly, concise, credible, and strategy-proof. For any non-terminal history h and any bidder i :*

1. *If $W_i^h \neq \emptyset$, then $\Theta_i^h \supseteq \{\theta_i \mid \tilde{y}(\theta_i, \underline{\theta}_{-i}) = i\}$.*
2. *For all $\theta_i, \theta'_i \in L_i^h \setminus W_i^h$, for all $\theta_{-i} \in \Theta_{-i}^h$, $\tilde{y}(\theta_i, \theta_{-i}) = \tilde{y}(\theta'_i, \theta_{-i})$.*

Moreover, if $W_{P(h)}^h \cap L_{P(h)}^h = \emptyset$ and $W_{P(h)}^h \neq \emptyset$, then $P(h)$ faces a posted price at h .

PROOF: We will prove by induction. The step- κ inductive hypothesis is: For any non-terminal history h with no more than κ predecessors,

1. *If $W_i^h \neq \emptyset$, then $\Theta_i^h \supseteq \{\theta_i \mid \tilde{y}(\theta_i, \underline{\theta}_{-i}^h) = i\}$.*
2. *For all $\theta_i, \theta'_i \in L_i^h \setminus W_i^h$, for all $\theta_{-i} \in \Theta_{-i}^h$, $\tilde{y}(\theta_i, \theta_{-i}) = \tilde{y}(\theta'_i, \theta_{-i})$.*

The hypothesis holds for $\kappa = 0$, since for the initial history h_\emptyset , for all i , $\Theta_i^{h_\emptyset} = \Theta_i$, and the protocol is orderly.

³⁹Since (G, S_N) is orderly, we must eventually learn either that $\theta_i \triangleright \theta_{\psi(\hat{i})}$ or vice versa, since this information is necessary to determine whether \hat{i} or $\psi(\hat{i})$ should win when the other bidders' types are $\underline{\theta}_{N \setminus \{\hat{i}, \psi(\hat{i})\}}^{h^*}$. Thus, reaching Step 4 of Stage 1 or Step 7 of Stage 2 implies that $\beta \triangleright \theta_i$.

Suppose the inductive hypothesis holds for some κ . We will prove it holds for $\kappa + 1$. Let h be a non-terminal history with κ predecessors. We will prove that the property holds for any immediate successor, that is, any $h' \in \text{succ}(h)$.

There are four cases to consider, and we will make repeated use of Proposition 9.

Case 1: Consider any bidder i such that $W_i^h = \emptyset$. Take any $h' \in \text{succ}(h)$. Then $W_i^{h'} = \emptyset$. Moreover, $L_i^h \setminus W_i^h = \Theta_i^h \supseteq \Theta_i^{h'} = L_i^{h'} \setminus W_i^{h'}$. Thus, by the inductive hypothesis for κ , for all $\theta_i, \theta'_i \in L_i^{h'} \setminus W_i^{h'}$, for all $\theta_{-i} \in \Theta_{-i}^{h'}$, $\tilde{y}(\theta_i, \theta_{-i}) = \tilde{y}(\theta'_i, \theta_{-i})$.

Case 2: Consider any bidder $i \neq P(h)$ such that $W_i^h \neq \emptyset$. Take any $h' \in \text{succ}(h)$. Then, applying the inductive hypothesis for κ and (G, S_N) orderly yields

$$\Theta_i^{h'} = \Theta_i^h \supseteq \{\theta_i \mid \tilde{y}(\theta_i, \underline{\theta}_{-i}^h) = i\} \supseteq \{\theta_i \mid \tilde{y}(\theta_i, \underline{\theta}_{-i}^{h'}) = i\}. \quad (17)$$

If $W_i^h = W_i^{h'}$, then $L_i^h \setminus W_i^h = L_i^{h'} \setminus W_i^{h'}$, so part 2 of the inductive hypothesis holds for i at h' . If $W_i^h \neq W_i^{h'}$, then $W_i^h \supset W_i^{h'}$. Thus, since $\Theta_j^h = \Theta_j^{h'}$ for $j \neq P(h)$ and (G, S_N) is orderly, for all $\theta_i \in L_i^{h'} \setminus W_i^{h'}$, $\theta_i \triangleleft_{P(h')} \underline{\theta}_{-i}^{h'}$. Again by (G, S_N) orderly, part 2 of the inductive hypothesis holds for i at h' .

