Shill-Proof Auctions*

Andrew Komo† Scott Duke Kominers‡ Tim Roughgarden§

November 10, 2025

Abstract

We characterize single-item auction formats that are *shill-proof* in the sense that a profit-maximizing seller has no incentive to submit shill bids. We distinguish between *strong* shill-proofness, in which a seller with full knowledge of bidders' valuations can never profit from shilling, and *weak* shill-proofness, which requires only that the expected equilibrium profit from shilling is non-positive. The Dutch auction (with a suitable reserve) is the unique (revenue-)optimal and strongly shill-proof auction. Any deterministic auction can satisfy only two properties in the set {static, strategy-proof, weakly shill-proof}. Our main results extend to settings with affiliated and interdependent values.

*The authors thank Mohammad Akbarpour, Eric Budish, Jeremy Bulow, Peter Cramton, Jacob Leshno, Shengwu Li, Joshua Mollner, Stephen Morris, Marek Pycia, Alex Rees-Jones, Roberto Saitto, Clayton Thomas, and seminar audiences at a16z crypto, EC '25, MIT, SITE, and Stony Brook for helpful comments. Komo gratefully acknowledges support from the NSF Graduate Research Fellowship Program. Kominers gratefully acknowledges support from the Ng Fund and the Mathematics in Economics Research Fund of the Harvard Center of Mathematical Sciences and Applications. Roughgarden's research at Columbia University is supported in part by NSF awards CCF-2006737 and CNS-2212745, and research awards from the Briger Family Digital Finance Lab and the Center for Digital Finance and Technologies. Part of this work was conducted during the Simons Laufer Mathematical Sciences Institute Fall 2023 program on the Mathematics and Computer Science of Market and Mechanism Design, which was supported by NSF Grant No. DMS-1928930 and by the Alfred P. Sloan Foundation under grant G-2021-16778. Kominers and Roughgarden are members of the Research Team at a16z crypto (for general a16z disclosures, see https://www.a16z.com/disclosures/). Komo was an intern on the a16z crypto Research Team during part of this research. Notwithstanding, the ideas and opinions expressed herein are those of the authors, rather than of a16z or its affiliates. Kominers and Roughgarden also advise companies on marketplace and incentive design, including auction design. Any errors or omissions remain the sole responsibility of the authors.

†MIT Economics, Email: komo@mit.edu

‡Harvard Business School, Harvard University, and a16z crypto, Email: kominers@fas.harvard.edu

§a16z crypto and Columbia University, Email: tim.roughgarden@gmail.com

1 Introduction

1.1 Shill Bidding in Auctions

Shill Bidding in Practice. Auction theory typically assumes that an auction is carried out as described (by the seller or a third party) and focuses solely on the bidders' incentives. Reality is often different. For example, while major auction houses like Christie's or Sotheby's may appear to be carrying out textbook English (ascending) auctions, a degree of skullduggery is often afoot. According to a *New York Times* article from 2000:

> Some tricks of the trade, like an auctioneer's drumming up excitement by acknowledging nonexistent bids only he hears and potential buyers who bid with nearly imperceptible secret signals, have been around for decades. Making up bids, for instance, is known as "bidding off the chandelier" from an era when the grand auction rooms were adorned with ornate lighting.[1]

The practice continues to this day: Christie's Conditions of Sale for their flagship New York location, in a section titled "Auctioneer's Discretion," states (among other things) that "The auctioneer can. . . move the bidding backwards or forwards in any way he or she may decide. . . ."[2]

Such chandelier bids or *shill* bids—bids submitted by the seller in order to manipulate the final selling price—appear to be particularly common in online auctions. For example, eBay has long gone out of its way to emphasize that shill bidding is forbidden and will be punished:

> We want to maintain a fair marketplace for all our users, and as such, shill bidding is prohibited on eBay. [. . .] eBay has a number of systems in place to detect and monitor bidding patterns and practices. If we identify any malicious behavior, we'll take steps to prevent it.[3]

According to many eBay users, however, shill bidding remains rampant. Here's a sample quote from the eBay discussion forums:

> The Sellers post a Buy Now price 3–4 times the actual cost of the item. Then they place the item on an auction at $0.01. This is to get as many views as possible. The shill comes in shortly after the auction starts and . . . is there to prevent the item from being sold below their profit margin.[4]

Chen et al. (2020) find that nearly 10% of all bidders on eBay Motors are shill bidders.

Shill-Proof Auctions. Much of auction theory to date encourages truthful bidding through careful auction design, while punting on challenges like seller deviations and collusion via appeal to unmodeled concepts such as the out-of-mechanism enforcement of rules.

[1] See "Genteel Auction Houses Turning Aggressive," *New York Times*, April 24, 2000.

[2] See https://www.christies.com/help/buying-guide-important-information/conditions-of-sale.

[3] See https://www.ebay.com/help/policies/selling-policies/selling-practices-policy/shill-bidding-policy?id=4353.

[4] See https://community.ebay.com/t5/Buying/My-experience-with-Shill-bidders/td-p/30402514.

Anecdotes about eBay and other online platforms suggest that such methods are only partially effective at deterring seller deviations. Thus, it makes sense to ask: To what extent can these deviations instead be disincentivized through an auction's design?

The goal of this paper is to understand which auction formats are "shill-proof" in the sense that a seller cannot profit through the submission of shill bids. We show that shill bidding can matter even in private-value auctions. The reader might wonder why shill bids can have an impact in the private-values case—assuming that the choice of reserve price does not affect participation (as it does in the eBay example), isn't a shill bid the same thing as a reserve price?

The answer depends on when the seller has an opportunity to shill and the information available to them at that time. For example, consider an English auction in which the seller also participates via shill bidding. Suppose the valuations of the (real) bidders are private and drawn i.i.d. from a regular distribution F and that the opening bid of the auction is set optimally (for revenue), to the monopoly price θ^{ρ^*} of F. As the auction proceeds, with the offered price p starting at θ^{ρ^*} and increasing from there (in increments of ϵ, say), the seller can shill bid at any time. Suppose that the only additional information known to the seller at a given round of the auction is that the remaining bidders are willing to pay at least p. Then, the seller asks himself: "now that I know how many bidders are willing to pay at least p, do I want to shill and reset the reserve price to $p+\epsilon$?" Under our assumption that F is regular, the answer is "no," and an expected revenue–maximizing seller will never shill.[5]

Now suppose that the seller has full knowledge of bidders' realized valuations. In this scenario, the seller will certainly, in some cases, want to shill in an English auction to push the price up to just below the highest of the bidders' valuations. Lest this informational assumption—that the seller knows the full valuation profile—seem impossibly demanding, consider the Dutch (descending) auction (with an arbitrary reserve price). Here, any shill bid by the seller terminates the auction immediately, leaving the seller holding the item and earning zero revenue. Therefore, even if the seller knows the full valuation profile, he would not want to shill bid.

We map out a theory of "shill-proof" auctions, focusing on the following basic questions:

- Which auction formats are "strongly shill-proof" in the sense of the Dutch auction, i.e., with shill bidding being unprofitable even with full knowledge of bidders' realized valuations?

- Which auction formats are "weakly shill-proof" in the sense of the English auction (with bidders' valuations drawn i.i.d. from a regular distribution and an optimally chosen reserve price), i.e., with shill bidding being unprofitable in expectation at equilibrium?

- To what extent are strong and weak shill-proofness compatible with other desirable properties such as optimality, efficiency, ex-post incentive compatiblility, and sealed-bid implementation? To what extent is shill-proofness dependent on the bidders' valuation structure?

[5] Auction theory experts will now immediately recognize that the English auction with an optimally chosen reserve price is *not* generally shill-proof in this sense when the valuation distribution is not regular or the values are not private. See, for example, Footnote 18 of Milgrom and Weber (1982) for more discussion of the latter point.

1.2 Overview of Results

Iterative auction formats like Dutch and English auctions play a central role in our theory, and accordingly we study (real and shill) bidding in the extensive-form game that is induced by a choice of auction format, relying on a framework for extensive-form auction analysis developed by Li (2017) and Akbarpour and Li (2020). We consider single-item auctions with N bidders. A subset of these are shill bidders, which we model as bidders with zero private value for the item and with utility equal to the seller's revenue.[6] For most of the paper, we assume that non-shill bidders have private valuations drawn i.i.d. from a known distribution. In Section 5, we generalize to any affiliated type distribution and any interdependent value function satisfying some curvature assumptions (Assumption 5.1). We also assume that shill bidders observe all actions. An auction is then *weakly shill-proof* (Definition 2.1) if there exists an equilibrium of the induced extensive-form game in which the shill bidders never shill (i.e., always bid their true private value of 0). An auction is *strongly shill-proof* (Definition 2.2) if, moreover, shill bidders' equilibrium strategies are ex-post strategies. In our first result, we focus on *public* auctions (Definition 3.4), meaning auctions in which every bidder's action is publicly observable. This is arguably the most natural model for the analysis of typical iterative auctions such as Dutch and English auctions. We then turn to arbitrary information structures to prove our other results.

Next, we summarize the main results of this paper; see also Figure 1.

Strongly Shill-Proof Auctions. Our main result (Theorem 3.6) uniquely characterizes strongly shill-proof auctions: the Dutch auction (with consistent tie-breaking and monopoly reserve price) is strongly shill-proof and optimal (i.e., maximizes the seller's expected revenue), and is the *only* such auction in the public setting. In particular, strongly shill-proof optimal auctions cannot avoid using a large number of rounds, and they cannot be ex-post incentive compatible (for real bidders). The rough intuition for the proof of this result is that: (i) No matter the information structure, strongly shill-proof auctions must be pay-as-bid (Lemma 3.1); (ii) for any auction format other than a Dutch auction, there exists a history in which some bidder i can indicate that her value is higher than 0 without the auction ending immediately; (iii) incentive compatibility in tandem with the public setting then implies that this information effectively induces the auction to revise its reserve price upward or increases perceived competitiveness for the item being sold, which reduces bid shading; and thus (iv) there exist valuations for the bidders such that, if bidder i is a shill bidder, shilling will increase the seller's revenue.[7]

Weakly Shill-Proof and Ex-Post Incentive Compatible Static Auctions. We next turn to investigating the richer design space of weakly shill-proof auctions. Our main result about weakly shill-proof auctions (Theorem 4.3), focuses on *single-action* auctions, meaning auction formats that induce extensive-form games in which each bidder moves exactly once. In our framework, a single-action auction is *static* when each bidder moves

[6] The prior literature has sometimes modeled shill bidding via an unknown number of bidders, with some subset of the bidders who end up participating in the auction being shills. Our framework is essentially equivalent: we can take N to be large and require 0 to be in the support of the valuation distribution; and a bidder with value 0 is equivalent (in terms of outcomes) to a bidder not arriving.

[7] We focus on single-item auctions, but our uniqueness results a fortiori provide an upper bound on what is possible for multiunit auctions, as well. We leave formal study of multi-unit auctions and other settings, such as sequential auctions, to future work.

simultaneously. We prove that no weakly shill-proof, single-action auction can satisfy even a very weak ex-post incentive compatiblility condition (Definition 4.2). Thus, an auction can satisfy two and only two of the properties in the set {single-action, ex-post incentive compatible, weakly shill-proof}.[8]

Weakly Shill-Proof and Efficient Auctions. We then investigate *efficient* (and weakly shill-proof) public auctions. We prove that a Dutch auction with a reserve price of 0 is the unique prior-independent auction (in the sense of Dhangwatnotai et al. (2015), with no dependence whatsoever on the valuation distribution) that is both efficient and weakly shill-proof (Corollary 4.8). In the regular, IPV case, we show (in Proposition 4.6) that fixing the monopoly price, a prior-independent, weakly shill-proof and efficient auction must conclude with a Dutch auction when all bidders' valuations are known to be below said price. A format such as beginning with an English auction at the monopoly price and then, should there be no takers, concluding with a Dutch auction is an example of an auction that is weakly shill-proof and efficient given the monopoly price.[9]

Weakly Shill-Proof and Ex-Post Incentive Compatible Optimal Auctions. The previous two results imply that ex-post incentive compatible auctions cannot be both strongly shill-proof and optimal, nor can they be (robustly) weakly shill-proof and efficient. The English auction (with an optimal reserve price) is, as we've noted, weakly shill-proof, optimal, and ex-post incentive compatible in the regular, IPV case. Is it the unique such auction? Does this combination of properties require a potentially large number of rounds? Our next result (Proposition 4.12) shows that, in general, the answer is no: In fact, although we can never find a weakly shill-proof, optimal, and ex-post incentive compatible auction that finishes in one round, we can always find a weakly shill-proof auction and a valuation distribution such that the worst-case number of auction rounds required is an arbitrarily small fraction of the number of rounds needed in the English auction.

Affiliated and Interdependent Values. Our last theorem, Theorem 5.4, shows that as the type distribution becomes less affiliated (Definition 5.2) and the values become less commonly valued (Definition 5.3), then both the set of strongly shill-proof auctions and the set of weakly shill-proof auctions expand. We then apply Theorem 5.4 to show most of our characterization results (Theorem 3.6, Theorem 4.3, and Corollary 4.8) hold in a more general, affiliated environment. Theorem 5.4 means that shill-proofness admits a partial order with respect to both the value function and the affiliation structure. Our theorem holds for any affiliation structure, most commonly used interdependent value functions (the function must satisfy Assumption 5.1), and any extensive-form auction where an optimal transfer rule is used. The last restriction is necessary because we consider discrete types and so there are a multitude of incentive compatible transfers that could yield different revenue. We generalize Lemma 3.1 in Proposition 5.7 to show that even with affiliation, a strongly shill-proof and optimal auction must be pay-as-bid with a reserve structure on the allocation.

[8] Assuming a regular valuation distribution and a corresponding optimal reserve price, a second-price auction is single-action and ex-post incentive compatible; a first-price auction is single-action and weakly shill-proof; and an English auction is ex-post incentive compatible and weakly shill-proof.

[9] In fact, this auction format closely resembles the Honolulu–Sydney fish auction documented by Hafalir et al. (2023).

	Static	Not Static
Strategy-Proof	Impossible (Theorem 4.3)	Ascending, Screening Auction (Proposition 4.12)
Not Strategy-Proof	First-Price Auction	Dutch Auction (Theorem 3.6)

(a) Weakly shill-proof and optimal auctions

	Static	Not Static
Strategy-Proof	Impossible (Theorem 4.3)	Not Robustly (Proposition 4.6)
Not Strategy-Proof	Not Robustly (Proposition 4.6)	Dutch Auction (Robustly Unique, Proposition 4.6)

(b) Weakly shill-proof and efficient auctions

	Static	Not Static
Strategy-Proof	Impossible (Theorem 3.6)	Impossible (Theorem 3.6)
Not Strategy-Proof	Impossible (Theorem 3.6)	Dutch Auction (Unique, Theorem 3.6)

(c) Strongly shill-proof and optimal auctions

	Static	Not Static
Strategy-Proof	Impossible (Proposition 4.6)	Impossible (Proposition 4.6)
Not Strategy-Proof	Impossible (Proposition 4.6)	Dutch Auction (Unique, Proposition 4.6)

(d) Strongly shill-proof and efficient auctions

Figure 1: **Summary of Results.** Characterization of single-item auction formats that are strongly or weakly shill-proof, along with other properties such as optimality, efficiency, ex-post incentive compatiblility, and sealed-bid implementations.

1.3 Related Work

While the idea and practice of shill bidding by a seller have long been well known, the auction theory literature on the topic is surprisingly thin. Chakraborty and Kosmopoulou (2004) consider common value auctions and focus on technological barriers (as opposed to auction formats) that can mitigate shill bidding. Lamy (2009) studies shill bidding specifically in English auctions in which bidders' valuations are affiliated in the sense of Milgrom and Weber (1982), and proves that shill bidding effectively cancels out the effects of affiliation in equilibrium due to real bidders conditioning on bids being fake (see also Izmalkov (2004)). We also consider shill bidding in affiliated environments, but consider a much larger class of possible extensive-form games.

Porter and Shoham (2005) consider a model similar to a second-price auction, motivated by "cheating" by online platforms that can announce a manipulated auction outcome subsequent to collecting all of the bidders' bids. More recently, a number of works (e.g., Roughgarden (2021); Lavi et al. (2022); Basu et al. (2023); Chung and Shi (2023)) have considered shill bidding in the context of blockchain transaction fee mechanism design, with an emphasis on knapsack auctions that are ex-post incentive compatible, shill-proof, and robust to various forms of collusion. Ausubel and Milgrom (2006) and Day and Milgrom (2008) consider shill bids by *bidders* in a multi-item auction, who are looking to exploit complementarities to lower their payments in VCG-type mechanisms—as opposed to shill bids by a seller looking to increase revenue, as is the case of this paper.[10] Contemporaneous work by Shinozaki (2024) and Zeng (2024) also study shill bidding, but take a different approach to the problem. While we study ex-interim deterrence against shill bidding in extensive form games, they group auctions into equivalence classes based on the outcome and primarily focus on ex-ante deterrence against the seller inserting additional bidders into the auction to increase perceived competition. The results in those papers are complementary to ours: While we show that the dynamics of the auction are important for preventing shill bidding, they show that when the shill bidders can increase perceived competition outside

[10]More broadly there is a literature on sybil-resistance referred to as "false-name proofness" (see, e.g., Conitzer et al. (2010) for a reference).

of taking actions in the auction, only the posted-price mechanism is non-manipulable (see also Footnote 13).

Our theory of shill-proof auctions is similar in spirit to the theory of credible mechanisms developed by Akbarpour and Li (2020), and leverages their framework for extensive-form auction analysis. That said, shill-proofness differs conceptually from credibility as shill-proofness focuses on the auctioneer's incentives to insert fake bids, whereas credibility focuses on the auctioneer's incentive to truthfully report the actions of a bidder to other bidders. Furthermore, the results in this paper are also qualitatively different. For example, there are a multitude of credible auctions, but only one strongly shill-proof auction and there are a multitude of strategy-proof and weakly shill-proof auctions, but only one strategy-proof and credible auction (see Section 6.1 for more discussion). In Proposition 6.1, we prove that strong shill-proofness is a stronger condition than credibility and weak shill-proofness is a weaker condition in the sense that if an auction is strongly shill-proof, then it is credible, which in turn implies that it is weakly shill-proof.

More recent research on credible mechanisms, usually with a focus on evading the impossibility results of Akbarpour and Li (2020) under extra assumptions (such as adding cryptographic tools), includes the work of Essaidi et al. (2022), Ferreira and Weinberg (2020), and Chitra et al. (2023). More distantly related papers include that of Haupt and Hitzig (2021), who prove a uniqueness result for the Dutch auction under contextual privacy constraints.

1.4 Outline of the Paper

In Section 2, we present our formal model of shill bidding in auctions. Section 3 studies strong shill-proofness and shows the ways in which Dutch auctions are uniquely optimal at disincentivizing shill bidding. Section 4 explores which formats are weakly shill-proof and provides a trilemma result. Section 5 generalizes the model to a setting with affiliation and interdependent values, proves that shill-proofness admits, and ordering with respect to affiliation and interdependent values, and generalizes our results. In Section 6, we conclude the paper by discussing extensions.

2 Model

In this paper, we consider extensive-form, single item auctions. An extensive-form game G is a tuple of possible histories H, and, for each history $h \in H$, functions mapping h to: (i) a player taking an action, $P(h)$; (ii) a set of possible actions, $A(h)$; (iii) an information set, $\mathcal{I}(h)$; and (iv) the most recent action taken, $\mathcal{A}(h)$. As further notation, we denote the starting history of the game by $h_\varnothing$ and the set of terminal histories as Z; we say $h' \prec h$ if h' precedes h, i.e., there exists a sequence of actions that lead from h' to h.

We restrict attention to single-item auctions, which means that for every terminal history $z \in Z$, we can associate an allocation and transfer vector: $z = (x, t)$, with $\sum_{i=1}^{N} x_i \leqslant 1$ and $x_i \in \{0, 1\}$ for all i. Abusing notation slightly, we use $x(z), t(z)$ to mean the vectors (x, t) associated with the terminal history z. We also assume perfect recall and finite depth. (Definition E.1 in the Supplemental Appendix gives a formal, thorough, and standard definition of extensive-form games.)

2.1 Bidders – Real and Shill

In the auction, there is a set of potential bidders B, with $|B| = N$, who might participate. We assume that the seller's value is commonly known to be 0. Of these potential bidders, a set of real bidders R actually participate. Each bidder $i \in B$ has an independent probability p of participating, $\mathbb{P}[i \in R] = p$.[11] The other bidders, $S = B \backslash R$, are shill bidders whose incentives are completely aligned with the seller/auctioneer's, i.e., their utility is defined by the sum of real bidders' transfers: for $i \in S, u_i(z) = \sum_{j \in R} t_j(z)$. Each real bidder $i \in R$ has value θ_i for the item being sold where $\theta_i \sim F$ independently for each i. Each real bidder has quasi-linear utility: for $i \in R$, $u_i(z) = x_i(z)\theta_i - t_i(z)$. We assume F is discrete, with support ϑ consisting of the ordered atoms $0 = \theta^1 < \theta^2 < \ldots < \theta^M$, and we define $f(\theta^k) = \mathbb{P}_{w \sim F}\left[w = \theta^k\right]$ to be the probability mass function (pmf) of the distribution. As notation, for each shill bidder $i \in S$, we assign $\theta_i = 0$ and let $\theta = (\theta_1, \ldots, \theta_N)$. The choice of values for shill bidders does not affect their incentives, and by supposing that their values are 0, we can define efficiency and optimality (revenue-maximization) in terms of only θ instead of θ and R.[12] Observe that given how θ is generated, we are in the standard, symmetric, single-item independent private values (IPV) setting.

Note that (i) whether a given bidder is real and (ii) the bidders' values for the item are not built into the extensive-form game G. Instead, bidders' strategies are a function of their types. Like most papers in the extensive-form auction literature, we study games with a finite type space because defining auctions with a continuum of types requires defining a general class of continuous-time games. To the authors' knowledge, there is no theory of continuous-time games that rivals the generality and flexibility of extensive-form games.

Real bidders have no ex-ante information about who else is a real bidder; they only know that each other bidder is real with probability p.[13] However, throughout the course of the auction, they can update their beliefs about which bidders are real and adjust their actions accordingly. We assume that shill bidders know the set of shill bidders and observe all previous actions taken. Formally, for any history h, if $P(h) \in S$, then $\mathcal{I}(h) = \{h\}$. This assumption rules out games with simultaneity (including static games) from the perspective of the shill bidders, but not real bidders.[14]

Our equilibrium concept is pure-strategy Perfect Bayesian Equilibrium; a formal definition of the auction equilibrium (G, σ) can be found in Definition A.1. We write $\sigma(\theta; R)$ for the strategy profile when the value profile is θ and the realized set of real bidders is R. Perfect Bayesian equilibria are defined without consideration of group deviations by the

[11] This randomness plays little role in our analysis—we impose it only so that the overarching structure of our model has bidders with ex-ante, symmetric, independent private values. See also Footnote 6.

[12] When considering optimal auctions, we naturally assume the seller only cares about raising revenue from real bidders rather than shill bidders, since shill bidders are proxies for the seller themself.

[13] Unlike the assumption that N is fixed, this assumption is an economically substantive one: The only way for shill bidders to manipulate the outcome of the auction is for shill bidders to take actions in the auction. If the bidders instead knew who were the real bidders, shill bidders could have an incentive to appear as real bidders in order to increase perceived competition.

[14] In extensive form games, simultaneity is modeled as the information set of a player having multiple elements. Without cryptography or other unmodeled technologies, we view it as reasonable to assume that while the auction may appear simultaneous to the real bidders, actions are taking place sequentially and the seller can observe those actions.

shill bidders. However, with this setup, it is without loss to focus only on individual shill bidder's incentives as opposed to any group actions because all the shill bidders have the same objective function and information available to them.

2.2 Auction Environment

Throughout the paper, we focus on auction equilibria that are ex-post monotone and individually rational: The auction equilibrium (G, σ) is **monotone** if, for all i, j, θ_{-j}, and $\theta_j > \theta'_j$, $\left[t_i\left(\sigma(\theta; B)\right) > 0 \implies t_i\left(\sigma(\theta; B)\right) \geqslant t_i\left(\sigma(\theta'_j, \theta_{-j}; B)\right)\right]$, and is **individually rational** (IR) if, for all v and $i \in B$, $x_i(\sigma(\theta; B))v_i - t_i(\sigma(\theta; B)) \geqslant 0$. The monotonicity condition on transfers is satisfied by all standard single-item auction formats such as English auctions, Vickrey auctions, Dutch auctions, and sealed first-price auctions. Moreover, note that we impose the monotonicity condition primarily for convenience; all our results continue to hold if we instead just assume that total transfers are weakly higher whenever any bidder reports a value higher than 0. The ex-post IR condition rules out all-pay auctions and ensures that only the winner pays the seller.[15] We also make this assumption primarily for convenience, versions of all our main results would still hold if we were to relax to ex-interim IR.

2.3 Shill-Proofness

Next, we define our key shill-proofness desiderata. We are interested in auction equilibria in which shill bidders do not shill. Formally, this corresponds to requiring that shill bidders always act like real bidders who have value 0 for the item—since real bidders who have value 0 will never enter non-trivial bids in equilibrium, requiring shill bidders to have the same actions in equilibrium in effect means that shilling does not occur.

Definition 2.1. An auction equilibrium (G, σ) is **weakly shill-proof** if σ is invariant to the realization of S, i.e., for all θ and $S, S' \subseteq \{i : \theta_i = 0\}$: $\sigma\left(\theta; B \backslash S\right) = \sigma\left(\theta; B \backslash S'\right)$.

Note that Definition 2.1 is a statement about an equilibrium of an auction—it is possible (although we have not found an example of this) that an auction may have both shill-proof equilibria and non-shill-proof equilibria.

We can also strengthen our no-shilling criterion from equilibrium to ex-post strategy:

Definition 2.2. An auction equilibrium (G, σ) is **strongly shill-proof** if it is weakly shill-proof and σ is an ex-post strategy profile for shill bidders, i.e., for all σ', S, and θ_{-S},

$$\sum_{j \in R} t_j(\sigma(0, \theta_{-S}; R)) \geqslant \sum_{j \in R} t_j\left(\sigma'_S, \sigma_{-S}(0, \theta_{-S}; R)\right).$$

Strong shill-proofness is obviously preferable to weak shill-proofness (all else equal), especially if there are concerns about a seller acquiring information about real bidders' valuations beyond what is encoded by the prior. As we will see, however, the design space of weak shill-proof auctions is meaningfully larger than that of strong shill-proof auctions.

[15] Unlike with continuous types, with discrete types there is positive probability of multiple bidders having the same highest value and so an optimal or efficient auction make randomize between bidders. For many of our results, we will assume Definition 3.2 to rule out this case.

2.4 Revelation Principle

In order to make progress in understanding shill-proof auction formats, the following revelation principle will be helpful: for every auction equilibrium (G, σ), there exists a direct auction that can be summarized by a direct allocation rule $\tilde{x} : \vartheta^N \to [0,1]^N$, a direct transfer rule $\tilde{t} : \vartheta^N \to \mathbb{R}^N$, a menu rule

$$\mu : \mathcal{P}\left(\vartheta^N\right) \times B \to \bigcup_{L=2}^{M} \left(\{\mathcal{T} \in \Pi(\vartheta) \mid |\mathcal{T}| = L\} \times B^L\right)$$

where $\Pi(\vartheta)$ is the set of all possible partitions over the type space, and a starting player $\xi_0 \in B$.[16] The first input to the menu rule μ is a set Θ of valuation profiles of the form $\Theta = \Theta_1 \times \Theta_2 \times \cdots \times \Theta_N$ with $\Theta_i \subseteq \vartheta$ for all i—intuitively, the valuation profiles that are, in equilibrium, consistent with a particular history. The second input is a player ξ who is to move next. The output of the rule is a collection $\left\{\left(W_\ell, \vec{\xi}_\ell\right)\right\}_{\ell \in \{1,\ldots,L\}}$, where the W_ℓ's are a partition of Θ_ξ (the equilibrium strategy σ determines the partition; player ξ will truthfully choose a subset based on her valuation) and $\vec{\xi}_\ell$ indicates the next player to move should player ξ choose W_ℓ. Under σ, the player ξ will always select the partition W_ℓ such that $\theta_\xi \in W_\ell$. If $\vec{\xi}_\ell = \varnothing$, then the game ends should choice ℓ be selected by the bidder ξ. For a typical iterative auction, one generally has $L = 2$ with the two sets corresponding to types above and below some value, respectively. Or, for a single-action auction, the W_ℓ's are generally singletons, with one per type in Θ_i.

We show that for any implementable outcome $(\tilde{x}, \tilde{t})$ of the auction, we can always find a menu rule that is "informative"—the set of possible outcomes differs across partition selections[17]—that also implements the same outcome. So, without loss of generality, we restrict menu rules in this way and then describe an auction equilibrium by $(\tilde{x}, \tilde{t}, \mu, \xi_0)$. (See Lemma A.4 in the Appendix for a more formal treatment.) We refer to $(\tilde{x}, \tilde{t}, \mu, \xi_0)$ as an *auction* when convenient. As is always the case with direct mechanisms, the auction encompasses both the game form and the equilibrium, i.e., by appealing to the revelation principle we have implicitly selected the equilibrium.

Finally, as notation for later sections, for a set Θ of valuation profiles, define $\overline{\Theta}_i = \max\{\theta_i : \theta_i \in \Theta_i\}$ to be the maximum possible value of bidder i; $\overline{\Theta}_{-i}$ to be the maximum possible value of bidders $j \neq i$; and $\overline{\Theta} = \max_i\left\{\overline{\Theta}_i\right\}$. We define $\underline{\Theta}_i, \underline{\Theta}_{-i}$, and $\underline{\Theta}$ as the corresponding values for minima instead of maxima.

3 Strongly Shill-Proof Auctions

3.1 Direct Mechanisms

In this subsection, we first show that all strongly shill-proof auctions must be pay-as-bid and then show that under an assumption that real bidders observe all past actions, we can precisely pin down the Dutch auction as the the only strongly shill-proof auction.

[16] This revelation principle is similar to those found in, for example, Ashlagi and Gonczarowski (2018); Mackenzie (2020); Mackenzie and Zhou (2022); Pycia and Troyan (2023).

[17] This notion of informativeness is very similar to the pruned condition from Akbarpour and Li (2020).

Lemma 3.1 (Pay-as-bid)**.** *If an auction $(\tilde{x}, \tilde{t}, \mu, \xi_0)$ is strongly shill-proof, then it must be a pay-as-bid auction. Formally, for all ξ, θ_ξ, and $\theta_{-\xi}, \theta'_{-\xi}$,*

$$\tilde{x}_\xi\left(\theta_\xi, \theta_{-\xi}\right) = \tilde{x}_\xi\left(\theta_\xi, \theta'_{-\xi}\right) \implies \tilde{t}_\xi\left(\theta_\xi, \theta_{-\xi}\right) = \tilde{t}_\xi\left(\theta_\xi, \theta'_{-\xi}\right).$$

The proof of Lemma 3.1 and all other results can be found in the Appendix. Observe that Lemma 3.1 holds (as does Theorem 3.6) even if we relax the assumption on shill bidders' information sets because strong shill-proofness means that shill bidders want to report 0 even if they know the precise valuations of other bidders ex-ante. To see why Lemma 3.1 is true, consider the case where $R = \{\xi\}$. Then, the shill bidders will report whichever values maximize $\tilde{t}_\xi$ and so $\tilde{t}_\xi$ must be constant across all outcomes with the same allocation.

Recall that with continuous types, two bidders have the same type with probability 0 and the optimal allocation rule is uniquely defined up to measure-zero events. However, since types are discrete, the probability of ties is non-zero and so we have to define a tie-breaking rule. We adopt the notion of orderliness introduced by Akbarpour and Li (2020): there exists a fixed priority order—independent of values—over which bidder wins an item if there is a tie.[18] We clarify their definition by explicitly linking the priority order to the ex-post allocation rule.

Definition 3.2. An auction equilibrium (G, σ) is **orderly** if there exists a total ordering $\rhd$ over (θ_i, i) with the following properties: For all i and j,

(i) $\theta_i > \theta_j \implies (\theta_i, i) \rhd (\theta_j, j)$;

(ii) if there exists m such that $(\theta^m, i) \rhd (\theta^m, j)$, then for all k, $(\theta^k, i) \rhd (\theta^k, j)$; and

(iii) for all θ, if $\tilde{x}_i(\theta) \in \{0, 1\}$ and $(\theta_i, i) \rhd (\theta_j, j)$, then $\tilde{x}_i(\theta) > \tilde{x}_j(\theta)$.

We can now define the orderly allocation rule for any auction with reserve type θ^ρ as

$$\tilde{x}_i^\rho(\theta) = \mathbb{1}\left\{\theta_i \geqslant \theta^\rho, (\theta_i, i) = \max_{\rhd}\left\{(\theta_j, j)\right\}_{j \in B}\right\}.$$

Then, the orderly, efficient allocation rule is $\tilde{x}^E = \tilde{x}^1$ and the orderly, optimal allocation rule is $\tilde{x}^* = \tilde{x}^{\rho^*}$ for some ρ^*.

We can now explicitly define the revenue-maximizing pay-as-bid bidding functions. Note that we need to select among multiple IC bidding functions because the discrete type space implies that the IC constraints need not bind with equality. Using Lemma B.2 from the Appendix, the bidding function for bidder i is given by

$$b_i^1(\theta^m) = \theta^m - \sum_{k:\theta^k < \theta^m} (\theta^{k+1} - \theta^k) \frac{\left(F(\theta^k)\right)^{i-1}\left(F(\theta^{k-1})\right)^{N-i-1}}{\left(F(\theta^m)\right)^{i-1}\left(F(\theta^{m-1})\right)^{N-i-1}}.$$

The transfer rule is then $\tilde{t}^E = b^1 \cdot \tilde{x}^E$ and $\tilde{t}^* = b^1 \cdot \tilde{x}^*$ for the efficient and optimal auctions, respectively. Observe that the asymmetry in the bidding and transfer functions arises purely from the discrete types and the tie-breaking rule, not from the definition of shill-proofness.

[18] For example, if ties are broken lexicographically, then the auction is orderly.

3.2 Indirect Mechanisms

Now that we have pinned down the direct mechanism, we turn to the indirect implementation. The information structure available to real bidders matters for which mechanisms are strongly shill-proof. If there is no information leakage between bidders, then the classic first-price auction is strongly shill-proof. However, if any bidder can learn any outcome-relevant information about another bidder's actions, then the first-price auction is not strongly shill-proof. In Section 3.3, we return to the primary setting studied in this work, public auctions.

To expand, the single-action, first-price auction where each bidder sequentially makes a bid will be strongly shill-proof when the game is static, i.e., when each bidder takes her single action with no information about what other bidders have done. However, we can show that the single-action, first-price auction is strongly shill-proof only in static games—if there is ever a chance that a bidder learns some outcome-informative information about what actions previous bidders have taken, then the auction is no longer strongly shill-proof. Formally, in an orderly auction, for some i, j, and θ with $(\theta^M, j) \rhd (\theta_i, i) \rhd (\theta^{\rho^*}, j)$, let h_θ be the (unique) history reached in equilibrium by $\sigma(\theta; B)$ where $P(h) = j$ and let h_0 be the corresponding history reached by types $(0, \theta_{-i})$. We say that an information structure is **strictly more informative** than a static game if there exist i, j, and θ such that $\mathcal{I}(h_\theta) \neq \mathcal{I}(h_0)$ and $\tilde{x}_j(\theta) = 1$. Note that if there exists such a type vector, then there must exist a type vector where $\tilde{x}_j(\theta) = 0$.

Proposition 3.3. *Under a static information structure, the single-action, first-price auction is strongly shill-proof. However, for any information structure strictly more informative than a static game, the single-action, first price auction is not strongly shill-proof.*

A core intuition underlying Proposition 3.3 is that the stronger a bidder perceives her competition, the less she shades her bid in a first-price auction. So, if a shill bidder i is able to distinguish that she has a value that could change the outcome of the auction to bidder j, then there must be some value θ_j under bidder j perceives the auction as more competitive; this is occurs if and only if there is outcome-relevant information leaked from bidder i to j. This will increase revenue if j wins and so the first-price auction is not strongly shill-proof.

As noted in Section 2, implementing no information leakage (or, equivalently, simultaneity) between bidders is often an unrealistic assumption. Even with some technology that might deter information leakage in the actual mechanism, information leakage through side channels is inevitable in many settings. So, we take Proposition 3.3 as evidence therefore that the first-price auction is, in many settings, potentially subject to manipulation by shill bidders.

3.3 Dutch Auctions

Now that we have shown that the extensive form of an auction affects whether or not it is strongly shill-proof, we turn to studying which extensive forms are always strongly shill-proof. We show that Dutch auctions are strongly shill-proof, no matter the information structure, and show that that there exists a natural information structure—the public setting—such that Dutch auctions are uniquely strongly shill-proof.

To begin, we formally define the public information structure:

Definition 3.4. An auction equilibrium (G, σ) is **public** if the information set at any history is all previous actions taken. Formally, for any history h, $\mathcal{I}(h) = \{h\}$.

Public auctions are common in practice: from open air fish markets, to auctions on eBay, participants often can see every action other bidders take before choosing what to do.[19] A second interpretation of an auction being public is that all information is totally leakable: If an auction is strongly shill-proof and public, then shill bidders can credibly signal to all other bidders any actions they have taken in the auction and it will not affect the outcome. In general, we believe that it is natural to assume that shill bidders actions are visible in settings where considering shill bidding is important. Otherwise, shill bidding can only manipulate the outcome ex-post, which differs from the normal interpretation of shill bidding and its applications.