Case 3: Consider $i = P(h)$, such that $W_i^h \cap L_i^h \neq \emptyset$. Let I_i be such that $h \in I_i$. By Proposition 11, there exists $a^* \in A(I_i)$ such that, for all $\theta_i \in W_i^h$, $S_i(I_i, \theta_i) = a^*$. Let h' be any immediate successor. Since $\Theta_{-i}^h = \Theta_{-i}^{h'}$, $L_i^h \setminus W_i^h \supseteq L_i^{h'} \setminus W_i^{h'}$, so by the inductive hypothesis for κ , the second part of the inductive hypothesis holds for i at h' . If h' is an immediate successor that follows play of a^* , then $\Theta_i^{h'} \supseteq W_i^h = \{\theta_i \mid \tilde{y}(\theta_i, \underline{\theta}_{-i}^h) = i\} = \{\theta_i \mid \tilde{y}(\theta_i, \underline{\theta}_{-i}^{h'}) = i\}$, where the first equality follows from the inductive hypothesis for κ and (G, S_N) orderly, and the second equality follows from $\Theta_{-i}^h = \Theta_{-i}^{h'}$. If h' is an immediate successor that follows play of any action other than a^* , then $W_i^{h'} = \emptyset$.

Case 4: Consider $i = P(h)$, such that $W_i^h \cap L_i^h = \emptyset$, $W_i^h \neq \emptyset$. By the inductive hypothesis for κ , $\Theta_i^h \supseteq \{\theta_i \mid \tilde{y}(\theta_i, \underline{\theta}_{-i}^h) = i\}$. Let $\theta'_i = \min\{\theta_i \mid \tilde{y}(\theta_i, \underline{\theta}_{-i}^h) = i\}$. $\theta'_i \in W_i^h$, so $\theta'_i \notin L_i^h$. Hence, for all $\theta_{-i} \in \Theta_{-i}^h$, $\min\{\theta_i \mid \tilde{y}(\theta_i, \theta_{-i}) = i\} \leq \theta'_i$. By (G, S_N) orderly, for all $\theta_{-i} \in \Theta_{-i}^h$, $\min\{\theta_i \mid \tilde{y}(\theta_i, \theta_{-i}) = i\} \geq \theta'_i$. Thus, by threshold pricing (Proposition 8), i faces a posted price $\tau_h = \theta'_i$ at h . (G, S_N) is concise, so i is not called to play at any successor of h . By the definition of W_i^h , L_i^h , and the induced allocation rule $\tilde{y}(\cdot)$, it follows that, for any $h' \in \text{succ}(h)$, either $L_i^{h'} = \emptyset$ or $W_i^{h'} = \emptyset$. If $L_i^{h'} = \emptyset$, then conditional on reaching h' , i wins at price τ_h , and by (G, S_N) pruned, h' is terminal and the inductive hypothesis holds trivially. If $W_i^{h'} = \emptyset$, then $L_i^h \setminus W_i^h = L_i^h \supseteq L_i^{h'} = L_i^{h'} \setminus W_i^{h'}$, so the inductive hypothesis for κ implies that part 2 of the inductive hypothesis holds for i at h' .

Our analysis of Case 4 also proves the final claim in Proposition 12. Q.E.D.

We now have the pieces in place to complete the characterization. Assume F_N is symmetric and regular, and (G, S_N) is optimal, orderly, concise, credible, and strategy-proof. Clause 1 of Definition 14 holds by assumption. Clause 2 holds by Proposition 8. We will now show that we can label offered prices, accepting actions, and quitting actions, so that the rest of the definition holds.

Take any non-terminal history h , and let $i = P(h)$, let $h \in I_i$. There are three cases to consider:

Case 1: $W_i^h = \emptyset$. In this case, $\Theta_i^h = L_i^h$, so by Proposition 12, for all $\theta_i, \theta'_i \in \Theta_i^h$, and all $\theta_{-i} \in \Theta_{-i}^h$, $\tilde{y}(\theta_i, \theta_{-i}) = \tilde{y}(\theta'_i, \theta_{-i})$. By threshold pricing, for all $j \in N$, all $\theta_i, \theta'_i \in \Theta_i^h$, and all $\theta_{-i} \in \Theta_{-i}^h$, $\tilde{t}_j(\theta_i, \theta_{-i}) = \tilde{t}_j(\theta'_i, \theta_{-i})$. That is, conditional on reaching h , i 's type no longer affects the outcome. By (G, S_N) pruned, no histories fall into this case.

Case 2: $W_i^h \cap L_i^h \neq \emptyset$. Then, by Proposition 11, there exists action a^* available at h such that, for all $\theta_i \in W_i^h$, $S_i(I_i, \theta_i) = a^*$. This action is the unique **accepting** action. All other

actions **quit**. The **offered price** is $\min\{\theta_i \in \Theta_i^h \mid S_i(I_i, \theta_i) = a^*\}$. By defining the offered price in this way, we have ensured (by Proposition 9) that the offered price is no less than i 's current bid. Observe that for any history h' that follows play of $a \neq a^*$ at h , $W_i^{h'} = \emptyset$. Thus, if i plays $a \neq a^*$, then he does not win, pays zero, and (by our analysis of Case 1) is not called to play again. Proposition 12 implies that $\theta_i^K \in W_i^h$, so $S_i(I_i, \theta_i^K) = a^*$. Thus, since i 's strategy is measurable with respect to his types and information sets, i knows which action accepts, which actions reject, and the price offered.

Case 3: $W_i^h \neq \emptyset$, $W_i^h \cap L_i^h = \emptyset$. By Proposition 12, i faces a posted price τ_h at h . The **offered price** is τ_h . By (G, S_N) concise, I_i is a singleton set, and i is not called to play at any successor of h . Since i faces a posted price at most once along the path-of-play, if he was called to play at any predecessor h' of h , then that h' was in Case 2. Thus, by Proposition 12 and threshold pricing, the offered price τ_h is at least his current bid. By the definition of W_i^h , L_i^h , and the induced allocation rule $\tilde{y}(\cdot)$, it follows that for any action a available at h , either $\{\theta_i \mid S_i(I_i, \theta_i) = a\} \subseteq W_i^h$ or $\{\theta_i \mid S_i(I_i, \theta_i) = a\} \subseteq L_i^h$. Actions of the first kind **accept**, and actions of the second kind **reject**. I_i is singleton, so i knows which actions accept, which actions reject, and the offered price. $W_i^h \cap L_i^h = \emptyset$, so if i plays any action that accepts, then he wins the object and pays the offered price.

It remains to prove that if i wins the object at some terminal history z , then he pays his current bid. Let h be the latest predecessor of z at which i was called to play. If h is in Case 3, then the conclusion follows trivially. Suppose h is in Case 2. By Proposition 12, $\Theta_i^h \supseteq \{\theta_i \mid \tilde{y}(\theta_i, \underline{\theta}_{-i}^h) = i\}$. By construction, all types in $\{\theta_i \mid \tilde{y}(\theta_i, \underline{\theta}_{-i}^h) = i\}$ play the accepting action a^* at h . Thus,

$$\min\{\theta_i \in \Theta_i^h \mid S_i(I_i, \theta_i) = a^*\} \leq \min\{\theta_i \mid \tilde{y}(\theta_i, \underline{\theta}_{-i}^h) = i\} \leq \min\{\theta_i \mid \tilde{y}(\theta_i, \underline{\theta}_{-i}^z) = i\}. \tag{18}$$

i wins at z , so z must follow play of a^* , and i 's current bid at z is $\min\{\theta_i \in \Theta_i^h \mid S_i(I_i, \theta_i) = a^*\}$. Moreover, all types who played a^* at h win upon reaching z , so

$$\min\{\theta_i \in \Theta_i^h \mid S_i(I_i, \theta_i) = a^*\} \geq \min\{\theta_i \mid \tilde{y}(\theta_i, \underline{\theta}_{-i}^z) = i\}. \tag{19}$$

Thus, $\min\{\theta_i \in \Theta_i^h \mid S_i(I_i, \theta_i) = a^*\} = \min\{\theta_i \mid \tilde{y}(\theta_i, \underline{\theta}_{-i}^z) = i\}$. By threshold pricing, at z , i pays $\min\{\theta_i \mid \tilde{y}(\theta_i, \underline{\theta}_{-i}^z) = i\}$, so i pays his current bid.

This completes the proof of Theorem 4.

B.6. Theorem 6

B.6.1. Virtual Ascending \rightarrow Credible, Strategy-Proof

Suppose F_N is regular, and (G, S_N) is optimal and a virtual ascending auction. (G, S_N) is strategy-proof. The proof requires only small modifications to the proof of Theorem 4; at each point we relied on $\tilde{y}(\cdot)$ being orderly, we instead rely on $\tilde{y}(\cdot)$ maximizing virtual value. We omit the details to avoid repetition.