Next, we define the Dutch auction with reserve price θ^ρ. The Dutch auction is defined as the auction which begins by offering each bidder i the item at $b_i^1(\theta^M)$, and then if no bidder chooses to buy the item at that price, the item is offered for $b_i^1(\theta^{M-1})$ and so on until either a bidder has chosen to buy the item or the price to be offered drops below $b_i^1(\theta^\rho)$. Note that the optimal Dutch auction is the Dutch auction with reserve price θ^{ρ^*}. We consider only orderly auctions and therefore, at each price level, bidders are offered the opportunity to buy the item in priority order. Formally:

Definition 3.5. The **Dutch auction with reserve price** θ^ρ is defined by the allocation rule $\tilde{x}^\rho$, first-price transfer rule $\tilde{t}^1 = \tilde{x}^\rho \cdot b^1$, initial player $(\cdot, \xi_0) = \max_\triangleright \{(0, i)\}$, and menu

$$\mu_\rho^D(\Theta, \xi) = \left\{ \left(W_{\mathsf{L}}, \vec{\xi}_{\mathsf{L}} \right), \left(W_{\mathsf{H}}, \vec{\xi}_{\mathsf{H}} \right) \right\},$$

$$\text{where } W_{\mathsf{H}} = \left\{ \overline{\Theta}_\xi \right\}, W_{\mathsf{L}} = \Theta_\xi \backslash \left\{ \overline{\Theta}_\xi \right\}, \vec{\xi}_{\mathsf{H}} = \varnothing, \text{ and}$$

$$\vec{\xi}_{\mathsf{L}} = \begin{cases} \left(\cdot, \tilde{\xi} \right) = \max_\triangleright \left\{ \left(\overline{\Theta}_i, i \right) : i \neq \xi \right\} & \exists i \neq \xi \text{ such that } |\Theta_i| > 1 \text{ and } \overline{\Theta}_i \geqslant \theta^\rho \\ \varnothing & \text{otherwise} \end{cases}.$$

Theorem 3.6. *A Dutch auction with any reserve price is strong shill-proof. Furthermore, if a public, orderly and optimal auction is strongly shill-proof, then it is the Dutch auction with reserve price* θ^{ρ^*}.

Proof Sketch. The Dutch auction is strongly shill-proof because any shill bid immediately ends the auction—and in that case there would be no transfers from other bidders. To gain an intuition for why uniqueness holds, note that the key defining property of the Dutch auction is that any bid immediately ends the auction. Then, we prove the result in four cases. First, we consider an auction where the player rotation differs from a Dutch auction; we then show that if a shill bidder indicates she has the highest possible remaining value, this has the effect of ex-interim increasing the effective reserve price. Since the auction is public and only the highest value is allocated, when the types are realized so that some bidder with a higher priority order has the same value (and non-shill bidders have a low value),

[19] Non-examples of public auctions include the FCC spectrum auctions, where bidders typically only learn information on other bidders' actions in rounds (see Milgrom and Segal (2017) for more information).

raising the reserve must increase revenue from the winning bidder.[20] The next case considers what happens when bidders are queried in priority order, and the minimum of a partition is greater than the reserve price, but less than the highest value of other bidders. This raises the ex-interim reserve price and so by the same argument as above, shill bidding is profitable. In the third case, we consider what happens when the minimum in the partition is equal to the highest value of other bidders. By construction, this can only occur when a bidder is the lowest priority and so, applying our orderliness assumption, this raises the ex-interim reserve to be the highest type still possible in equilibrium for other bidders without immediately allocating the good to the shill bidder. We can then apply the same argument as the first two cases to conclude that shill bidding would be profitable in this case. In the final case, we consider an auction where the minimum in the partition is below the reserve price. In this case, the ex-interim reserve has not changed, so the logic from the previous cases does not apply. Instead, shill bidding in this case makes it appear as if there is more competition for the item, and we prove that this causes "less bid shading," i.e., higher final transfers from the winner. □

Observe that although the standard intuition about shill bidding is that its' purpose is to influence other bidders' perception about the common value for an item, Theorem 3.6 shows that just being able to influence other bidders perception about the probability they win the item without directly winning the item *in a single case* is enough to limit the number of possible auction formats to just Dutch auctions. Moreover, under public information, such a case always exists for non-Dutch auctions. The public information structure is the most informative, while a static information structure is the least informative. We show in the Supplemental Appendix that as the information structure becomes less informative, the number of strongly shill-proof auctions weakly increases (Proposition E.3).

4 Weakly Shill-Proof Auctions

In Lemma 3.1, we demonstrate that strong shill-proofness uniquely pins down the transfer rule; and in Theorem 3.6, we demonstrate that once we consider public auctions, we uniquely pin down the extensive form as well. In this section, we turn to studying our more permissive definition, weak shill-proofness, and first show the limits of possible extensive forms that are weakly shill-proof and then exhibit extensive-form auctions that are weakly shill-proof and strategy-proof, including a new auction format that can be arbitrarily faster than the English auction.

4.1 Trilemma

In this subsection, we prove that it is impossible to find an auction that is weakly shill-proof, strategically simple for even one of the real bidders, and *fast* in the sense that each bidder takes exactly one action. To make this trilemma as general as possible, we do not even

[20] If the auction were not public, then real bidders' beliefs might not change ex-interim and so their transfers might not change either. For example, in a sealed first-price auction, a shill bidder's actions have no effect on the transfers of other bidders.

require the auction to be optimal or efficient. Instead, we are only using condition (iii) of our orderliness definition (Definition 3.2), which requires that a bidder who does not have the highest type must not win the item. Note that such a condition is without loss for any symmetric and deterministic mechanism.

We begin by defining a **single-action auction**. An auction is considered single-action when each bidder takes precisely one action in the auction (under all possible histories). More formally, for any $h_N \in Z$, let $h_\varnothing \prec h_1 \prec \dots h_{N-1} \prec h_N$ be the sequence of preceding histories. Then, an auction is **single-action** if for all $i \in B$, there exists a unique $n \leqslant N$ such that $i = P(h_n)$. Without loss, we label the bidders $1, \dots, N$, in the order that they move and label the action taken by bidder i as a_i.[21]

For expositional purposes, instead of tracking the information set $\mathcal{I}_i$ of a bidder $i \in R$, we assume that the information i has when taking an action is a signal $s_i \in \mathcal{S}_i$. This signal is generated via a deterministic function $\psi_i : \left(\times_{j<i} A_j\right) \to \mathcal{S}_i$ called an **experiment**.[22] For notational convenience, we assume that ψ_i is surjective for all i. We can think of the experiment as a garbling of the previous bidders' actions—the experiment can pool together multiple actions from previous bidders to a single signal and so a signal is not always perfectly informative of previous actions. We can recover the public setting with a fully informative experiment, i.e., $\psi = \text{Id}$, the identity mapping. We can capture classical static game settings via an uninformative experiment that always returns the same output, $\psi = \varnothing$. We use $\psi_i^{-1}(s_i)$ to denote the set of θ_{-i} that are possible from the perspective of bidder i given its signal.

A revelation principle holds in this setting: for any single-action auction, we can define the direct allocation and transfer rules as $\tilde{x}(\theta)$ and $\tilde{t}(\theta)$, respectively, with the appropriate incentive compatibility and individual rationality constraints for real bidders (Lemma C.1 in the Appendix), and appropriate IC constraints for weak and strong shill-proofness (Lemmatta C.2 and E.13, respectively, in the Appendix).

To finish defining all the terms necessary for our main result of this section, we formally define ex-post incentive compatibility and then give a weaker notion of ex-post incentive compatiblility in the single-action auction setting: ex-post incentive compatiblility for at least a single bidder.

Definition 4.1. An auction $(\tilde{x}, \tilde{t}, \mu, \xi_0)$ is **ex-post incentive compatible** if it is an ex-post strategy for real bidders to report their values truthfully: for all real bidders $i \in R$, θ, and θ_i',

$$\tilde{x}(\theta) \cdot \theta_i - \tilde{t}(\theta) \geqslant \tilde{x}(\theta_i', \theta_{-i}) \cdot \theta_i - \tilde{t}(\theta_i', \theta_{-i}).$$

Definition 4.2. A single-action auction $(\tilde{x}, \tilde{t}, \mu, \xi_0)$ is **mildly ex-post incentive compatible** if there exists a real bidder $i < N$ such that truthfulness is an ex-post strategy for i conditional on the realization of her signal: there exists bidder $i < N$, such that for all θ_i, θ_i', s_i, and $\theta_{-i}, \theta_{-i}' \in \psi_i^{-1}(s_i)$: $\tilde{x}_i(\theta) \cdot \theta_i - \tilde{t}_i(\theta) \geqslant \tilde{x}_i(\theta') \cdot \theta_i - \tilde{t}_i(\theta')$.[23]

[21] The bidder ordering and therefore the labeling can be endogenous to actions taken.

[22] Abusing notation, we also sometimes take $\psi_i : \vartheta^{i-1} \to \mathcal{S}_i$, i.e., the experiment maps values to signals instead of actions.

[23] We exclude the last bidder who takes an action from our definition because a take-it-or-leave-it offer to that bidder can be optimal and ex-post incentive compatible. Theorem 4.3 would still hold if we instead defined mild ex-post incentive compatiblility to mean ex-post incentive compatible for at least two bidders.

Theorem 4.3. *There exists no orderly auction that is single-action, mildly ex-post incentive compatible, and weakly shill-proof.*

Proof Sketch. Consider any real bidder $i < N$. By weak shill-proofness, the transfer from bidder i, conditional on winning (or losing) the auction, is invariant to the values of bidders who take actions after her (Lemma C.3 in the Appendix). If this were not the case, then if every bidder $j > i$ is a shill bidder, the shill bidders would report the values that would maximize the transfer from the winning bidder. By mild ex-post incentive compatiblility, the transfer from bidder i, conditional on winning the auction, is invariant to her value (Lemma C.4 in the Appendix). This is because if there were multiple winning reports with different transfer amounts, only the smallest transfer amount would make truthful reporting of the value an ex-post strategy. So, in every single-action, optimal auction, the transfer from the winning bidder i, can depend only on the values reported by bidders before i. But, this means that if a bidder has positive utility for winning the item (as would be the case if $\theta_i > \theta_j$ for all $j < i$), then she should report θ^M to maximize the probability of winning (without changing the transfer paid upon winning). Thus, the auction must treat bidder i as if she reported θ^M, which violates the allocation rule of an optimal auction. □

Observe that while there are no auctions at the intersection of all three conditions in Theorem 4.3, we can find multiple auctions at the intersection of any two of those conditions. For example, assuming an IPV and regular (Definition 4.4) environment, the second-price auction is single-action and mildly ex-post incentive compatible; the first-price auction is single-action and weakly shill-proof; and the English auction is mildly ex-post incentive compatible and weakly shill-proof. None of the examples just described are unique; there exist other single-action and mildly ex-post incentive compatible auctions, single-action and weakly shill-proof auctions, and mildly ex-post incentive compatible and weakly shill-proof auctions.

4.2 Other Weakly Shill-Proof Auctions

Now that we have shown an impossibility result regarding weakly shill-proof auctions, in this subsection, we will provide some characterization results for weakly shill-proof auctions. Recall from the introduction that even in an IPV environment, an English auction need not be weakly shill-proof when the type distribution is irregular. So, for tractability, for the rest of this subsection, we suppose that bidders have IPV types and regular type distributions. We use the definition of regularity for discrete types found in Elkind (2007):

Definition 4.4. A distribution F is **regular** if for all k, the virtual value $\varphi^k = \theta^k - (\theta^{k+1} - \theta^k)\frac{1-F(\theta^k)}{f(\theta^k)}$ is non-decreasing.

4.2.1 Weakly Shill-Proof and Efficient Auctions

To further analyze when concerns about shill-bidding lead to Dutch auctions, we now examine the case of efficient auctions. Shill bidding is important to consider in the efficient auction context because the designer may be interested in allocating goods efficiently even while sellers are trying maximize revenue. For example, in two-sided marketplaces like eBay,

the auctioneer/market designer places positive welfare weight on both buyers and sellers, whereas sellers typically do not—and sellers may certainly try to shill bid in these settings, as discussed in the Introduction. Our next result shows that in order for an auction to be robustly *weakly* shill-proof and efficient, part of its game tree must be a Dutch auction. If an auction is not *robustly* weakly shill-proof, we mean that we can find a value distribution such that the auction is not weakly shill-proof. More formally, if the auction $(\tilde{x}, \tilde{t}, \mu, \xi_0)$ is parameterized by the optimal reserve θ^{ρ^*}, the number of atoms below the reserve $\underline{M}$ and the number of atoms weakly above the reserve $\overline{M}$,[24] then the following result holds.

Definition 4.5. An efficient auction $(\tilde{x}^E, \tilde{t}, \mu, \xi_0)$ is a **semi-Dutch auction with cutoff** θ^{ρ^*} if for any v such that $\max_i \{v_i\} < \theta^{\rho^*}$:

(i) $\breve{\Theta} = \left\{w : w < \theta^{\rho^*}\right\}^N$ is reached; and

(ii) $\mu(\Theta, \xi) = \mu^D(\Theta, \xi)$ for any player ξ and possible values $\Theta \subseteq \breve{\Theta}$ where μ^D is the menu rule for the Dutch auction with reserve price 0 from Definition 3.5.

Proposition 4.6. *For every public and efficient auction that is not a semi-Dutch auction with cutoff θ^{ρ^*}, there exists a regular value distribution with optimal reserve θ^{ρ^*} under which the auction is not weakly shill-proof.*

The key step in the proof of Proposition 4.6 resembles the proof of Theorem 3.6—in any non-Dutch auction, shill bidders can ex-interim "raise the reserve price" by changing their actions. However, given that we are interested in weak shill-proofness instead of strong shill-proofness, we have to examine shill bidders' incentives when we take expectations over real bidders' values instead of conditioning directly on their values. Regularity implies that above θ^{ρ^*}, shill bidders do not have an incentive to shill bid in auction formats such as the English auction (see Section 1.1). However, below θ^{ρ^*}, we can always find a regular distribution such that the ex-interim expected value of raising the reserve price is always positive; in particular, we can find a value distribution where the atoms are far enough apart that raising the reserve price a single "level" generates a large amount of additional revenue. In the Appendix, we construct the claimed sub-class of regular distributions (see Definition C.5). So, below θ^{ρ^*}, the auction must resemble the Dutch auction; the class of all such auctions is precisely all semi-Dutch auctions with cutoff θ^{ρ^*}. In the Supplemental Appendix (Example E.5), we provide an example of an auction format that is not semi-Dutch and that is weakly shill-proof for some value distributions but not others; and the following example presents a real-world setting that roughly fits the premises of Proposition 4.6 where a semi-Dutch auction (that is not a Dutch auction) is used.

Example 4.7. The Honolulu-Sydney fish auctions and Istanbul flower auctions documented by Hafalir et al. (2023) blend elements of the Dutch and English auctions: The auction begins at some intermediate price and if anyone bids, then the price ascends like in the English auction. If no one bids, the price descends until someone bids like in the Dutch auction.[25]

[24] Here, $\underline{M} + \overline{M} = M$.

[25] Honolulu-Sydney auction, once someone bids, other bidders can counter-bid and raise the price once more. However, in practice there is little counter-bidding. On the theoretical side, in an IPV setting, there exists an equilibrium where there is no counter-bidding. Counter-bidding once the Dutch auction starts is not allowed in the Istanbul flower auction.

The Honolulu-Sydney auction plausibly fits the technical assumptions made in Proposition 4.6: The auctions are public, as they take place in person and all bidders can see other bidders' actions. Market participants are interested in efficient outcomes because the goods are perishable and there are positive disposal costs for the sellers. We do not mean to imply that the Honolulu-Sydney auction was instituted precisely because it is shill-proof, but we highlight it as further evidence that in markets where it is difficult to monitor shill bidding, shill-proof mechanisms may arise.

Note that Proposition 4.6 relies on the auction format being able to condition on the true, optimal reserve. If we instead require the auction format to be completely prior-independent, then the only public, efficient, and weakly shill-proof auction is the Dutch auction.

Corollary 4.8. *For any public and efficient auction that is not a Dutch auction, there exists a regular value distribution under which the auction is not weakly shill-proof.*

Proof. Observe that a semi-Dutch auction with cutoff $\theta^{\rho^*} = \theta^M$ is simply a Dutch auction. Then, apply Proposition 4.6 for a regular distribution with optimal reserve $\theta^{\rho^*} = \theta^M$. □

4.2.2 Weakly Shill-Proof and Strategy-Proof Auctions

We have shown that the only optimal auction with a feasible equilibrium for real bids and an ex-post strategy for shill bidders not to shill (strong shill-proofness) is the Dutch auction. We now investigate the dual question: what optimal auctions have an ex-post strategy for real bidders (ex-post incentive compatiblility) and an equilibrium under which no shill bidding occurs (weak shill-proofness)?

An optimal auction is ex-post incentive compatible if and only if has the second-price transfer rule (Akbarpour and Li, 2020, Proposition 8):

$$\tilde{t}_i^2(\theta) = \tilde{x}_i^*(\theta) \cdot \max\left\{\theta^{\rho^*}, \text{second-highest value in } \{\theta_1, \ldots, \theta_N\}\right\}.$$

Note that if a shill bidder knew the valuations of all other bidders, then shill bidding would turn a second-price auction into a first price auction, which bounds the expected profit for a shill bidder from shilling. So, in order to find a ex-post incentive compatible and weakly shill-proof auction, we must find a menu rule that implements a second-price auction where the expected gain from shill bidding is sufficiently small at shill bidders' information sets. As discussed in Section 1.1, for regular value distributions, one weakly shill-proof, ex-post incentive compatible, and optimal auction is the English auction. We formalize the English auction in our framework as follows:

Definition 4.9. The **English auction with reserve price** θ^{ρ^*} is defined as the auction with the optimal allocation rule $\tilde{x}^*$, second-price transfer rule $\tilde{t}^2$, initial player $(\cdot, \xi_0) = \min_{\triangleright}\{(0,i)\}_{i\in B}$, and menu

$$\mu^E(\Theta,\xi) = \left\{\left(W_{\mathsf{L}}, \vec{\xi}_{\mathsf{L}}\right), \left(W_{\mathsf{H}}, \vec{\xi}_{\mathsf{H}}\right)\right\},$$

$$\text{where } W_{\mathsf{L}} = \left\{\theta \in \Theta_\xi : \theta < \theta^{\rho^*}\right\} \cup \left\{\underline{\Theta}_\xi\right\}, W_{\mathsf{H}} = \Theta_\xi \backslash W_{\mathsf{L}},$$

$$\vec{\xi}_{\mathsf{L}} = \vec{\xi}_{\mathsf{H}} = \begin{cases} \left(\cdot, \tilde{\xi}\right) = \min_{\triangleright}\left\{(\underline{\Theta}_i, i) : i \neq \xi, \overline{\Theta}_i = \theta^M, |\Theta_i| > 1\right\} & \overline{\Theta}_{-\xi} = \theta^M \\ \varnothing & \text{otherwise} \end{cases}.$$

Remark 4.10. The English auction with reserve price θ^{ρ^*} is weakly shill-proof, ex-post incentive compatible, and optimal (when the value distribution is regular). Depending on the information bidders have when taking actions, that auction can also be made dominant-strategy.[26]

The English auction is not the only ex-post incentive compatible and weakly shill-proof auction. While the English auction is used frequently, one drawback is that it is "slow"—each bidder can be queried on their willingness-to-pay on the order of M times. Specifically, let $Q^E(F) = M - \rho^* + 1$ be the worst-case number of times a bidder must be queried. To explore whether there are weakly shill-proof and ex-post incentive compatible auctions that require fewer rounds of communication, we introduce a natural "compression" of the English auction that comprises the following two phases:

1. An English auction is run from θ^{ρ^*} to some θ^Y.

2. If necessary, a second-price auction is then run among players who have not dropped out before the value level of θ^Y.