Again by a parallel argument to Theorem 4, S_i is a best response to any (S'_0, S_{-i}) for $S'_0 \in \mathcal{S}_0^*(S_0^G, S_N)$. Thus, if (G, S_N) is not credible, then there exists (G', S_N) is BIC and has voluntary participation, but yields strictly higher expected revenue for the auctioneer, which implies that (G, S_N) is not optimal. Thus, if (G, S_N) is optimal and a virtual ascending auction, then (G, S_N) is credible.

B.6.2. Credible, Strategy-Proof \rightarrow Virtual Ascending

Propositions 7, 8, 9, and 10 pin down some details even when F_N is not symmetric. We start by proving an analogue to Proposition 11.

PROPOSITION 13: Assume F_N is regular and interleaved, and (G, S_N) is optimal and strategy-proof. If (G, S_N) is credible, then (G, S_N) is winner-pooling.

PROOF: As before, we will show that if (G, S_N) is not winner-pooling, then the auctioneer has a profitable safe deviation, so (G, S_N) is not credible. Let h^* be some history at which the winner-pooling property does not hold; we pick h^* such that, for all $h < h^*$, h is not a counterexample to winner-pooling. Since (G, S_N) is regular and interleaved, and the winner-pooling property held at all predecessors to h^* , Proposition 4 implies that for all i , either $W_i^{h^*} = \emptyset$ or $W_i^{h^*} = \{\theta_i \mid \eta_i(\theta_i) > \max(0, \max_{j \neq i} \eta_j(\underline{\theta}_j^{h^*}))\}$. Let us define i^* , $\theta_{i^*}^*$ and h^{**} as before.

The proof of Proposition 11 works here with the following modifications: First, we define

$$\psi(i) = \arg \max_{j \in N \setminus \{i\}} \{\eta_j(\theta_j^{K_j}) \mid W_j^{h^*} \neq \emptyset\}. \quad (20)$$

Second, we say $\theta_{\psi(i)}$ *i-separates* at $\gamma \in \mathbb{R}$ if

$$\{\theta_i \mid \eta_i(\theta_i) \geq \gamma\} = \{\theta_i \mid \eta_i(\theta_i) \geq \eta_j(\theta_{\psi(i)})\}. \quad (21)$$

Third, we initialize $\beta := \min\{\eta_{i^*}(\theta_{i^*}^*), \eta_{\psi(i^*)}(\theta_{\psi(i^*)}^{K_{\psi(i^*)}})\}$ and specify the algorithm as:

Stage 1

1. Pick $\theta_{\psi(i^*)}$ that i^* -separates at β .
2. Simulate $(\theta_{\psi(i^*)}, \underline{\theta}_{N \setminus \{i^*, \psi(i^*)\}}^{h^*})$ against i^* starting from \hat{h}_{i^*} , until either $\eta_{i^*}(\underline{\theta}_{i^*}^{h^*}) \geq \beta$ or $\hat{h}_{i^*} \in Z$.
3. If $\eta_{i^*}(\underline{\theta}_{i^*}^{h^*}) \geq \beta$, then set $\beta := \underline{\theta}_{i^*}^{h^*}$ and go to Stage 2.
4. Else, set $\hat{N} := \hat{N} \setminus \{i^*\}$,

$$\beta := \min_{i \neq i^*, \theta_i} \eta_i(\theta_i) \mid \theta_i \in W_i^{h^*} \quad (22)$$

and go to Stage 2.

Stage 2

1. If $\hat{N} = 1$, go to Stage 3.
2. Set $\hat{i} := \{i \in \hat{N} \mid \eta_i(\underline{\theta}_i^{h^*}) < \beta\}$.
3. Pick $\theta_{\psi(\hat{i})}$ that \hat{i} -separates at β .
4. If $(\theta_{\psi(\hat{i})}, \underline{\theta}_{N \setminus \{\hat{i}, \psi(\hat{i})\}}^{h^*}) \notin \Theta_{-\hat{i}}^{\hat{h}_i}$, set $\hat{h}_i := \text{cousin}(\hat{h}_i, (\theta_{\psi(\hat{i})}, \underline{\theta}_{N \setminus \{\hat{i}, \psi(\hat{i})\}}^{h^*}))$.
5. Simulate $(\theta_{\psi(\hat{i})}, \underline{\theta}_{N \setminus \{\hat{i}, \psi(\hat{i})\}}^{h^*})$ against \hat{i} starting from \hat{h}_i , until either $\eta_i(\underline{\theta}_i^{h^*}) \geq \beta$ or $\hat{h}_i \in Z$.
6. If $\eta_i(\underline{\theta}_i^{h^*}) \geq \beta$, set $\beta := \eta_i(\underline{\theta}_i^{h^*})$ and go to Step 1 of Stage 2.
7. Else, set $\hat{N} := \hat{N} \setminus \{\hat{i}\}$ and go to Step 1 of Stage 2.