Definition 4.11. The **ascending, screening auction** with screen level θ^Y is defined by the optimal allocation rule $\tilde{x}^*$, second-price transfer rule $\tilde{t}^2$, initial player $(\cdot, \xi_0) = \min_{\triangleright} \left\{ \left(\overline{\Theta}_i, i \right) \right\}$, and menu

$$\mu(\Theta, \xi) = \begin{cases} \mu^E(\Theta, \xi) & \exists i \text{ such that } \underline{\Theta}_i \leqslant \theta^Y \text{ and } \overline{\Theta}_i = \theta^M \\ \left\{ \left(\left\{ \theta^k \right\}, \vec{\xi}^k \right) \right\}_{k \in \{Y+1, \ldots, M\}} & \text{otherwise} \end{cases},$$

$$\text{where } \left(\cdot, \vec{\xi}^k \right) = \max_{\triangleright} \left\{ (0, i) : |\Theta_i| > 1, \underline{\Theta}_i = \theta^Y \right\} \text{ for } k < M \text{ and } \vec{\xi}^M = \varnothing.$$

The ascending, screening auction reduces the maximum number of times each bidder can be queried to $Q^{AS,Y}(F) = Y - \rho^* + 2$. Because the transfer rule is $\tilde{t}^2$, the ascending, screening auction is ex-post incentive compatible and optimal.

We use the ascending, screening auction format to explore how fast a weakly shill-proof, ex-post incentive compatible and optimal auction can be, as a function of the underlying value distribution. Our next result shows that the ascending, screening auction can be weakly shill-proof, ex-post incentive compatible, optimal, and take an arbitrarily small fraction of queries as compared to the English auction depending on the value distribution:

Proposition 4.12. *For all $\varepsilon > 0$, there exists a value distribution F and screen level θ^Y such that $Q^{AS,Y}(F)/Q^E(F) < \varepsilon$ and the ascending, screening auction with screening level θ^Y is weakly shill-proof, ex-post incentive compatible, and optimal.*

The ascending, screening auction is orderly and optimal by construction; and it is ex-post incentive compatible because the English auction phase and the second-price phase both induce the same (ex-post incentive compatible) allocation and transfer rule. The larger θ^Y is, the less that can be extracted in expectation from shill bidding and the more likely it is that a shill bidder will win the item if she shill bids. We provide a sufficient minimum

[26] See Section 6.2 for a discussion of how our results translate into dominant-strategy contexts.

bound on θ^Y based on a few moments of a distribution (not its number of atoms), such that for distributions with "thin-enough" right tails—in particular monotone hazard rate distributions—the ascending, screening auction is weakly shill-proof (Lemma C.10). We can then construct a sequence of distributions with increasing numbers of atoms and constant θ^Y to complete the proof. The distributional requirements we need here differ from the class we find to prove Proposition 4.6, demonstrating that designing weakly shill-proof auctions has rich interactions with properties of the prior distribution.

5 Affiliation, Interdependent Values, and Shill-Proofness

In this section, we generalize our model to allow for interdependent values and a more general distribution of shill bidders. We will then show that shill-proofness admits an order with respect to the type distribution and value function: the higher the level of affiliation and level of common values that is present in the environment, the harder it is for an auction format to be shill-proof, either weakly or strongly. Next, we show that Lemma 3.1 holds under our general model and pins down the optimal allocation rule for strongly shill-proof auctions in affiliated environments. Finally, we show that our main results (Theorems 3.6 and 4.3) still hold under our more general model.

5.1 Generalized Model

We maintain the assumptions that there is a set of potential bidders B, with $|B| = N$, who might participate and that the seller's value is commonly known to be 0. We also maintain our previous assumptions on real and shill bidder utilities and information sets. However, instead of assuming that each potential bidder has probability p of participating, we instead assume that the set of real bidders is drawn symmetrically and randomly, i.e., for any $R', R'' \subset B$, if $|R'| = |R''|$, then $\mathbb{P}[R = R'] = \mathbb{P}[R = R'']$. For parsimony, we also assume full support, but the probability of certain sets of real bidders can be arbitrarily small; this rules out mechanisms that condition on large coalition of bidders in unusual ways in order to become weakly shill-proof.[27] Letting the probability of certain realizations go to zero (or relaxing the full support assumption entirely), this treatment of the set of shill bidders nests many natural models of shill bidders. For example, it nests a model of shill bidding where there is at most or exactly one shill bidder. It also nests our model of shill bidding from Section 2 where each bidder has an i.i.d. chance of being a real bidder.

In our general model, we interpret the type $\theta_i \in \vartheta$ to flexibly represent both a private value for the item and a signal of the common value for the item. We will expand on exactly how the type enters the valuation for real bidders below. As notation, for each shill bidder $i \in S$, we assign their type to be $\theta_i = 0$ to represent them having the lowest possible valuation for the item. Then, we can say types are jointly drawn from a full-support, symmetric F and we define $f(\theta) = \mathbb{P}_{w \sim F}[w = \theta]$ to be the pmf of the distribution. We assume that the type distribution is **affiliated** in the Milgrom and Weber (1982) sense: For all $\theta, \theta' \in \vartheta^N$,

$$\log f(\theta \vee \theta') + \log f(\theta \wedge \theta') \geqslant \log f(\theta) + \log f(\theta').$$

[27] Our results for strong shill-proofness do not require a full support distribution over shill bidder sets.

Affiliation captures the idea that types are positively correlated: the probability of types/signals all being high or low is higher than the probability of some signals being high and some being low.

There is a value function $v : \vartheta \times \vartheta^{N-1} \to \mathbb{R}$ such that for each bidder i, her value for the item is $v\left(\theta_i, \{\theta_j\}_{j \neq i}\right)$; note that this form implies that bidders have symmetric preferences over the types of other bidders. We also make a few functional form restrictions on v below.

Assumption 5.1. The value function v for a bidder satisfies the following properties:

(i) For all k, $v(\theta^k, 0) = \theta^k$;

(ii) v is weakly increasing in θ_{-i};

(iii) If $\theta_i \geqslant \theta_j$, then $v\left(\theta_i, \left\{\theta_j, \theta_{-(i,j)}\right\}\right) \geqslant v\left(\theta_j, \left\{\theta_i, \theta_{-(i,j)}\right\}\right)$; and

(iv) For all θ such that $\theta_i \geqslant \max_{j \neq i} \{\theta_j\}$, v is weakly super-modular and has weakly decreasing differences.

Condition (i) is a (without loss) normalization conditional on the other parts of the assumption. Condition (ii) means that θ encodes common value preferences—a bidder values an item more when other bidders value it more. Condition (iii) means that bidders value their own high-type realizations more than they value high-type realizations for other bidders. (This is the standard single-crossing assumption needed for a responsive, incentive compatible mechanism to exist.) Condition (iv) means that conditional on having the highest ex-ante signal of value, a bidder has higher value for other bidders' signals whenever her type is high and has diminishing marginal returns to high signals.

Note that the standard, private values setting where for all θ, $v(\theta_i, \theta_{-i}) = \theta_i$ satisfies Assumption 5.1. Most other value functions in the literature also satisfy these conditions; for example, generalized, additive, interdependent values $v(\theta) = \theta_i + \kappa \sum_{j \neq i} \theta_j$ with $\kappa \leqslant 1$ and maximum common values $v(\theta) = \max_i \{\theta_i\}$.

5.2 Shill-Proof Order

To prove that increasing the level of affiliation and interdependence of values makes it more difficult for an auction to be shill-proof, let us begin by formally defining what it means for the value distribution to be "more affiliated" and for the value function to be "more commonly valued."

Definition 5.2. Consider two distributions $F, F' \in \Delta(\vartheta^N)$ such that all marginal distributions are the same: $F_i = F'_i$, for all i. A distribution F' is **more affiliated** than F, $F' \geq_{\text{Aff}} F$, if for all $x, y \in \vartheta^N$, $\log f'(x \vee y) - \log f'(x) \geqslant \log f(y) - \log f(x \wedge y)$.

Our definition of the affiliation order is standard and from Karlin and Rinott (1980). When $F' \geq_{\text{Aff}} F$, values between bidders are more highly correlated under F' than F. We fix the marginal distributions to emphasize that our definition focuses on the correlation between bidders and not the relative strength of the distributions. All affiliated distributions are more affiliated than the type distribution where bidders are independent: If F is affiliated, then $F \geq_{\text{Aff}} \prod_i F_i$.

Definition 5.3. A value function v' is **more commonly valued** than v, $v' \succeq_{\text{Com}} v$, if v' is weakly more super-modular than v and for all θ, $v'(\theta) \geqslant v(\theta)$.

The point-wise comparison in Definition 5.3 makes sense because v, v' are normalized so that $v(\theta^k, 0) = v'(\theta^k, 0) = \theta^k$—and therefore the point-wise comparison means that other bidders' signals are valued more as the value function becomes more commonly valued. The increase in super-modularity means that we also require types to be more complementary when we say that the value function is more commonly valued. We view this as reasonable because if signals became less complementary, there would exist $v' \succeq_{\text{Com}} v$ such the marginal change in bidder i's value based on her own type would be less under v' than under v. Any value function v satisfying Assumption 5.1 is more commonly valued than the value function $v_{\text{private}}(\theta_i, \theta_{-i}) = \theta_i$ arising under the private values model.

Letting $\tilde{t}^*(\tilde{x}, \mu, \xi_0, v, F)$ be defined as some optimal transfer rule conditional on the extensive-form game and the primitives of the environment, we are now in position to formally state our main result for this section:

Theorem 5.4. *Consider affiliated type distributions F and F', and value functions v and v' satisfying Assumption 5.1. Suppose $(\tilde{x}, \tilde{t}^*(\tilde{x}, \mu, \xi_0, v, F), \mu, \xi_0)$ is orderly and strongly shill-proof. Then, if $v' \succeq_{\text{Com}} v$ and $F \succeq_{\text{Aff}} F'$, it is the case that $(\tilde{x}, \tilde{t}^*(\tilde{x}, \mu, \xi_0, v', F'), \mu, \xi_0)$ is strongly shill-proof. The same statement holds for weak shill-proofness.*

Proof Sketch. To begin the proof sketch, let us first observe that the ex-interim transfers from each bidder is pinned down based on the ex-interim allocation rule via a bidder's incentive compatibility constraint.

Lemma 5.5. *For every auction $(\tilde{x}, \tilde{t}^*(\tilde{x}, \mu, \xi_0, v, F), \mu, \xi_0)$, the ex-interim transfer rule for bidder i is*

$$T_i(\theta_i; \Theta) = X_i(\theta_i; \Theta) v^*(\theta_i, \theta_i) - \sum_{m : \theta^{j_m} < \theta_i} X_i(\theta^{j_m}; \Theta) \cdot (v^*(\theta^{j_m}, \theta^{j_{m+1}}; \Theta) - v^*(\theta^{j_m}, \theta^{j_m}; \Theta)),$$

$$\text{where } v_i^*(\theta^k, \theta^{k'}; \Theta) = \mathbb{E}\left[v(\theta^k, \theta_{-i}) \mid \theta_i = \theta^{k'}, \theta_{-i} \in \Theta_{-i}, \tilde{x}_i(\theta^k, \theta_{-i}) = 1 \right]$$

is the expected value of the item for bidder i of type $\theta^{k'}$ conditional on winning the item and playing as if she were type θ^k and $\{\theta^{j_m}\}_m$ are the ordered atoms of Θ_i.

Note that feasibility normally pins down the transfer rule via the envelope theorem. However, we are working with discrete types instead of continuous types and so the incentive compatibility constraint is slackened as the change in allocation probability from misreporting is discrete. So, we use optimality to select among this multitude of ex-interim transfers that maintain incentive compatibility. (This is an adaptation of the standard envelope theorem formulation of ex-interim transfers adapted to the extensive-form game by noting that if, at any point in the game, T_i does not take this form, there exists a profitable deviation by bidder i to commit to acting as a lower type throughout the rest of the game.)

Next, we prove the following technical lemma.

Lemma 5.6. *Consider sets S and S' such that $S \cap S' = \varnothing$ and $\min_{v \in S} \{v\} > \min_{v \in S'} \{v\}$. Then, for any weakly super-modular, non-decreasing function g with decreasing differences, if $F' \succeq_{\text{Aff}} F$, then for all i,*

$$\mathbb{E}_{\theta \sim F'}\left[g(\theta) \mid \theta_i \in S \right] - \mathbb{E}_{\theta \sim F'}\left[g(\theta) \mid \theta_i \in S' \right] \geqslant \mathbb{E}_{\theta \sim F}\left[g(\theta) \mid \theta_i \in S \right] - \mathbb{E}_{\theta \sim F}\left[g(\theta) \mid \theta_i \in S' \right]. \quad (1)$$

The proof then proceeds by contrapositive for affiliation and then common values. We consider that the auction is not weakly shill-proof for distribution F and consider $F' \geq_{\text{Aff}} F$. We show that if there exists a profitable deviation under F, then that same deviation must be profitable under F'. Any profitable deviation either manipulates a bidder's ex-interim expected transfer or ex-post manipulates her payment. Given that we have fixed the extensive form, any deviation that is possible under (v, F) is possible under (v', F'). So, ex-post manipulation will be profitable under both and so we only need to consider ex-interim manipulations. We can re-write expected profit of deviation as the difference between the expectation of a sum taken over the true and manipulated type reports. We show that the sum satisfies the conditions of Lemma 5.6 and so we can conclude that deviating increases profits. Note that the probability that bidder i wins the item is constant between the two type distributions because the allocation rule is held constant and so we have shown that the auction is not weakly shill-proof under F'. The argument when strongly shill-proof is the exact same. The proof when considering common values is also by contrapositive. In this case, we apply the fact that v' is more super-modular and point-wise larger to directly conclude that the profit from the deviation is larger. □

5.3 Extending our Characterizations

We can now use Theorem 5.4 to generalize our main characterization results from previous sections. To begin, we observe that while we know the optimal allocation rule in an IPV environment and the revenue equivalence theorem implies that any transfer rule yields the same revenue, the same need not hold with interdependent values and affiliation. So, we must simultaneously search for an allocation and transfer rule that satisfies our desiderata.

In order to find the optimal allocation and transfer rules, we first pin down the transfer rule and then the allocation rule. Our proof of Lemma 3.1 holds in our general environment and so strong shill-proofness implies that the transfer rule must be pay-as-bid. Then, having pinned down the transfer rule, we can see first that the unconstrained efficient allocation is feasible, and also that in the optimal auction we will never want to allocate to a bidder who does not have the highest type. Thus, we can solve the monopolist screening problem to find the optimal reserve in this auction. The formal result is as follows.

Proposition 5.7. *Under any affiliated distribution and value function satisfying Assumption 5.1, if an auction $(\tilde{x}, \tilde{t}, \mu, \xi_0)$ is strongly shill-proof, then it is pay-as-bid: For all ξ, v_ξ, $\theta_{-\xi}$, and $\theta'_{-\xi}$,*

$$\tilde{x}_\xi(\theta_\xi, \theta_{-\xi}) = \tilde{x}_\xi\left(\theta_\xi, \theta'_{-\xi}\right) \implies \tilde{t}_\xi(\theta_\xi, \theta_{-\xi}) = \tilde{t}_\xi\left(\theta_\xi, \theta'_{-\xi}\right).$$

All efficient mechanisms have an allocation rule such that

$$\sum_i \tilde{x}_i^E(\theta) = 1 \text{ and } i \notin \operatorname*{argmax}_j \{\theta_j\} \implies \tilde{x}_i^E(\theta) = 0.$$

All optimal mechanisms have an allocation rule such that

$$\sum_i \tilde{x}_i^*(\theta) = \begin{cases} 1 & \max_j \{\theta_j\} \geqslant \theta^{\rho^*} \\ 0 & \text{otherwise} \end{cases} \quad \text{and } i \notin \operatorname*{argmax}_j \{\theta_j\} \implies \tilde{x}_i^E(\theta) = 0$$

for some θ^{ρ^*}.

We can now generalize Theorem 3.6. Modifying the transfer rule in the Dutch auction to be the appropriate pay-as-bid rule, the following corollary holds.

Corollary 5.8. *Under any affiliated type distribution and value function satisfying Assumption 5.1, the statement of Theorem 3.6 holds.*

Proof. The proof that any Dutch auction is strongly shill-proof is the exact same as in the IPV case. By Proposition 5.7, we know what the orderly and optimal allocation rule is. By Theorem 3.6, we know that in the IPV case, no non-Dutch auction will be strongly shill-proof. We then apply Theorem 5.4 to conclude that all non-Dutch auctions are not strongly shill-proof and optimal in general. □

Many of our results for weakly shill-proof auctions continue to hold in our general model when we restrict to optimal transfer rules. The generalization of Theorem 4.3 is as follows.

Corollary 5.9. *Under any affiliated type distribution F and value function v satisfying Assumption 5.1, there exists no orderly auction $(\tilde{x}, \tilde{t}^*(\tilde{x}, \mu, \xi_0, v, F), \mu, \xi_0)$ that is single-action, mildly ex-post incentive compatible and weakly shill-proof.*

Proof. Applying Theorem 5.4 to Theorem 4.3 with a restriction on possible transfer rules completes the proof. □

While regularity is not as applicable with affiliation, our robustness result (Corollary 4.8) for the Dutch auction still hold.

Corollary 5.10. *Under any affiliated type distribution F and value function v satisfying Assumption 5.1 and for any public and efficient auction $(\tilde{x}, \tilde{t}^*(\tilde{x}, \mu, \xi_0, v, F), \mu, \xi_0)$, if the auction is not a Dutch auction with reserve price 0, there exists a value distribution under which the auction is not weakly shill-proof.*

Proof. Since the Dutch auction with reserve price 0 is strongly shill-proof, it must be weakly shill-proof. We then apply Theorem 5.4 to Corollary 4.8 to rule out any other auction format being robustly weakly shill-proof and complete the proof. □

We conclude this section by observing that Proposition 4.12 does not necessarily hold in the general model. In fact, as we mentioned in the introduction, even the English auction is not necessarily weakly shill-proof under an IPV type distribution if the distribution is not regular. We leave as a future research question to find conditions on type distributions such that some strategy-proof and weakly shill-proof mechanism exists.