Stage 3

1. Set $\hat{i} := i \mid i \in \hat{N}$.
2. Pick $\theta_{\psi(\hat{i})}$ that \hat{i} -separates at β .
3. If $(\theta_{\psi(\hat{i})}, \underline{\theta}_{N \setminus \{\hat{i}, \psi(\hat{i})\}}^{h^*}) \notin \Theta_{-\hat{i}}^{\hat{h}_i}$, set $\hat{h}_i := \text{cousin}(\hat{h}_i, (\theta_{\psi(\hat{i})}, \underline{\theta}_{N \setminus \{\hat{i}, \psi(\hat{i})\}}^{h^*}))$.
4. Simulate $(\theta_{\psi(\hat{i})}, \underline{\theta}_{N \setminus \{\hat{i}, \psi(\hat{i})\}}^{h^*})$ against \hat{i} starting from \hat{h}_i , until $\hat{h}_i \in Z$.

5. Choose the outcome that corresponds to that terminal history, $x = g(\hat{h}_i)$, and terminate.

This deviating algorithm does not change the allocation; the object is kept if $\max_i \eta_i(\theta_i) \leq 0$, and allocated to $\arg \max_i \eta_i(\theta_i)$ otherwise (where $\arg \max_i \eta_i(\theta_i)$ is singleton since F_N is interleaved). Revenue is at least as high as under S_0^G , and strictly higher when $\theta_N = (\theta_{i^*}^*, \underline{\theta}_{-i^*}^{h^*})$.

It remains to check that the various steps of the algorithm are well-defined. We can pick separating types in Step 1 of Stage 1, because either $\beta = \eta_{\psi(i^*)}(\theta_{\psi(i^*)}^{K_{\psi(i^*)}})$ or $\beta = \eta_{i^*}(\theta_{i^*}^*) < \eta_{\psi(i^*)}(\theta_{\psi(i^*)}^{K_{\psi(i^*)}})$. In the first case, $\theta_{\psi(i^*)}^{K_{\psi(i^*)}}$ will i^* -separate at β . In the second case, since $\eta_{i^*}(\theta_{i^*}^*) > \eta_{\psi(i^*)}(\theta_{\psi(i^*)}^1)$, by F_N interleaved there exists $\theta_{\psi(i^*)}$ that will i^* -separate at β .

When we pick separating types in Step 3 of Stage 2 and Step 2 of Stage 3, β is equal to $\eta_j(\theta_j)$ for some bidder j where $\theta_j \in W_j^{h^*}$. Consider $\theta'_i = \min\{\theta_i \mid \eta_i(\theta_i) \geq \beta\}$. Since $\theta_j \in W_j^{h^*}$, it follows (by F_N regular and interleaved) that $\eta_i(\theta'_i) > \eta_{\psi(\hat{i})}(\theta_{\psi(\hat{i})}^1)$. If $\eta_i(\theta'_i) < \eta_{\psi(\hat{i})}(\theta_{\psi(\hat{i})}^{K_{\psi(\hat{i})}})$, then, by F_N interleaved, there exists $\theta_{\psi(\hat{i})}$ that will \hat{i} -separate at β . If $\eta_i(\theta'_i) \geq \eta_{\psi(\hat{i})}(\theta_{\psi(\hat{i})}^{K_{\psi(\hat{i})}})$, then since β never exceeds $\min\{\eta_i(\theta_i) \mid \eta_i(\theta_i) \geq \eta_{\psi(\hat{i})}(\theta_{\psi(\hat{i})}^{K_{\psi(\hat{i})}})\}$, it follows that $\theta_{\psi(\hat{i})}^{K_{\psi(\hat{i})}}$ will \hat{i} -separate at β .

We can choose cousins (in Step 4 of Stage 2 and Step 3 of Stage 3) because F_N is regular and (G, S_N) is strategy-proof and optimal, by the same argument as in the proof of Theorem 4 that invokes Proposition 10. Thus, the algorithm is well-defined, and produces a profitable safe deviation, which completes the proof. Q.E.D.