6 Discussion

6.1 Shill-Proofness vs. Credibility

As discussed in the review of literature, another notion of "cheating" by the auctioneer is that of *in-credibility* introduced by Akbarpour and Li (2020). An auction is credible if a revenue-maximizing auctioneer has no incentive to lie about what other players are doing. That information environment differs from this paper because we assume that bidders

correctly (though perhaps not fully) perceive the actions of other players and where in the game tree they are. In the Online Appendix, we formally define credibility in our setting (Definition E.8) and prove the following implications:

Proposition 6.1. *Suppose (G, σ) is an orderly auction. If the auction is strongly shill-proof, then it must be credible. If the auction is credible, then it must be weakly shill-proof.*

We also define *ψ-credibility* for single-action auctions (Definition E.10) as a generalization of credibility allowing for bidders to have exogenous signals (where ψ is our notation from Section 4.1) about the actions of other bidders as well as additional communication from the auctioneer. Recall that we defined $\psi = \text{Id}$ to mean that the signals reveal the actions of previous bidders and $\psi = \varnothing$ to be static auction setting. We prove that the implications of Proposition 6.1 still hold and give conditions under which credibility coincides with strong and weak shill-proofness:

Proposition 6.2. *Suppose (G, σ) is a single-action, orderly auction. If the auction is strongly shill-proof, then it must be ψ-credible. If it is $(\psi = \varnothing)$-credible, then it is strongly shill-proof. If the auction is ψ-credible, then it must be weakly shill-proof. If it is weakly shill-proof, then it is $(\psi = \text{Id})$-credible.*

Proposition 6.2 implies that Theorem 4.3 is a generalization of the credibility trilemma (Akbarpour and Li, 2020, Theorem 1).

6.2 Dominant Strategies

This paper focuses on ex-post strategies. However, all our results can be extended to dominant strategies as well. To extend Theorem 3.6 from an ex-post strategy not to shill to a dominant strategy is straight-forward: shill bidding in the Dutch auction always leads to 0 revenue, which means it is a weakly dominated strategy, regardless of what other bidders do. Further, since there exist no other auctions besides the Dutch auction that have an ex-post strategy not to shill bid, there can exist no other auction with a dominant strategy not to shill bid.

To extend Proposition 4.12 from an ex-post equilibrium to a dominant strategy equilibrium for real bidders, some care must be taken in considering the information sets of different bidders when they take actions. However, we can provide dominant-strategy equilibria versions of the English and ascending, screening auctions by assuming bidders move simultaneously each round of the English auction, as well in the second-price auction phase of the ascending, screening auction. Theorem 4.3 holds if we were to instead consider dominant strategies for real bidders as dominant strategy incentive compatibility is a stronger condition than ex-post incentive compatibility.

References

Akbarpour, Mohammad and Shengwu Li (2020): "Credible auctions: A trilemma," *Econometrica*, 88 (2), 425–467.

Ashlagi, Itai and Yannai A. Gonczarowski (2018): "Stable matching mechanisms are not obviously strategy-proof," *Journal of Economic Theory*, 177, 405–425.

Ausubel, Lawrence M. and Paul Milgrom (2006): *The lovely but lonely Vickrey auction*, 22–26.

Bahrani, Maryam, Pranav Garimidi, and Tim Roughgarden (2024): "Centralization in Block Building and Proposer-Builder Separation," ArXiv:2401.12120.

Basu, Soumya, David Easley, Maureen O'Hara, and Emin Gün Sirer (2023): "StableFees: A Predictable Fee Market for Cryptocurrencies," *Management Science*, 69 (11), 6508–6524.

Chakraborty, Indranil and Georgia Kosmopoulou (2004): "Auctions with shill bidding," *Economic Theory*, 24, 271–287.

Chen, Kong-Pin, Ting-Peng Liang, Ted Chang, Yi-chun Liu, Shou-Yung Yin, and Ya-Ting Yu (2020): "How Serious is Shill Bidding in Online Auctions? Evidence from eBay Motors," Working Paper.

Chitra, Tarun, Matheus V. X. Ferreira, and Kshitij Kulkarni (2023): "Credible, optimal auctions via blockchains," ArXiv:2301.12532.

Chung, Hao and Elaine Shi (2023): "Foundations of Transaction Fee Mechanism Design," in *Proceedings of the 34th Annual ACM-SIAM Symposium on Discrete Algorithms (SODA)*, 3856–3899.

Conitzer, Vincent, Nicole Immorlica, Joshua Letchford, Kamesh Munagala, and Liad Wagman (2010): "False-name-proofness in social networks," in *Internet and Network Economics: 6th International Workshop, WINE 2010, Stanford, CA, USA, December 13-17, 2010. Proceedings 6*, Springer, 209–221.

Day, Robert and Paul Milgrom (2008): "Core-selecting package auctions," *International Journal of Game Theory*, 36, 393–407.

Dhangwatnotai, Peerapong, Tim Roughgarden, and Qiqi Yan (2015): "Revenue Maximization with a Single Sample," *Games and Economic Behavior*, 91, 318–333.

Elkind, Edith (2007): "Designing and learning optimal finite support auctions," in *Proceedings of the Eighteenth Annual ACM-SIAM Symposium on Discrete Algorithms (SODA)*, SIAM, 736–745.

Essaidi, Meryem, Matheus V. X. Ferreira, and S. Matthew Weinberg (2022): "Credible, strategyproof, optimal, and bounded expected-round single-item auctions for all distributions," ArXiv:2205.14758.

Ferreira, Matheus V. X. and S. Matthew Weinberg (2020): "Credible, truthful, and two-round (optimal) auctions via cryptographic commitments," in *Proceedings of the 21st ACM Conference on Economics and Computation*, 683–712.

Hafalir, Isa, Onur Kesten, Katerina Sherstyuk, and Cong Tao (2023): "When Speed is of Essence: Perishable Goods Auctions," Working Paper.

Haupt, Andreas and Zoë Hitzig (2021): "Contextually Private Mechanisms," ArXiv:2112.10812.

Izmalkov, Sergei (2004): "Shill Bidding and Optimal Auctions," MIT Working Paper.

Karlin, Samuel and Yosef Rinott (1980): "Classes of orderings of measures and related correlation inequalities. I. Multivariate totally positive distributions," *Journal of Multivariate Analysis*, 10 (4), 467–498.

Lamy, Laurent (2009): "The shill bidding effect versus the linkage principle," *Journal of*

Economic Theory, 144 (1), 390–413.
Lavi, Ron, Or Sattath, and Aviv Zohar (2022): "Redesigning Bitcoin's fee market," *ACM Transactions on Economics and Computation*, 10 (1), Art. 5.
Li, Shengwu (2017): "Obviously strategy-proof mechanisms," *American Economic Review*, 107 (11), 3257–3287.
Mackenzie, Andrew (2020): "A revelation principle for obviously strategy-proof implementation," *Games and Economic Behavior*, 124, 512–533.
Mackenzie, Andrew and Yu Zhou (2022): "Menu mechanisms," *Journal of Economic Theory*, 204, 105511.
Milgrom, Paul and Ilya Segal (2017): *Designing the US incentive auction*, Cambridge University Press, 803–812.
Milgrom, Paul R. and Robert J. Weber (1982): "A Theory of Auctions and Competitive Bidding," *Econometrica*, 50 (5), 1089–1122.
Porter, Ryan and Yoav Shoham (2005): "On Cheating in Sealed-Bid Auctions," *Decision Support Systems*, 39 (1), 41–54.
Pycia, Marek and Peter Troyan (2023): "A theory of simplicity in games and mechanism design," *Econometrica*, 91 (4), 1495–1526.
Roughgarden, Tim (2021): "Transaction fee mechanism design," *ACM SIGecom Exchanges*, 19 (1), 52–55.
Shinozaki, Hiroki (2024): "Shutting-out-proofness in object allocation problems with money," *Working Paper*.
Zeng, Haoyuan (2024): "Identity-Proof Auctions," Working Paper.

A Model (Section 2) Appendix

Definition A.1. Consider any set of real bidders R and tuple (G, σ). We restrict the set of potential deviations for shill bidders to

$$\Sigma_S = \{\sigma'_S : \forall \theta_{-S}, \exists \theta_S \text{ such that } (\sigma'_S, \sigma_{-S}(\theta_{-S})) = \sigma(\theta_S, \theta_{-S}; R = B)\}.$$

Then, the tuple (G, σ) is an **auction equilibrium** if for all $i \in R$ and deviating strategies σ'_i,

$$\mathbb{E}_{\theta'_{-i}, \tilde{R}}\left[u_i\left(\sigma\left(\theta_i, \theta'_{-i}; \tilde{R}\right)\right)\right] \geqslant \mathbb{E}_{\theta'_{-i}, \tilde{R}}\left[u_i\left(\sigma'_i\left(\theta_i, \theta'_{-i}; \tilde{R}\right), \sigma_{-i}\left(\theta_i, \theta'_{-i}; \tilde{R}\right)\right)\right],$$

and for all $i \in S$ and $\sigma'_S \in \Sigma_S$,

$$\mathbb{E}_{\theta'_{-S}}\left[u_i\left(\sigma\left(0, \theta'_{-S}; R\right)\right)\right] \geqslant \mathbb{E}_{\theta'_{-S}}\left[u_i\left(\sigma'_S\left(0, \theta'_{-S}; R\right), \sigma_{-S}\left(0, \theta'_{-S}; R\right)\right)\right].$$

In Definition A.1, the strategy space Σ_S is restricting shill bidders to only take actions that could be played on-path by a real bidder. In other words, shill bidders must act "as-if" they are real bidders and not take any actions that would definitively prove that they are shill bidders. This restriction allows us to move to a direct mechanism where shill-proofness is defined as it being an equilibrium (ex-interim for weak shill-proofness, ex-post for strong shill-proofness) for all shill bidders to report 0. Note that if we were to enlarge the

strategy space Σ_S to be the set $\hat{\Sigma}_S$ of all strategy profiles, are main results would not change. For Theorem 3.6, Proposition 4.6, and Proposition 4.12, we are focused on augmented direct games and in such games $\Sigma_S = \hat{\Sigma}_S$. For Theorem 4.3, we know that $\Sigma_S \subset \hat{\Sigma}_S$ and our impossibility result must still hold if the set of possible deviations by shill bidders is larger; thus, the theorem still holds.

Before we state our revelation principle in this context, we recall (with slight modification of notation) a definition and result from Akbarpour and Li (2020) that will be helpful in the proof.

Definition A.2 (Akbarpour and Li (2020), Definition 2)**.** A game equilibrium (G, σ) is **pruned** if, for any history h:

(i) There exists θ such that $h \preceq z(\sigma(\theta; B))$.

(ii) If $h \notin Z$, then $|\text{succ}(h)| \geqslant 2$.

(iii) If $h \notin Z$, then for $i = P(h)$, there exists θ_i,θ_i', and θ_{-i} such that

(a) $h \prec z(\sigma(\theta; B))$,

(b) $h \prec z(\sigma(\theta_i', \theta_{-i}; B))$, and

(c) $(x, t)(\sigma(\theta; B)) \neq (x, t)(\sigma(\theta_i', \theta_{-i}; B))$.

Lemma A.3 (Akbarpour and Li (2020), Proposition 1)**.** *If (G, σ) is a game equilibrium, then there exists a game equilibrium (G', σ') that is pruned and for all θ, $(x, t)(\sigma(\theta; B)) = (x', t')(\sigma'(\theta; B))$.*

Lemma A.4 (Augmented Revelation Principle)**.** *For every game equilibrium (G, σ) there exists an auction $(\tilde{x}, \tilde{t}, \mu, \xi_0)$ that meets the following conditions:*

(i) There exists a direct mechanism $(\tilde{x}, \tilde{t})$: for all θ, $\tilde{x}(\theta) = x(\sigma(\theta; B))$ and $\tilde{t}(\theta) = t(\sigma(\theta; B))$.

(ii) There exists a choice menu rule μ that is a function of the potential values $\Theta = \Theta_1 \times \cdots \times \Theta_N$ and bidder ξ. This rule has an output of $L \geqslant 2$ choices characterized as $\left\{\left(W_\ell, \vec{\xi}_\ell\right)\right\}_{\ell \in \{1,\ldots,L\}}$ where:

(a) $\{W_\ell\}_{\ell \in \{1,\ldots,L\}}$ forms a partition of Θ_ξ, $\vec{\xi}_\ell \in (B \cup \{\varnothing\}) \setminus \{\xi\}$, and $\vec{\xi} = \varnothing$ signifies the game has ended.

(b) For any ℓ such that $\vec{\xi}_\ell \neq \varnothing$, let $\hat{\Theta}^\ell = (\Theta_1, \ldots, \Theta_{\xi-1}, W_\ell, \Theta_{\xi+1}, \ldots, \Theta_N)$. Then, for any such ℓ, there exists $\theta_{\xi_\ell}, \theta'_{\xi_\ell} \in \hat{\Theta}^\ell_{\xi_\ell}$, and $\theta_{-\xi_\ell} \in \hat{\Theta}^\ell_{-\xi_\ell}$ such that $(\tilde{x}, \tilde{t})(\theta_{\xi_\ell}, \theta_{-\xi_\ell}) \neq (\tilde{x}, \tilde{t})\left(\theta'_{\xi_\ell}, \theta_{-\xi_\ell}\right)$. If $\theta_\xi \in W_\ell$, then the next player in the game is $\vec{\xi}_\ell$ and the menu presented to her is $\mu(\hat{\Theta}^\ell, \vec{\xi}_\ell)$.

(c) If ℓ is such that $\vec{\xi}_\ell = \varnothing$, then for all $\theta, \theta' \in \hat{\Theta}_\ell$, $\left(\tilde{x}, \tilde{t}\right)(\theta) = \left(\tilde{x}, \tilde{t}\right)(\theta')$.

(d) The first player to take an action is ξ_0, who is presented the menu $\mu(\vartheta^N, \xi_0)$.

Proof. To prove Condition i, we simply construct $(\tilde{x}, \tilde{t})$ by iterating over all possible θ and defining $(\tilde{x}, \tilde{t})$ as the outcome of $\sigma(\theta; B)$ in G.

To prove Condition ii, we first observe that by Definition A.1, shill bidders must act "as-if" they were real bidders and that we have restricted to pure strategies. Thus, we can always label actions as classes $\left(W_\ell, \vec{\xi}_\ell\right)$ of a partition of the remaining possible values for the current player ξ and satisfy Condition iia. The fact that $L \geqslant 2$ is equivalent to Conditions i and ii of Definition A.2, and Condition iii of Definition A.2 is equivalent to Condition iib here. We can then apply Lemma A.3 to find a game that satisfies these properties. Condition iic follows from the fact that G is well-defined (with each terminal history associated with a single outcome). Condition iid is simply mapping the first player in G to ξ_0 and the auctioneer has no information on bidders' values yet. □

B Strongly Shill-Proof Auctions (Section 3) Appendix

Lemma B.1. *An optimal auction (G, σ) is **winner-paying**: For all i and v,*

$$x_i\left(\sigma(\theta; B)\right) = 0 \implies t_i\left(\sigma(\theta; B)\right) = 0.$$

Proof. By the ex-post IR constraint, when $x_i\left(\sigma(\theta; B)\right) = 0$, we have $t_i\left(\sigma(\theta; B)\right) \leqslant 0$. It then follows from optimality that $t_i\left(\sigma(\theta; B)\right) = 0$. To see this, note that for bidder $j \neq i$, equilibrium constraints on bidder j slacken when moving from $t_i < 0$ to $t_i = 0$ and so her play will remain the same. Meanwhile the transfer from bidder i strictly increases moving from $t_i < 0$ to $t_i = 0$. □

Proof of Lemma 3.1.

Towards contradiction, suppose there exists a strongly shill-proof auction $(\tilde{x}, \tilde{t}, \mu, \xi_0)$, player ξ, and values $\theta_\xi, \theta_{-\xi}, \theta'_{-\xi}$ such that $\tilde{x}_\xi(\theta) = \tilde{x}_\xi(\theta_\xi, \theta'_{-\xi})$, but $\tilde{t}_\xi(\theta) \neq \tilde{t}_\xi(\theta_\xi, \theta'_{-\xi})$. WLOG, suppose $\tilde{t}_\xi(\theta) > \tilde{t}_\xi(\theta_\xi, \theta'_{-\xi})$. Now by Lemma B.1, ξ can only have two different transfers if that player wins the item under the allocation. Then, take $R = \{\xi\}$ and by monotonicity, $\tilde{t}_\xi(\theta) > \tilde{t}_\xi(\theta_\xi, \theta'_{-\xi}) \geqslant \tilde{t}(\theta_\xi, 0)$ and thus shilling increases revenue and the auction is not strongly shill-proof. □

Given our auction $(\tilde{x}, \tilde{t}, \mu, \xi_0)$, we define $X_i(\theta_i; \Theta)$ and $T_i(\theta_i; \Theta)$ to be the ex-interim quantity and transfer rules, respectively, when bidder i has value θ_i and the set of potential values for all bidders is Θ.

Lemma B.2. *For every optimal auction $(\tilde{x}^*, \tilde{t}, \mu, \xi_0)$, the ex-interim transfer rule for bidder i is*

$$T_i(\theta_i; \Theta) = X_i(\theta_i; \Theta)\theta_i - \sum_{m:\theta^{j_m} < \theta_i} \left[X_i(\theta^{j_m}; \Theta) \cdot (\theta^{j_{m+1}} - \theta^{j_m})\right],$$

where $\{\theta^{j_m}\}_m$ are the ordered atoms of Θ_i.

Proof. To prove that T has the claimed form, we will consider a specific non-truthful reporting: if a bidder has value θ^m, she commits to mis-reporting (selecting partitions) $\theta^{m'}$ for the

rest of the game. We now follow the proof of Theorem 1 of Elkind (2007). Since our direct mechanism is an equilibrium for real bidders, we must have that

$$X_i(\theta^{j_m};\Theta)\theta^{j_m} - T_i(\theta^{j_m};\Theta) \geqslant X_i(\theta^{j_{m-1}};\Theta)\theta^{j_m} - T_i(\theta^{j_{m-1}};\Theta), \text{ and}$$
$$X_i(\theta^{j_{m-1}};\Theta)\theta^{j_{m-1}} - T_i(\theta^{j_{m-1}};\Theta) \geqslant X_i(\theta^{j_m};\Theta)\theta^{j_{m-1}} - T_i(\theta^{j_m};\Theta).$$

Defining U_i to be the ex-interim utility for bidder i, the preceding expressions become:

$$U_i(\theta^{j_m};\Theta) \geqslant U_i(\theta^{j_{m-1}};\Theta) + (\theta^{j_m} - \theta^{j_{m-1}})X_i(\theta^{j_{m-1}};\Theta), \text{ and}$$
$$U_i(\theta^{j_{m-1}};\Theta) \geqslant U_i(\theta^{j_m};\Theta) - (\theta^{j_m} - \theta^{j_{m-1}})X_i(\theta^{j_m};\Theta).$$

Thus, $(\theta^{j_m} - \theta^{j_{m-1}})X_i(\theta^{j_{m-1}};\Theta) \leqslant U_i(\theta^{j_m};\Theta) - U_i(\theta^{j_{m-1}};\Theta) \leqslant (\theta^{j_m} - \theta^{j_{m-1}})X_i(\theta^{j_m};\Theta)$. Hence, any IC mechanism is such that

$$U_i(\theta^{j_m};\Theta) = U_i(\theta^{j_1};\Theta) + \sum_{k=2}^{m}(\theta^{j_m} - \theta^{j_{m-1}})\tilde{X}_i(\theta^{j_m};\Theta)$$
$$\text{where } \tilde{X}_i(\theta^{j_m};\Theta) \in \left[X_i(\theta^{j_{m-1}};\Theta), X_i(\theta^{j_m};\Theta)\right].$$

Therefore, we have that

$$T_i(\theta^{j_m};\Theta) = X_i(\theta^{j_m};\Theta)\theta^{j_m} - U_i(\theta^{j_1};\Theta) - \sum_{k=2}^{m}(\theta^{j_m} - \theta^{j_{m-1}})\tilde{X}_i(\theta^{j_m};\Theta). \tag{2}$$

By the ex-post IR condition, we have $U_i(\theta^{j_1};\Theta) \geqslant 0$ for all Θ. So, solving for the optimal transfer rule from Equation (2),

$$T_i^*(\theta^{j_m};\Theta) = \max_{U_i,\tilde{X}}\left[X_i(\theta^{j_m};\Theta)\theta^{j_m} - U_i(\theta^{j_1};\Theta) - \sum_{k=2}^{m}(\theta^{j_m} - \theta^{j_{m-1}})\tilde{X}_i(\theta^{j_m};\Theta)\right]$$
$$\text{such that } U_i(\theta^{j_1};\Theta) \geqslant 0 \text{ and } \tilde{X}_i(\theta^{j_m};\Theta) \in \left[X_i(\theta^{j_{m-1}};\Theta), X_i(\theta^{j_m};\Theta)\right].$$

The solution to this maximization is $U_i(\theta^{j_1};\Theta) = 0, \tilde{X}_i(\theta^{j_m};\Theta) = X_i(\theta^{j_{m-1}};\Theta)$. Thus, Equation (2) becomes

$$T_i(\theta_i;\Theta) = X_i(\theta_i;\Theta)\theta_i - \sum_{m:\theta^{j_m}<\theta_i}\left[X_i(\theta^{j_m};\Theta)\cdot(\theta^{j_{m+1}} - \theta^{j_m})\right]. \qquad \square$$

For any value choice $(W,\cdot) \in \mu(\cdot,\cdot)$, let us define $\underline{W} = \min_{w\in W}\{w\}$ and $\overline{W} = \max_{w\in W}\{w\}$, respectively.

Lemma B.3 (Extended Pay-as-Bid)**.** *Consider a strongly shill-proof and optimal auction $(\tilde{x}^*, \tilde{t}, \mu, \xi_0)$. Fix Θ, ξ and consider any $(W, \vec{\xi}) \in \mu(\Theta,\xi)$. If there exists $\theta, \theta' \in \Theta$ such that $\theta_\xi, \theta'_\xi \in W$ and $\mu(\Theta,\xi)$ is the last action ξ takes, then,*

$$\tilde{x}^*_\xi(\theta) = \tilde{x}^*_\xi(\theta') = 1 \implies \tilde{t}_\xi(\theta) = \tilde{t}_\xi(\theta') = \frac{T_\xi(\underline{W};\Theta)}{X_\xi(\underline{W};\Theta)},$$

i.e., transfers are constant conditional on allocation and are pinned down by the ex-interim outcome functions from the lowest type in the partition.

Proof. Since an auction cannot distinguish between values in the same choice set, we can apply Lemma 3.1 to conclude that if $\mu(\Theta,\xi)$ is the last action ξ takes, then $\tilde{t}(\theta) = \tilde{t}(\theta')$. To conclude the proof, we note that ξ wins no matter what her value is in W and then apply Lemma B.1 to conclude that $T_\xi(\underline{W};\Theta) = \tilde{t}_\xi(\theta)\cdot X_\xi(\underline{W};\Theta)$. $\square$

Proof of Proposition 3.3

Under a static information structure, the single-action, first-price auction directly corresponds to the direct mechanism. In a pay-as-bid mechanism, a shill bidder's actions have no way to influence the payments except by changing the allocation and by the definition of orderly, if a shill bidder changes the allocation she wins the auction, which is dominated. To prove that any more informative structure is not strongly shill-proof, consider some i, j and θ where $\mathcal{I}(h) \neq \mathcal{I}(h_0)$. By definition, information sets form a partition and so $\mathcal{I}(h) \cap \mathcal{I}(h_0) = \varnothing$. Observe that j can have any type at both h and h_0 given that each bidder takes a single action. Now, consider the minimal m such that j believes ex-interim that she could win with type θ^m. By assumption, $\theta^m > \theta^{\rho^*}$. Because we are considering strong shill-proofness, we can assume without loss that j does in fact win the auction at θ^m. So, at h, we can then apply Lemma B.3 to complete the proof that j must pay more at h than at h_0 and so if $i \in S$, shill bidding is profitable at this type realization. □

Proof of Theorem 3.6.

We first show that the Dutch auction with reserve price θ^ρ is a well-defined, i.e., that the stopping rule allows for the auction to be orderly.[28] We then show that it is strongly shill-proof. Finally, we show that there are no other public, strongly shill-proof, orderly, optimal auctions.

The Dutch Auction is Orderly, Optimal, and Strongly Shill-Proof. The Dutch auction quantity rule is orderly and the transfer rule is ex-post IR and monotone. Indeed, by construction, the next player $\vec{\xi}$ is always the player with the potentially highest value (including for tie-breaking). So, if that player indicates that she is of the highest possible type, the outcome (allocation and transfer) is fully determined and the auction ends. The auction ends once there are no players who could have values weakly greater than θ^ρ.

We now prove that the Dutch auction is strongly shill-proof. Towards contradiction, suppose not. So, there must exist some S, $\xi \in S$, and Θ such that $\left\{\overline{\Theta}_\xi\right\}$ is selected from the menu $\mu(\Theta, \xi)$. But by construction, this means that the auction immediately ends and the good is allocated to the shill bidder who misreported. By Lemma B.1, the revenue from this deviation is 0, which must be weakly less than any other possible transfer.

Uniqueness. Towards contradiction, suppose there exists a menu rule $\tilde{\mu} \neq \mu^D_{\rho^*}$ that is associated with a public, strongly shill-proof, orderly, optimal auction. Therefore, there exists Θ and ξ such that $\tilde{\mu}(\Theta, \xi) \neq \mu^D_{\rho^*}(\Theta, \xi)$. Without loss, we will suppose that Θ is the first time in the game tree that $\tilde{\mu}$ differs from $\mu^D_{\rho^*}$. Formally, for all $\hat{\Theta} \supsetneq \Theta$, $\tilde{\mu}\left(\hat{\Theta}, \xi\right) = \mu^D\left(\hat{\Theta}, \xi\right)$. We now proceed in cases.

Case 1 (Different Next Player Choice). Suppose $\tilde{\mu}(\Theta, \xi) = \left\{\left(W_{\mathsf{L}}, \tilde{\vec{\xi}}_{\mathsf{L}}\right), \left(W_{\mathsf{H}}, \tilde{\vec{\xi}}_{\mathsf{H}}\right)\right\}$, where $\tilde{\vec{\xi}}_{\mathsf{L}} \neq \vec{\xi}_{\mathsf{L}}$ or $\tilde{\vec{\xi}}_{\mathsf{H}} \neq \vec{\xi}_{\mathsf{H}}$. If $\tilde{\vec{\xi}}_{\mathsf{L}} = \vec{\xi}_{\mathsf{L}}$, then $\tilde{\vec{\xi}}_{\mathsf{H}} = \vec{\xi}_{\mathsf{H}}$ because the outcome is fully resolved once a bidder selects the high partition. So, we need only consider the case where $\tilde{\vec{\xi}}_{\mathsf{L}} \neq \vec{\xi}_{\mathsf{L}}$. By Definition 3.2, $\tilde{\vec{\xi}}_{\mathsf{H}} \neq \varnothing$ (even if the ξ chooses $\left\{\overline{\Theta}_\xi\right\}$). Now, there must exist some bidders b_1 and b_2 where b_1 is called before b_2 but $(\cdot, b_2) \rhd (\cdot, b_1)$ as otherwise the auction calls players in

[28] The auction is public by definition.

the same order as the Dutch auction, which we assumed was not the case. Let $R = B \setminus \{b_1\}$, and for some m, $\theta_{b_2} = \theta^m$ and $\theta_i = \theta^{m-2}$ for all bidders $i \notin \{b_1, b_2\}$. Taking the expression from Lemma B.2 and dividing both sides by X_{b_2}, we get

$$\frac{T_{b_2}(\theta^m;\Theta)}{X_{b_2}(\theta^m;\Theta)} = \theta^m - \sum_{k:\theta^{j_k}<\theta^m} \frac{X_{b_2}(\theta^{j_k};\Theta)}{X_{b_2}(\theta^m;\Theta)} \cdot (\theta^{j_{k+1}} - \theta^{j_k}) < \theta^m.$$

(Note that there must be at least one such k in the summation because otherwise $\Theta_{b_2} = \{\theta^m\}$ and b_2 would not take an action.)

The last choice b_2 makes is to select a partition W such that $\underline{W} \geqslant \theta^{m-2}$. We can therefore apply Lemma B.3 to conclude that the transfer if b_1 reports 0 must be $\frac{T_{b_2}(\underline{W};\Theta)}{X_{b_2}(\underline{W};\Theta)} \leqslant \frac{T_{b_2}(\theta^m;\Theta)}{X_{b_2}(\theta^m;\Theta)} < \theta^m$. If bidder b_1 instead reports θ^m, then bidder b_2 will win and the revenue will be θ^m and so the auction will not be strongly shill-proof.

Case 2 (Different Partitions). Suppose there exists $(W, \cdot) \in \tilde{\mu}(\Theta, \xi)$ such that $W \notin \{W_{\mathsf{L}}, W_{\mathsf{H}}\}$. We need not consider the case where $\underline{W} = 0$ because in that case either we can consider some other choice $W' \notin \{W_{\mathsf{L}}, W_{\mathsf{H}}\}$ or $W = \Theta_\xi$ which would violate Lemma A.4. We need not consider $\underline{W} > \overline{\Theta}_{-\xi}$ because Θ is the first time $\tilde{\mu}$ differs from $\mu^D_{\rho^*}$ and for all Θ, ξ, and $(\tilde{W}, \cdot) \in \mu^D_{\rho^*}(\Theta, \xi)$, it is the case that $\underline{\tilde{W}} \leqslant \overline{\Theta}_{-\xi}$. So, there are now only three sub-cases we must consider: $\underline{W} \in [\theta^{\rho^*}, \overline{\Theta}_{-\xi})$, $\underline{W} = \overline{\Theta}_{-\xi}$, and $\underline{W} \in (0, \theta^{\rho^*})$.

Case 2a ($\underline{W} \in [\theta^{\rho^*}, \overline{\Theta}_{-\xi})$). In this sub-case, there exists m^* such that $\theta^{\rho^*} \leqslant \underline{W}_\ell \leqslant \theta^{m^*} < \overline{\Theta}_{-\xi}$. Since Θ is the first time that $\tilde{\mu}$ differs from $\mu^D_{\rho^*}$, we can suppose there exists i such that $(\theta^{m^*}, \xi) \rhd (\theta^{\rho^*}, i)$ because otherwise the outcome would already be resolved or the player rotation would be the only difference (Case 1). Then, suppose bidder i is such that $i \in R$ and $\theta_i \geqslant \theta^{m^*+1}$. Take bidder $\xi \in S$ to shill θ^{m^*}; and for $k \notin \{i, \xi\}$, take $\theta_k = \underline{\Theta}_k < \theta^{m^*}$. Therefore, by Lemma B.2, observe that for the last action i takes, her ex-interim transfer must be higher when shill ξ reports θ^{m^*} than when she reports 0. Thus by Lemma B.3, when $\theta^{m^*} < \overline{\Theta}_{-\xi}$, there exists a valuation vector θ such that a shill bidder would want to deviate away from reporting 0—and therefore such an auction is not strongly shill-proof.

Case 2b ($\underline{W} = \overline{\Theta}_{-\xi}$). In this sub-case, we know $\theta^{\rho^*} \leqslant \underline{W}_\ell = \overline{\Theta}_{-\xi}$. Since Θ has been generated via a Dutch auction so far, the current player has the lowest tie-breaking priority, i.e., ξ is such that for all $j \neq \xi$, $(\overline{\Theta}_j, j) \rhd (\overline{\Theta}_j, \xi)$. Letting $j \in R$ and $\theta_j = \overline{\Theta}_{-\xi}$, take bidder $\xi \in S$ to report $\overline{\Theta}_{-\xi}$; and for all $k \notin \{j, \xi\}$, take $\theta_k = \underline{\Theta}_k < \overline{\Theta}_{-\xi}$. As noted, $(\overline{\Theta}_j, j) \rhd (\overline{\Theta}_j, \xi)$ and so bidder j is allocated the item and not shill bidder ξ. Therefore, by the same argument as Case 2a, shill bidding will increase revenue.

Case 2c ($\underline{W} \in (0, \theta^{\rho^*})$). By Condition iib of Lemma A.4, there must be some chance that ξ could win the auction in order to affect outcomes. In particular, by definition of $\tilde{x}^*$, there exists j and $\theta_j \in \Theta_j$ such that $(\theta_j, j) \rhd (\theta^{\rho^*}, \xi)$ and by Lemma 3.1, there exists $\theta'_j \in \Theta_j$ such that $(\overline{W}, \xi) \rhd (\theta'_j, j)$. If any bidder is ever offered a choice with $\underline{W} \geqslant \theta^{\rho^*}$, then the previous cases imply that a shill bidder can profitably deviate and so we only have to consider the instances where no such choice is offered. Now, suppose $R = \{j\}$. Suppose all shill bidders play the strategy of selecting the partition $\tilde{W}$ such that $0 < \underline{\tilde{W}} < \theta^{\rho^*}$ if such a choice is available. Let the final move that j takes to be $W^{0,j,\text{last}}$ and $W^{S,j,\text{last}}$ under the shill bidders' strategy of selecting 0 and not, respectively. Similarly, define $\Theta^{0,\text{last}}, \Theta^{S,\text{last}}$ as the possible values and $\tilde{t}^0_j, \tilde{t}^S_j$ as the transfers under these respective strategies. Observe that it is without

loss to assume that $\theta^\omega \equiv \underline{W}^{0,j,\text{last}} \leqslant \underline{W}^{S,j,last}$ because in the latter case the shill bidders are always acting as if they have higher values than in the former case. Next, for $c \in \{0, S\}$, let $p_i^{m,c} = \mathbb{P}\left[(\theta^m, j) \rhd (\theta_i, i) \mid \Theta_i^{c,\text{last}}\right]$ and define $\zeta_{i,m} = \frac{p_i^{m,0} p_i^{\omega,S}}{p_i^{m,S} p_i^{\omega,0}}$. Observe that for all $i \neq j$ and $m \leqslant \omega$, it is the case that $\zeta_{i,m} \geqslant 1$ with at least one strict inequality because $\overline{\Theta}_i^{S,\text{last}} \geqslant \overline{\Theta}_i^{0,\text{last}}$ for all i with strict inequality for at least one i and m. So, by Lemmatta B.2 and B.3,[29]

$$\begin{aligned}
\tilde{t}_j^S - \tilde{t}_j^0 &\geqslant \frac{T_\xi(\theta^\omega; \Theta^{S,\text{last}})}{X_\xi(\theta^\omega; \Theta^{S,\text{last}})} - \frac{T_\xi(\theta^\omega; \Theta^{0,\text{last}})}{X_\xi(\theta^\omega; \Theta^{0,\text{last}})} \\
&= \sum_{\rho^* \leqslant k < \omega} \left((\theta^{k+1} - \theta^k) \prod_{i \neq j} \frac{p_i^{k,0}}{p_i^{\omega,0}} \right) - \sum_{\rho^* \leqslant k < \omega} \left((\theta^{k+1} - \theta^k) \prod_{i \neq j} \frac{p_i^{k,S}}{p_i^{\omega,S}} \right) \\
&= \frac{1}{\prod_{i \neq j} p_i^{\omega,0} p_i^{\omega,S}} \cdot \left(\sum_{\rho^* \leqslant k < \omega} (\theta^{k+1} - \theta^k) \cdot \left(\prod_{i \neq j} p_i^{k,S} p_i^{\omega,0} (\zeta_{i,k} - 1) \right) \right) > 0.
\end{aligned}$$

Thus, we have described a profitable shill bidding strategy in this sub-case. □

C Weakly Shill-Proof Auctions (Section 4) Appendix

C.1 Single-Action Auctions (Section 4.1) Appendix

Lemma C.1. *For any single-action auction, there exist unique $\tilde{x} : \vartheta^N \to \{0,1\}^N$ and $\tilde{t} : \vartheta^N \to \mathbb{R}^N$ such that:*

(i) (Correspondence) For all $v \in \vartheta^N$, $\tilde{x}^*(\theta) = x(\sigma(\theta; B))$ and $\tilde{t}(\theta) = t(\sigma(\theta; B))$.

(ii) (Individual Rationality) For all $i \in B$ and θ, $\tilde{x}_i^*(\theta)\theta_i - \tilde{t}_i(\theta) \geqslant 0$.

(iii) (Incentive Compatibility) For all R, $i \in R$, θ_i, and θ_i',

$$\begin{aligned}
&\mathbb{E}_{\theta_{-i},\tilde{R}}\left[\tilde{x}_i^*\left(\sigma(\theta;\tilde{R})\right)\theta_i - \tilde{t}_i(\sigma(\theta;\tilde{R})) \mid \theta_i, s_i = \psi(\theta_{j<i})\right] \\
&\qquad \geqslant \mathbb{E}_{\theta_{-i},\tilde{R}}\left[\tilde{x}_i^*\left(\sigma(\theta_i', \theta_{-i};\tilde{R})\right)\theta_i - \tilde{t}_i\left(\sigma(\theta_i', \theta_{-i};\tilde{R})\right) \mid \theta_i, s_i = \psi(\theta_{j<i})\right]. \quad (3)
\end{aligned}$$

Proof. To begin, let us note that we can uniquely define $(\tilde{x}, \tilde{t})$ point-wise based on the outcomes in G from playing $\sigma(\theta; B)$. Next, the IR constraint (Condition ii) follows immediately from the ex-post IR condition and our construction of $(\tilde{x}, \tilde{t})$. Finally, Equation (3) comes from Definition A.1 and recalling that we restrict shill bidders to actions that could have been taken by real bidders. □

[29] We may assume that the possible values for j are sequential above the reserve otherwise we could consider j as a shill bidder for some other bidder instead by the cases above.