We now state a claim that is analogous to Proposition 12.

PROPOSITION 14: *Assume F_N is interleaved and regular, and (G, S_N) is optimal, credible, and strategy-proof. For any non-terminal history h and any bidder i :*

1. *If $W_i^h \neq \emptyset$, then $\Theta_i^h \supseteq \{\theta_i \mid \tilde{y}(\theta_i, \underline{\theta}_{-i}^h) = i\}$.*
2. *For all $\theta_i, \theta'_i \in L_i^h \setminus W_i^h$, for all $\theta_{-i} \in \Theta_{-i}^h$, $\tilde{y}(\theta_i, \theta_{-i}) = \tilde{y}(\theta'_i, \theta_{-i})$.*

Moreover, if $W_{P(h)}^h \cap L_{P(h)}^h = \emptyset$ and $W_{P(h)}^h \neq \emptyset$, then $P(h)$ faces a posted price at h .

PROOF: The proof follows in parallel to the proof of Proposition 12. The only modifications are that we invoke Proposition 13 instead of Proposition 11, and that we rely on $\tilde{y}(\cdot)$ maximizing virtual value for a regular interleaved distribution, instead of $\tilde{y}(\cdot)$ being orderly. Q.E.D.

Having established Propositions 13 and 14, we then follow the construction at the end of the proof of Theorem 4 to label offered prices, accepting actions, and quitting actions, so that the rest of the definition holds. This completes the proof of Theorem 6.

B.7. Theorem 7

By inspection, first-price auctions are prior-free credible and static.

Suppose (G, S_N) is prior-free credible and static. Suppose there exist $\theta_i, \theta_{-i}, \theta'_{-i}$ such that i wins the object at (θ_i, θ_{-i}) and at (θ_i, θ'_{-i}) , but $t_i(\theta_i, \theta_{-i}) < t_i(\theta_i, \theta'_{-i})$. We now construct a deviation: If the action profile is consistent with (θ_i, θ_{-i}) , award the object to i and instead charge $t_i(\theta_i, \theta'_{-i})$. This deviation is always-profitable.

Consequently, there exists a function $\tilde{b}_i : \Theta_i \rightarrow \mathbb{R}$ such that if type θ_i wins, then i pays $\tilde{b}_i(\theta_i)$. Notably, this property holds everywhere, and not just almost everywhere.

We now partition i 's actions into bidding actions $B_i = \{\tilde{b}_i(\theta_i) \mid \theta_i \in \Theta_i \text{ and } \exists \theta_{-i} : \tilde{y}(\theta_i, \theta_{-i}) = i\}$, and actions that decline. The same steps as in the proof of Theorem 2 establish that (G, S_N) is a first-price auction.

B.8. Theorem 8

With finite type spaces, credible protocols are prior-free credible, so the first claim follows trivially.

In the proof of Theorem 4, we show that if (G, S_N) is optimal, orderly, concise, and strategy-proof but not an ascending auction, then there exists a safe deviation that is always-profitable, so (G, S_N) is not prior-free credible. Thus, the second claim follows.