Lemma C.2. *If a single-action auction is weakly shill-proof, then for all R, $\theta_{j<\min S}$,*[30] *and $\{\theta_i\}_{i\in S}$,*

$$\mathbb{E}_\theta\left[\sum_{k\in R}\tilde{t}_k\left(\{\theta_i\}_{i\in S},\{\theta_i\}_{i\notin S}\right)\mid\theta_{j<\min S}\right]\leqslant\mathbb{E}_\theta\left[\sum_{k\in R}\tilde{t}_k\left(0,\{\theta_i\}_{i\notin S}\right)\mid\theta_{j<\min S}\right].$$

Proof. Towards contradiction, suppose there exists R, $\theta_{j<\min S}$, and $\{\theta_i\}_{i\in S}$ such that

$$\mathbb{E}_\theta\left[\sum_{k\in R}\tilde{t}_k\left(\{\theta_i\}_{i\in S},\{\theta_i\}_{i\notin S}\right)\mid\theta_{j<\min S}\right]>\mathbb{E}_\theta\left[\sum_{k\in R}\tilde{t}_k\left(0,\{\theta_i\}_{i\notin S}\right)\mid\theta_{j<\min S}\right].$$

We now prove that the deviation by the coalition S where they report $\{\theta_i\}_{i\in S}$ is profitable and therefore that the auction is not weakly shill-proof. By assumption, a shill bidder observes actions by all bidders who take actions before her. So, $\{\theta_i\}_{i\in S}$ can condition on $\theta_{j<\min S}$ when making decisions. Then, the strategy by S of committing to report $\{\theta_i\}_{i\in S}$ regardless of what other bidders play after $\min_{i\in S}\{i\}$ must be strictly profitable compared to always reporting 0. Thus, we have found a strategy that does strictly better than always reporting 0: When the values before $\min_{i\in S}\{i\}$ are reported as $\theta_{j<\min S}$, report $\{\theta_i\}_{i\in S}$. Otherwise, report 0. This strategy in the direct game immediately translates to a profitable deviation in the auction by Definition A.1 and Lemma C.1 and thus the equilibrium is not weakly shill-proof. □

Now, when discussing single-action auctions, we focus on the direct mechanisms associated to weakly shill-proof auctions and so we will refer to an auction as $(\tilde{x},\tilde{t},\psi)$ without reference to R.

Lemma C.3. *Suppose a single-action, orderly auction $(\tilde{x},\tilde{t},\psi)$ is weakly shill-proof. Then, for all i, θ, and $\theta'_{j>i}$, $\left[\tilde{x}_i(\theta)=\tilde{x}_i\left(\theta_{j\leqslant i},\theta'_{j>i}\right)\implies\tilde{t}_i(\theta)=\tilde{t}_i\left(\theta_{j\leqslant i},\theta'_{j>i}\right)\right]$.*

Proof. Towards contradiction, suppose there exists i, θ, and $\theta'_{j>i}$, such that $\tilde{x}_i(\theta)=\tilde{x}_i\left(\theta_{j\leqslant i},\theta'_{j>i}\right)$, but $\tilde{t}_i(\theta;s)>\tilde{t}_i\left(\theta_{j\leqslant i},\theta'_{j>i};s\right)$. Because the auction is orderly, we can apply Lemma B.1 to conclude that $\tilde{x}_i(\theta)=\tilde{x}_i\left(\theta_{j\leqslant i},\theta'_{j>i}\right)=1$. Let $R=\{1,\ldots,i\}$. Then,

$$\mathbb{E}_\theta\left[\sum_{k\in R}\tilde{t}_k(\{\theta_i\}_{i\in S},\{\theta_i\}_{i\notin S})\mid\theta_{j\leqslant\min S}\right]=\tilde{t}_i(\theta)>\tilde{t}_i\left(\theta_{j\leqslant i},\theta'_{j>i}\right)\geqslant\tilde{t}_i\left(\theta_{j\leqslant i},0\right).$$

This violates Lemma C.2, and so we have reached a contradiction. □

Lemma C.4. *Suppose a single-action, orderly auction $(\tilde{x},\tilde{t},\psi)$ is mildly ex-post incentive compatible and weakly shill-proof. Then, there exists $i<N$ such that for all θ_i, θ'_i, and $\theta_{-i}\in\psi_i^{-1}(\theta_{j<i})$, $\left[\tilde{x}_i(\theta)=\tilde{x}_i(\theta'_i,\theta_{-i})\implies\tilde{t}_i(\theta)=\tilde{t}_i(\theta'_i,\theta_{-i})\right]$.*

Proof. Consider $i<N, R\ni i,\theta_i,\theta'_i$, and $\theta_{-i}\in\psi_i^{-1}(\theta_{j<i})$ such that $\tilde{x}_i(\theta)=\tilde{x}_i(\theta'_i,\theta_{-i})$. WLOG, suppose $\theta_i>\theta'_i$. By monotonicity, $\tilde{t}_i(\theta)\geqslant\tilde{t}_i(\theta'_i,\theta_{-i})$. Towards contradiction, suppose $\tilde{t}_i(\theta)>\tilde{t}_i(\theta'_i,\theta_{-i})$. By the winner-paying property, $\tilde{t}_i(\theta)>\tilde{t}_i(\theta'_i,\theta_{-i})$ implies that $\tilde{x}_i(\theta)=\tilde{x}_i(\theta'_i,\theta_{-i})=1$. However, note that $\tilde{t}_i(\theta)>\tilde{t}_i(\theta'_i,\theta_{-i})$ would mean that the utility of reporting θ'_i would be higher than truthful reporting under true value θ_i which would violate the mildly ex-post incentive compatiblility and thus $\tilde{t}_i(\theta'_i,\theta_{-i})=\tilde{t}_i(\theta'_i,\theta_{-i})$. □

[30] $\theta_{j<\min S}\equiv\{\theta_j:j<\min_{i\in S}\{i\}\}$.

Proof of Theorem 4.3.

Towards contradiction, suppose such an auction did exist. Fix $i < N$ and $\theta_{j<i}$ and then suppose that $\theta_i < \theta^M$. Combining Lemmatta C.3 and C.4, we can see that for all θ_i' and $\theta_{-i}, \theta_{-i}' \in \psi_i^{-1}(\theta_{j<i})$, if $\tilde{x}_i(\theta) = \tilde{x}_i(\theta')$, then $\tilde{t}_i(\theta) = \tilde{t}_i(\theta')$. So, define $\tilde{t}_i$ as the (constant) $\tilde{t}_i(\theta)$ for all θ such that $\tilde{x}_i(\theta) = 1$.

In order for bidder i to have an ex-post strategy, when $\tilde{x}(\theta) = 1$, it must also be the case that $\tilde{x}_i(\theta^M, \theta_{-i}) = 1$. So, applying the winner-paying property (and suppressing that the expectation is conditioned on $s_i = \psi_i(\theta_{j<i})$), we have

$$\begin{aligned}\mathbb{E}_{\theta_{-i}}\left[\tilde{x}_i\left(\theta^M, \theta_{-i}\right)\theta_i - \tilde{t}_i\left(\theta^M, \theta_{-i}\right)\right] \\ = \mathbb{E}_{\theta_{-i}}\left[\theta_i - \tilde{t}_i^* \mid \tilde{x}_i(\theta) = 1\right] + \mathbb{E}_{\theta_{-i}}\left[\theta_i - \tilde{t}_i^* \mid \tilde{x}_i(\theta) = 0, \tilde{x}_i\left(\theta^M, \theta_{-i}\right) = 1\right], \quad (4)\end{aligned}$$

$$\text{and } \mathbb{E}_{\theta_{-i}}\left[\tilde{x}_i(\theta)\theta_i - \tilde{t}_i(\theta)\right] = \mathbb{E}_{\theta_{-i}}\left[\theta_i - \tilde{t}_i^* \mid \tilde{x}_i(\theta) = 1\right]. \quad (5)$$

Taking the difference between Equation (4) and Equation (5), we see that

$$\begin{aligned}\mathbb{E}_{\theta_{-i}}\left[\tilde{x}_i\left(\theta^M, \theta_{-i}\right)\theta_i - \tilde{t}_i\left(\theta^M, \theta_{-i}\right)\right] - \mathbb{E}_{\theta_{-i}}\left[\tilde{x}_i(\theta)\theta_i - \tilde{t}_i(\theta)\right] \\ = \mathbb{E}_{\theta_{-i}}\left[\theta_i - \tilde{t}_i^* \mid \tilde{x}_i(\theta) = 0, \tilde{x}_i\left(\theta^M, \theta_{-i}\right) = 1\right] \quad (6)\end{aligned}$$

Now, by definition of orderly, $\tilde{x}$ is monotone, and by assumption $\tilde{t}$ is monotone. If there exists θ^m such that $\mathbb{P}[\tilde{x}_i(\theta^m, \theta_{-i}) = 1] < \mathbb{P}[\tilde{x}_i(\theta^M, \theta_{-i}) = 1]$, then for such m,

$$\begin{aligned}\mathbb{E}_{\theta_{-i}}\left[\tilde{x}_i(\theta^m, \theta_{-i})\theta^m - \tilde{t}_i(\theta^m, \theta_{-i})\right] \geqslant \mathbb{E}_{\theta_{-i}}\left[\tilde{x}_i(\theta^{m-1}, \theta_{-i})\theta^m - \tilde{t}_i(\theta^{m-1}, \theta_{-i})\right] \\ > \mathbb{E}_{\theta_{-i}}\left[\tilde{x}_i(\theta^{m-1}, \theta_{-i})\theta^{m-1} - \tilde{t}_i(\theta^{m-1}, \theta_{-i})\right] \geqslant 0,\end{aligned}$$

where the last inequality comes from the IR condition. Thus, the IR constraint does not bind for $\theta_i = \theta^m$. Since the good has to be allocated to the highest type, for all $i < N$, there exists θ_{-i} such that $\tilde{x}_i(\theta) = 0$ and $\tilde{x}_i(\theta^M, \theta_{-i}) = 1$. Thus,

$$\mathbb{E}_{\theta_{-i}}\left[\theta_i - \tilde{t}_i^* \mid \tilde{x}_i(\theta) = 0, \tilde{x}_i\left(\theta^M, \theta_{-i}\right) = 1\right] > 0. \quad (7)$$

Combining Equations (6) and (7), we see that

$$\mathbb{E}_{\theta_{-i}}\left[\tilde{x}_i\left(\theta^M, \theta_{-i}\right)\theta_i - \tilde{t}_i\left(\theta^M, \theta_{-i}\right)\right] > \mathbb{E}_{\theta_{-i}}\left[\tilde{x}_i(\theta)\theta_i - \tilde{t}_i(\theta)\right]. \quad (8)$$

We then apply the weak shill-proofness condition to simplify Equation (3) to

$$\mathbb{E}_{\theta_{-i}}\left[\tilde{x}_i(\theta)\theta_i - \tilde{t}_i(\theta)\right] \geqslant \mathbb{E}_{\theta_{-i}}\left[\tilde{x}_i(\theta_i', \theta_{-i})\theta_i - \tilde{t}_i(\theta_i', \theta_{-i})\right].$$

Taking $\theta_i' = \theta^M$, Equation (8) violates the IC constraint from Lemma C.1—and thus we have reached a contradiction. □

C.2 Efficient Auctions (Section 4.2.1) Appendix

In order to build towards a proof Proposition 4.6, we will prove that for a certain class of value distributions, every *weakly* shill-proof and efficient auction must be a semi-Dutch auction. Formally, we assume that the value distribution is sparse:

Definition C.5. A regular distribution F is **sparse** if for all $k < \rho^*$,

$$\theta^k - (\theta^{k+1} - \theta^k)\frac{f(\theta^{k+1})}{f(\theta^k)} < 0. \tag{9}$$

A distribution is sparse if the atoms are sufficiently far apart. Sparsity can also be a reasonable assumption if the auctioneer has preferences for the auction to be completed quickly, or otherwise finds it costly to distinguish between values that are close to each other.

Lemma C.6. *Consider an efficient auction and suppose F is regular and sparse. Let R, Θ such that $\Theta_i = \left\{w : w \in [\underline{\Theta}_i, \overline{\Theta}_i]\right\}$ for all $i \in R$, and consider $(\underline{W}, j)$ such that $\underline{W} < \theta^{\rho^*}, j \notin R$ and for all $i \in R, (\overline{\Theta}_i, i) \rhd (\underline{W}, j)$. Then, for all $\gamma < \underline{W}$,*

$$\mathbb{E}\left[\sum_{i\in R}\tilde{t}_i(\theta) \mid \Theta = (\Theta_{-j}, \{\gamma\})\right] < \mathbb{E}\left[\sum_{i\in R}\tilde{t}_i(\theta) \mid \Theta = (\Theta_{-j}, \{\underline{W}\})\right].$$

Thus, the following shilling strategy is profitable compared to always reporting 0: if there exists (Θ, ξ, W) such that $(W, \cdot) \in \mu(\Theta, \xi)$ and $\underline{W} \in (0, \theta^{\rho^*})$, then select W. Otherwise, select the partition containing 0.

Proof. Consider any $i \in R$ and Θ and let $C = \left(\sum_{\theta^k\in\Theta_i} f(\theta^k)\right)^{-1}$. Then, applying Equation (2),

$$\begin{aligned}
&\mathbb{E}\left[\tilde{t}_i(\theta) \mid \Theta\right] + U_i(\theta^{j_1}; \Theta) = \mathbb{E}\left[T_i(\theta_i; \Theta)\right] + U_i(\theta^{j_1}; \Theta) = C\sum_m f(\theta^{j_m})T_i(\theta^{j_m}; \Theta) \\
&= C\sum_{m:\theta^{j_m}\in\Theta_i} f(\theta^{j_m})\left(X_i(\theta^{j_m}; \Theta)\theta^{j_m} - \sum_{k<m}\left[\tilde{X}_i^k(\Theta)\cdot(\theta^{j_{k+1}} - \theta^{j_k})\right]\right) \\
&= C\left[\sum_{m:\theta^{j_m}\in\Theta_i} f(\theta^{j_m})X_i(\theta^{j_m}; \Theta)\theta^{j_m} - \sum_{m:\theta^{j_m}\in\Theta_i}\sum_{k<m} f(\theta^{j_m})\left[\tilde{X}_i^k(\Theta)\cdot(\theta^{j_{k+1}} - \theta^{j_k})\right]\right] \\
&= C\sum_{m:\theta^{j_m}\in\Theta_i}\left[\theta^{j_m}X_i(\theta^{j_m}; \Theta) - (\theta^{j_{m+1}} - \theta^{j_m})\frac{F(\overline{\Theta}_i) - F(\theta^{j_m})}{f(\theta^{j_m})}\tilde{X}_i^m(\Theta)\right] f(\theta^{j_m}).
\end{aligned}$$

Applying the definition of the efficient allocation rule $\tilde{x}^E$, we know that for $(\theta^m, i) \rhd (\gamma, j)$ and $(\theta^m, i) \rhd (\gamma', j)$, we can define $X_i(\theta^m; \Theta_{-j}) \equiv X_i(\theta^m; \Theta_{-j}, \{\gamma\}) = X_i(\theta^m; \Theta_{-j}, \{\gamma'\})$. Note that $\underline{W} \leqslant \min_i\left\{\overline{\Theta}_i\right\}$ by assumption and therefore, for $\underline{W} \in \left(\gamma, \theta^{\rho^*}\right)$,

$$\begin{aligned}
&\mathbb{E}\left[\tilde{t}_i(\theta) \mid \Theta = (\Theta_{-j}, \{\gamma\})\right] - \mathbb{E}\left[\tilde{t}_i(\theta) \mid \Theta = (\Theta_{-j}, \{\underline{W}\})\right] \\
&= C\sum_{m:\gamma\leqslant\theta^{j_m}<\underline{W}}\left[\theta^{j_m}X_i(\theta^{j_m}; \Theta) - (\theta^{j_{m+1}} - \theta^{j_m})\frac{F(\overline{\Theta}_i) - F(\theta^{j_m})}{f(\theta^{j_m})}\tilde{X}_i^m(\Theta)\right] f(\theta^{j_m}) \\
&\leqslant C\sum_{m:\gamma\leqslant\theta^{j_m}<\underline{W}}\left[\theta^{j_m} - (\theta^{j_m+1} - \theta^{j_m})\frac{f(\theta^{j_m+1})}{f(\theta^{j_m})}\right] f(\theta^{j_m})X_i(\theta^{j_m}; \Theta_{-j}) < 0
\end{aligned}$$

where the final inequality comes from sparsity. And so,

$$\mathbb{E}\left[\sum_{i\in R}\tilde{t}_i(\theta) \mid \Theta = (\Theta_{-j}, \{\gamma\})\right] < \mathbb{E}\left[\sum_{i\in R}\tilde{t}_i(\theta) \mid \Theta = (\Theta_{-j}, \{\underline{W}\})\right],$$

as claimed in the statement of the lemma. Thus, committing to misreport as $\underline{W}$ is strictly beneficial compared to any strategy that can only report $\gamma < \underline{W}$. □

Lemma C.7. *If F is regular and sparse, then every public, weakly shill-proof, and efficient auction is a semi-Dutch auction with cutoff θ^{ρ^*}.*

Proof. Suppose F is regular and sparse. Consider an arbitrary weakly shill-proof and efficient auction, $(\tilde{x}^E, \tilde{t}, \mu, \xi_0)$, and consider any θ such that $\max_i\{\theta_i\} < \theta^{\rho^*}$. We prove that both parts of the definition of a semi-Dutch auction are necessary.

Proof of Condition i of Definition 4.5. First we prove that if, for any player ξ and set of possible values Θ, there exists $(W, \cdot) \in \mu(\Theta, \xi)$ where $0 < \underline{W} < \theta^{\rho^*}$, then $\Theta \subseteq \breve{\Theta}$. Towards contradiction, suppose that there exists a (ξ, Θ, W) such that $\Theta \nsubseteq \breve{\Theta}$, $(W, \cdot) \in \mu(\Theta, \xi)$ $\underline{W} \in \left(0, \theta^{\rho^*}\right)$, and $\xi \in S$.