REFERENCES

- AKBARPOUR, M., AND S. LI (2020): "Supplement to 'Credible Auctions: A Trilemma'," *Econometrica Supplemental Material*, 88, <https://doi.org/10.3982/ECTA15925>. [445]
- ASHENFELTER, O. (1989): "How Auctions Work for Wine and Art," *Journal of Economic Perspectives*, 3, 23–36. [426]
- ATHEY, S., J. LEVIN, AND E. SEIRA (2011): "Comparing Open and Sealed Bid Auctions: Evidence From Timber Auctions," *The Quarterly Journal of Economics*, 126, 207–257. [427]
- BALIGA, S., L. C. CORCHON, AND T. SJÖSTRÖM (1997): "The Theory of Implementation When the Planner Is a Player," *Journal of Economic Theory*, 77, 15–33. [429]
- BERGEMANN, D., AND S. MORRIS (2005): "Robust Mechanism Design," *Econometrica*, 73, 1771–1813. [427]
- BERGEMANN, D., B. BROOKS, AND S. MORRIS (2017): "First-Price Auctions With General Information Structures: Implications for Bidding and Revenue," *Econometrica*, 85, 107–143. [430]
- (2019): "Revenue Guarantee Equivalence," *American Economic Review*, 109 (5), 1911–1929. [430]
- BERNHEIM, B. D., AND M. D. WHINSTON (1986): "Menu Auctions, Resource Allocation, and Economic Influence," *The Quarterly Journal of Economics*, 101, 1–31. [449]
- BESTER, H., AND R. STRAUZ (2000): "Imperfect Commitment and the Revelation Principle: The Multi-Agent Case," *Economics Letters*, 69, 165–171. [429]
- (2001): "Contracting With Imperfect Commitment and the Revelation Principle: The Single Agent Case," *Econometrica*, 69, 1077–1098. [429]
- BROOKS, B., AND S. DU (2018): "Optimal Auction Design With Common Values: An Informationally-Robust Approach," Working Paper. [436]
- BULOW, J., M. HUANG, AND P. KLEMPERER (1999): "Toeholds and Takeovers," *Journal of Political Economy*, 107, 427–454. [430]
- CARROLL, G. (2015): "Robustness and Linear Contracts," *American Economic Review*, 105, 536–563. [448]
- CASSADY, R. (1967): *Auctions and Auctioneering*. Univ. of California Press. [426,427]
- CHAKRABORTY, I., AND G. KOSMOPOULOU (2004): "Auctions With Shill Bidding," *Economic Theory*, 24, 271–287. [430]
- CHUNG, K.-S., AND J. C. ELY (2007): "Foundations of Dominant-Strategy Mechanisms," *The Review of Economic Studies*, 74, 447–476. [427]
- CREMER, J., AND R. P. MCLEAN (1988): "Full Extraction of the Surplus in Bayesian and Dominant Strategy Auctions," *Econometrica*, 1247–1257. [445]
- DEQUIEDT, V., AND D. MARTIMORT (2007): "Mechanism Design With Private Communication," Working Paper. [439]
- (2015): "Vertical Contracting With Informational Opportunism," *American Economic Review*, 105, 2141–2182. [429,439,445]
- DWORCZAK, P. (2020): "Mechanism Design With Aftermarkets: Cutoff Mechanisms," Working Paper. [426, 448]
- EDELMAN, B., M. OSTROVSKY, AND M. SCHWARZ (2007): "Internet Advertising and the Generalized Second-Price Auction: Selling Billions of Dollars Worth of Keywords," *American Economic Review*, 242–259. [427]

- ELKIND, E. (2007): "Designing and Learning Optimal Finite Support Auctions," in *Proceedings of the Eighteenth Annual ACM-SIAM Symposium on Discrete Algorithms*. Society for Industrial and Applied Mathematics, 736–745. [441,445]
- FANG, H., AND S. MORRIS (2006): "Multidimensional Private Value Auctions," *Journal of Economic Theory*, 126, 1–30. [430]
- GREEN, J., AND J.-J. LAFFONT (1977): "Characterization of Satisfactory Mechanisms for the Revelation of Preferences for Public Goods," *Econometrica*, 427–438. [427]
- HAFALIR, I., AND V. KRISHNA (2008): "Asymmetric Auctions With Resale," *American Economic Review*, 98, 87–112. [430]
- HAJLASZ, P., AND J. MALÝ (2002): "Approximation in Sobolev Spaces of Nonlinear Expressions Involving the Gradient," *Arkiv för Matematik*, 40, 245–274. [437]
- HOLMSTRÖM, B. (1979): "Groves' Scheme on Restricted Domains," *Econometrica*, 1137–1144. [427]
- HURWICZ, L. (1972): "On Informationally Decentralized Systems," in *Decision and Organization*, ed. by C. B. McGuire and R. Radner. Amsterdam: North-Holland, 297–336. [425,433]
- KRISHNA, V. (2010): *Auction Theory*. San Diego, USA: Academic Press. [427]
- LAMY, L. (2009): "The Shill Bidding Effect versus the Linkage Principle," *Journal of Economic Theory*, 144, 390–413. [430]
- LI, S. (2017): "Obviously Strategy-Proof Mechanisms," *American Economic Review*, 107, 3257–3287. [429,432, 454]
- LIU, Q., K. MIERENDORFF, X. SHI, AND W. ZHONG (2019): "Auctions With Limited Commitment," *American Economic Review*, 109 (3), 876–910. [429]
- LOERTSCHER, S., AND L. M. MARX (2017): "Optimal Clock Auctions," Working Paper. [430]
- LUCKING-REILEY, D. (2000): "Vickrey Auctions in Practice: From Nineteenth-Century Philately to Twenty-First-Century E-commerce," *The Journal of Economic Perspectives*, 14, 183–192. [427]
- MASKIN, E., AND J. RILEY (1984): "Optimal Auctions With Risk Averse Buyers," *Econometrica: Journal of the Econometric Society*, 1473–1518. [430,436,439]
- MCADAMS, D., AND M. SCHWARZ (2007a): "Credible Sales Mechanisms and Intermediaries," *American Economic Review*, 97, 260–276. [429,430]
- (2007b): "Who Pays When Auction Rules Are Bent?," *International Journal of Industrial Organization*, 25, 1144–1157. [430]
- MCAFEE, R. P., AND J. McMILLAN (1987): "Auctions and Bidding," *Journal of Economic Literature*, 25, 699–738. [425,427]
- MCAFEE, R. P., AND D. VINCENT (1997): "Sequentially Optimal Auctions," *Games and Economic Behavior*, 18, 246–276. [429]
- MEYER-NIEBERG, P. (1991): *Banach Lattices*. Springer-Verlag. [437]
- MILGROM, P. (1987): "Auction Theory," in *Advances in Economic Theory: Fifth World Congress*, ed. by T. Bewley. Cambridge: Cambridge University Press, 1–32. [429]
- MILGROM, P., AND I. SEGAL (2002): "Envelope Theorems for Arbitrary Choice Sets," *Econometrica*, 70, 583–601. [427]
- (2020): "Clock Auctions and Radio Spectrum Reallocation," *Journal of Political Economy*, 128 (1), 1–31. [443]
- MILGROM, P. R., AND R. J. WEBER (1982): "A Theory of Auctions and Competitive Bidding," *Econometrica: Journal of the Econometric Society*, 1089–1122. [430,436,440]
- MYERSON, R. B. (1981): "Optimal Auction Design," *Mathematics of Operations Research*, 6, 58–73. [427,435, 438]
- (1986): "Multistage Games With Communication," *Econometrica: Journal of the Econometric Society*, 323–358. [429]
- (2009): "Fundamental Theory of Institutions: A Lecture in Honor of Leo Hurwicz," *Review of Economic Design*, 13, 59. [448]
- PARKES, D. C., C. THORPE, AND W. LI (2015): "Achieving Trust Without Disclosure: Dark Pools and a Role for Secrecy-Preserving Verification," in *Third Conference on Auctions, Market Mechanisms and Their Applications*. [448]
- PORTER, R., AND Y. SHOHAM (2005): "On Cheating in Sealed-Bid Auctions," *Decision Support Systems*, 39, 41–54. [430]
- ROTHKOPF, M. H., AND R. M. HARSTAD (1995): "Two Models of Bid-Taker Cheating in Vickrey Auctions," *Journal of Business*, 257–267. [430]
- ROTHKOPF, M. H., T. J. TEISBERG, AND E. P. KAHN (1990): "Why Are Vickrey Auctions Rare?," *Journal of Political Economy*, 94–109. [425]

- ROUGHGARDEN, T., AND I. TALGAM-COHEN (2013): "Optimal and Near-Optimal Mechanism Design With Interdependent Values," in *Proceedings of the Fourteenth ACM Conference on Electronic Commerce*. ACM, 767–784. [446]
- SIMON, L. K., AND M. B. STINCHCOMBE (1989): "Extensive Form Games in Continuous Time: Pure Strategies," *Econometrica*, 1171–1214. [440]
- SKRETA, V. (2006): "Sequentially Optimal Mechanisms," *The Review of Economic Studies*, 73, 1085–1111. [429]
- (2015): "Optimal Auction Design Under Non-Commitment," *Journal of Economic Theory*, 159, 854–890. [429]
- VARTAINEN, H. (2013): "Auction Design Without Commitment," *Journal of the European Economic Association*, 11, 316–342. [429,430]
- VICKREY, W. (1961): "Counterspeculation, Auctions, and Competitive Sealed Tenders," *The Journal of Finance*, 16, 8–37. [425,427,447]
- VON NEUMANN, J., AND O. MORGENSTERN (1953): *Theory of Games and Economic Behavior* (Third Ed.). Princeton University Press. [429]
- WILSON, R. (1987): "Game Theoretic Analysis of Trading Processes," in *Adv. Econ. Theory*, ed. by T. Bewley. Cambridge University Press. [427]

Co-editor Dirk Bergemann handled this manuscript.

Manuscript received 9 December, 2017; final version accepted 11 November, 2019; available online 4 December, 2019.