Let us first prove it is without loss to assume (ξ, Θ, W) is such that for all i, $\underline{\Theta}_i = 0$ or $\underline{\Theta}_i \geqslant \theta^{\rho^*}$. If there exists (ξ, Θ, W) and i such that $\underline{\Theta}_i \in (0, \theta^{\rho^*})$, let us label that set as Θ^K and let $\Theta^0 \supset \Theta^1 \supset \ldots \supset \Theta^K$ be the sequence of on-path possible value sets preceding Θ^K. Let the players called along the path be $\xi^0, \xi^1, \ldots, \xi^K$ and the value partition selected by player k to be W^k. Note that $\Theta^0 = \vartheta^N$ and so for all i, $\underline{\tilde{\Theta}}_i = 0$ or $\underline{\tilde{\Theta}}_i \geqslant \theta^{\rho^*}$. So, the set $\mathcal{K} = \left\{k < K : \underline{W^k} \in (0, \theta^{\rho^*})\right\} \neq \varnothing$ and therefore $k^* = \min_{k\in\mathcal{K}}\{k\}$ is well-defined. If k is such that $\underline{W^k} \notin (0, \theta^{\rho^*})$ and for all i, $\underline{\Theta}_i^k \notin (0, \theta^{\rho^*})$, then it must be the case that for all i, $\underline{\Theta}_i^{k+1} \notin (0, \theta^{\rho^*})$. Since k^* is the first time in the game that a player selects a partition with $\underline{W^k} \in (0, \theta^{\rho^*})$, it must be the case that for all i, $\underline{\Theta}_i^k = 0$ or $\underline{\tilde{\Theta}}_i \geqslant \theta^{\rho^*}$. Since $\Theta^{k^*} \supset \Theta^K$, $\Theta^{k^*} \nsubseteq \breve{\Theta}$. Thus, $(\xi^{k^*}, \Theta^{k^*}, W^k)$ is such that $\Theta^{k^*} \nsubseteq \breve{\Theta}$, $\left(W^{k^*}, \cdot\right) \in \mu\left(\Theta^{k^*}, \xi^{k^*}\right)$, $\underline{W^{k^*}} \in \left(0, \theta^{\rho^*}\right)$, and for all i, $\underline{\Theta}_i^{k^*} = 0$ or $\underline{\Theta}_i^{k^*} \geqslant \theta^{\rho^*}$.

So, in order to have $\underline{W} \in (0, \theta^{\rho^*})$, it must be the case that $0 \in \Theta_\xi$. Thus it is possible for ξ to be a shill bidder while having so far only selected partitions that contain 0. Let $S = \left\{i : \overline{\Theta}_i < \theta^{\rho^*}\right\} \cup \{\xi\}$. By assumption that $\Theta \nsubseteq \breve{\Theta}$, there must exist i such that $\overline{\Theta}_i \geqslant \theta^{\rho^*}$ and thus we can suppose that $R \neq \varnothing$. By assumption that $\underline{W} \in \left(0, \theta^{\rho^*}\right)$, we can suppose that R is such that for all $i \in R$, $\overline{\Theta}_i > \underline{W}$. By Lemma C.6, this would contradict the hypothesis that the auction is weakly shill-proof and so we must have $\Theta \subseteq \breve{\Theta}$ when there exists $(W, \cdot) \in \mu(\Theta, \xi)$ such that $0 < \underline{W} < \theta^{\rho^*}$.

Proof of Condition ii of Definition 4.5. We now prove that for any player ξ and set of possible values $\Theta \subseteq \breve{\Theta}$, it is the case that $\mu(\Theta, \xi) = \mu^D_{\rho^*}(\Theta, \xi)$. Consider any option $(W, \cdot) \in \mu(\Theta, \xi)$. Observe that by Lemma C.6, it is not the case that $0 < \underline{W} < \overline{\Theta}_{-\xi}$. So, $\underline{W} \geqslant \overline{\Theta}_{-\xi}$. Since this is the case for all Θ, it must therefore be true that $\underline{W} = \overline{\Theta}_\xi$. This is because if $\overline{\Theta}_\xi > \underline{\Theta} \geqslant \overline{\Theta}_{-\xi}$, then there must have existed some earlier menu $\left(\tilde{W}, \cdot\right) \in \mu\left(\tilde{\Theta}, \tilde{\xi}\right)$ for which $\underline{\tilde{W}} < \overline{\tilde{\Theta}}_{\tilde{\xi}}$.

So far we have proven that $\mu(\Theta, \xi) = \left\{\left(W_{\mathsf{L}}, \tilde{\vec{\xi}}_{\mathsf{L}}\right), \left(W_{\mathsf{H}}, \tilde{\vec{\xi}}_{\mathsf{H}}\right)\right\}$. To complete the proof, we have to prove that $\tilde{\vec{\xi}}_{\mathsf{L}} = \vec{\xi}_{\mathsf{L}}$ and $\tilde{\vec{\xi}}_{\mathsf{H}} = \vec{\xi}_{\mathsf{H}}$. Towards contradiction, suppose $\tilde{\vec{\xi}}_{\mathsf{L}} \neq \vec{\xi}_{\mathsf{L}}$

or $\tilde{\vec{\xi}}_{\mathsf{H}} \neq \vec{\xi}_{\mathsf{H}}$. If $\tilde{\vec{\xi}}_{\mathsf{H}} \neq \vec{\xi}_{\mathsf{H}}$, then, by Lemma A.4, Condition (ii), there exists i such that $\overline{\Theta}_i = \overline{\Theta}_\xi$, $\big(\overline{\Theta}_i, i\big) \rhd \big(\overline{\Theta}_\xi, \xi\big)$. We can let $R = \{i\}$ and then apply Lemma C.6 to contradict the hypothesis that the auction is weakly shill-proof. If $\tilde{\vec{\xi}}_{\mathsf{L}} \neq \vec{\xi}_{\mathsf{L}}$, then, as argued in the proof of Theorem 3.6, the menu presented to $\tilde{\vec{\xi}}_{\mathsf{L}}$ must not have the auction end immediately, no matter what partition $\tilde{\vec{\xi}}_{\mathsf{L}}$ selects. Thus, our previous argument for the case where $\tilde{\vec{\xi}}_{\mathsf{H}} \neq \vec{\xi}_{\mathsf{H}}$ applies, and we can conclude that $\mu(\Theta, \xi) = \mu^D(\Theta, \xi)$. □

Proof of Proposition 4.6

The statement follows as a corollary of Lemma C.7. Consider any optimal reserve θ^{ρ^*}, $\underline{M}$ atoms below the optimal reserve, and $\overline{M}$ atoms above the optimal reserve. We construct a sparse (and regular) distribution $\tilde{F}$ with optimal reserve θ^{ρ^*}, $\underline{M}$ atoms below the optimal reserve, and $\overline{M}$ atoms above the optimal reserve. To begin, let δ such that $\underline{M}\delta \leqslant \theta^{\rho^*}$ and $(\underline{M}+1)\delta > \theta^{\rho^*}$. Then for all $k \leqslant \underline{M}$, let $\theta^{k+1} - \theta^k = \delta$ and $\tilde{f}(\theta^k) = e^{-\lambda(k-2)\delta} - e^{-\lambda(k-1)\delta}$. Note that

$$\tilde{\varphi}^k = \theta^k - (\theta^{k+1} - \theta^k)\frac{1 - \tilde{F}(\theta^k)}{\tilde{f}(\theta^k)} = (k-1)\delta - \delta\frac{e^{-\lambda(k-1)\delta}}{e^{-\lambda(k-2)\delta} - e^{-\lambda(k-1)\delta}},$$

and so $\tilde{\varphi}^{k+1} - \tilde{\varphi}^k = k\delta - (k-1)\delta = \delta > 0$; hence, $\tilde{F}$ satisfies the regularity condition for $k \leqslant \underline{M}$.

In order for θ^{ρ^*} to be an optimal reserve of $\tilde{F}$, it must be the case that for k^* such that $\tilde{\varphi}^{k^*} \geqslant 0$ and $\tilde{\varphi}^{k^*-1} < 0$, it is also the case that $\theta^{\rho^*} \in ((k^*-1)\delta, k^*\delta]$. Such a k^* must be equal to $\lceil \check{k} \rceil$, where $\check{k}\delta - \frac{\delta}{e^{\lambda\delta}-1} = 0$. Thus, $\underline{M} - 1 = k^* = \lceil \frac{1}{e^{\lambda\delta}-1} \rceil$.

In order for $\tilde{F}$ to be sparse, it must satisfy Equation (9), which here simplifies to

$$(k-1)\delta \cdot \left(1 + \frac{e^{-\lambda(k-2)\delta} - e^{-\lambda(k-1)\delta}}{e^{-\lambda(k-1)\delta} - e^{-\lambda(k)\delta}}\right) < k\delta \implies \frac{e^{2\lambda\delta} - 1}{e^{\lambda\delta} - 1} < \frac{k}{k-1}.$$

So, $\tilde{F}$ is sparse if

$$e^{2\lambda\delta} < \frac{2 + (e^{\lambda\delta} - 1)(2(k^* - \check{k}) - 1)}{1 - ((k^* - \check{k}) - 1)(e^{\lambda\delta} - 1)}. \tag{10}$$

Selecting λ such that $\underline{M} - 1 = \check{k} = k^*$, Equation (10) is satisfied.

Finally, to finish constructing $\tilde{F}$, we simply select atoms $\theta^{\underline{M}+2}, \ldots, \theta^{\underline{M}+\overline{M}}$ and respective probability weights to satisfy

$$\tilde{\varphi}^k \text{ is non-decreasing and } \sum_{k=\underline{M}+2}^{\underline{M}+\overline{M}} \tilde{f}(\theta^k) = e^{-\lambda(\underline{M}+1)\delta}.$$

This system of constraints has at most $\overline{M}$ constraints and $2(\overline{M}-1)$ free variables, so the system can be satisfied. Thus, we have constructed a regular and sparse $\tilde{F}$ that has the required values of θ^{ρ^*},$\underline{M}$, and $\overline{M}$. Then, we can apply Lemma C.7 to conclude the proof.

C.3 Strategy-Proof Auctions (Section 4.2.2) Appendix

Definition C.8. Let F be a discrete distribution with ordered atoms $0 = \theta^1 < \ldots < \theta^M$ and $\mathcal{F}$ be a continuous distribution with p.d.f. f. If $Y_{\mathcal{F}} \sim \mathcal{F}$, then F is a **discrete approximation** of $\mathcal{F}$ when $Y_F \sim F$ is defined as

$$Y_F = \begin{cases} \theta^1 & Y_{\mathcal{F}} \leqslant \theta^1 \\ \theta^k & Y_{\mathcal{F}} \in (\theta^{k-1}, \theta^k] \\ \theta^M & Y_{\mathcal{F}} > \theta^{M-1} \end{cases}. \tag{11}$$

For such a distribution F, let $\overline{\Delta} = \max_k \left\{\theta^k - \theta^{k-1}\right\}$. As convention, let F^{-1} be the left pseudo-inverse: $F^{-1}(x) = \max\left\{\theta^k : x \geqslant F(\theta^k)\right\}$.

Definition C.9. Let F be a discrete approximation of $\mathcal{F}$. The distribution F is a **monotone hazard rate (MHR) distribution** if $\frac{f(w^k)}{1-F(w^k)}$ is monotonically increasing in k and $h(x) = \frac{f(x)}{1-\mathcal{F}(x)}$ is monotonically increasing in x.

Lemma C.10. *If the value distribution is a discrete MHR distribution F, then for all*

$$\theta^Y \geqslant F^{-1}\left(F\left(\theta^{\rho^*}\right) + \max_{1 \leqslant n < N}\left\{\left(\max\left\{1 - \frac{\theta^{\rho^*}}{\theta^{\rho^*} + 2\overline{\Delta}}\left(\frac{f(\theta^{\rho^*})}{1 - F(\theta^{\rho^*})}\right)^n, 0\right\}\right)^{1/n}\right\}\right), \tag{12}$$

the ascending, screening auction with screening level θ^Y is a weakly shill-proof, ex-post incentive compatible, and optimal auction.

Proof of Lemma C.10.

The Ascending, Screening Auction is Orderly and Optimal. We first prove that the auction is well-defined, orderly, and optimal. The transfer and allocation function are orderly and optimal, so we only have to show that the menu rule can induce this outcome function. Let us examine the English auction phase first. The auction ends if and only if $\overline{\Theta} < \theta^M$. When that occurs, the auction has determined θ_i for all i given $\theta_i \geqslant \theta^{\rho^*}$. Thus, the outcome is fully determined. In the second-price auction phase, each value weakly greater than θ^Y is determined precisely (and there are at least two players with values weakly greater than θ^Y) and so the outcome rule is determined.

The Ascending, Screening Auction is Ex-Post Incentive Compatible. Observe that the definition of ex-post incentive compatiblility (Definition 4.1) is a function solely of the direct mechanism $(\tilde{x}, \tilde{t})$ and not of the menu rule μ. Both the English auction phase and the second-price auction use the same transfer function $\tilde{t}^2$. The entire auction has the optimal allocation rule $\tilde{x}^*$ and so the ascending, screening auction is ex-post incentive compatible.

The Ascending, Screening Auction is Weakly Shill-Proof. By assumption that the screen level is at least θ^{ρ^*}, any potential shill bidder will be asked to play at least once in the English auction before being able to play in the second-price auction. If the optimal shill bid is 0 in the first round of the English auction, then the auction is weakly shill-proof because once a bidder reports 0, she "drops out" and does not take another action.

Observe that given the form of the transfer rule, the maximum amount that a bidder i with value θ_i would have to pay is θ_i. Thus, the maximum possible gain in revenue from a shill bidder deviating is at most the difference between the first and second moment of F. Next, note that MHR distributions are regular. Regularity implies that if a shill bidder reports a non-zero value in the English auction stage and the auction concludes before reaching the second-price stage, the expected gain must be weakly less than 0. So, when considering the expected gain of misreporting, we can think of the expected gain from manipulating outcomes in the English auction component as at most 0 and can focus on manipulating outcomes in the second-price stage. Therefore, the total gains from misreporting as a shill bidder must be bounded above by the probability that a shill bidder is able to manipulate the outcome of the second-price auction multiplied by the expected difference between the first and second moments of the value distribution conditional on reaching the second-price auction stage.

Let $\mathcal{F}$ be the continuous distribution for which F is a discrete approximation. For an exponential distribution with rate λ, the expected difference between the first and second moments of T independent draws is $\frac{1}{\lambda}$. The exponential distribution, with its constant hazard rate, has the thickest right tail of any MHR distribution and so has the largest expected difference between its first and second moments (see proof of Theorem 5.1 in Bahrani et al. (2024)). In particular, since we are only interested in value draws above the reserve $\rho^*_{\mathcal{F}}$ and $\mathcal{F}$ has a non-decreasing hazard rate, we can take the rate $\lambda = h(\rho^*_{\mathcal{F}}) = \frac{1}{\rho^*_{\mathcal{F}}}$ and conclude that the maximum difference between the first and second moments of $\mathcal{F}$ must be bounded above by $\rho^*_{\mathcal{F}}$. Recall that $h(\rho^*_{\mathcal{F}}) = \frac{1}{\rho^*_{\mathcal{F}}}$ because $\mathcal{F}$ is regular and $\rho^*_{\mathcal{F}}$ is the optimal reserve of $\mathcal{F}$. Examining Equation (11), we can see that our discrete approximation pools draws from a continuous distribution upwards to atoms and so, if the absolute difference between two samples of the continuous distribution is κ, the absolute difference between the discrete approximation samples would be at most $\kappa + \overline{\Delta}$. Thus, the maximum possible expected difference between the first and second moments of F conditional on being above the reserve is at most $\rho^*_{\mathcal{F}} + \overline{\Delta}$. Note that we also know that $\left|\theta^{\rho^*} - \rho^*_{\mathcal{F}}\right| \leqslant \overline{\Delta}$.

Suppose bidder $i \in S$ and it is the first time she is taking an action. Then, under the rules of the auction, she has not indicated that her value is greater than θ^{ρ^*} yet. For any real bidder $j \neq i$, there are two cases: either bidder j has indicated her value is weakly greater than θ^{ρ^*} $\left(\mathbb{P}\left[\theta_j < \theta^Y\right] = F(\theta^Y) - F(\theta^{\rho^*})\right)$ or she has not yet taken an action $\left(\mathbb{P}\left[\theta_j < \theta^Y\right] = F(\theta^Y)\right)$. So, if $K \leqslant N$ real bidders have not dropped out yet (i.e., indicated that their value is less than θ^{ρ^*}), then the probability that the auction would continue to the second-price auction is at most $1 - \left(F\left(\theta^Y\right) - F\left(\theta^{\rho^*}\right)\right)^K$. Therefore, the maximum expected gain for a shill bidder from misreporting in her first action of the English auction phase when K bidders have not dropped is at most

$$\left(1 - \left(F\left(\theta^Y\right) - F\left(\theta^{\rho^*}\right)\right)^K\right)\left(\rho^*_{\mathcal{F}} + \overline{\Delta}\right). \tag{13}$$

We now turn to bounding the loss from reporting a non-zero value as a shill bidder. If shill bidder i misreports her value as θ^m at some point in the English auction phase and then she wins the item without taking another action, then the transfer the seller would have received had i not misreported is at least $\max\left\{\theta^{\rho^*}, \theta^{m-1}\right\} \geqslant \theta^{\rho^*}$, assuming at least one real bidder has value weakly above the reserve. To bound the probability that a real bidder j

would have won the item if not for shill bidder i's misreport, we can consider the probability that bidder j has indicated her value is at least $\underline{\Theta}_j \geqslant \theta^{\rho^*}$. By Definition C.9, the hazard rate of $\mathcal{F}$ is non-decreasing. So,

$$\mathbb{P}\left[\theta_j \leqslant \theta^m\right] \geqslant \frac{\sum_{\left\{k:\underline{\Theta}_j \leqslant \theta^k < \theta^m\right\}} f(\theta^k)}{1 - F(\underline{\Theta}_j)} \geqslant \frac{f(\underline{\Theta}_j)}{1 - F(\underline{\Theta}_j)} \geqslant \frac{f(\theta^{\rho^*})}{1 - F(\theta^{\rho^*})}.$$

Combining the preceding inequality with our hypothesis that K bidders have not dropped out yet, the expected loss for a shill bidder of misreporting is at least

$$\theta^{\rho^*} \cdot \left(\frac{f(\theta^{\rho^*})}{1 - F(\theta^{\rho^*})}\right)^K. \tag{14}$$

We conclude the proof by showing that θ^Y satisfying Equation (12) implies that the expected revenue loss from misreporting as a shill is weakly larger than the expected gain. Beginning with Equation (12), we can see that for all $K < N$,

$$\begin{aligned}
\theta^Y &\geqslant F^{-1}\left(F\left(\theta^{\rho^*}\right) + \max_{1 \leqslant n < N}\left\{\left(\max\left\{1 - \frac{\theta^{\rho^*}}{\theta^{\rho^*} + 2\overline{\Delta}}\left(\frac{f(\theta^{\rho^*})}{1 - F(\theta^{\rho^*})}\right)^n, 0\right\}\right)^{1/n}\right\}\right) \\
&\geqslant F^{-1}\left(F\left(\theta^{\rho^*}\right) + \left(\max\left\{1 - \frac{\theta^{\rho^*}}{\theta^{\rho^*} + 2\overline{\Delta}}\left(\frac{f(\theta^{\rho^*})}{1 - F(\theta^{\rho^*})}\right)^K, 0\right\}\right)^{1/K}\right) \\
&\geqslant F^{-1}\left(F\left(\theta^{\rho^*}\right) + \left(\max\left\{1 - \frac{\theta^{\rho^*}}{\rho^*_{\mathcal{F}} + \overline{\Delta}}\left(\frac{f(\theta^{\rho^*})}{1 - F(\theta^{\rho^*})}\right)^K, 0\right\}\right)^{1/K}\right).
\end{aligned}$$

This implies that

$$\theta^{\rho^*} \cdot \left(\frac{f(\theta^{\rho^*})}{1 - F(\theta^{\rho^*})}\right)^K \geqslant \left(1 - \left(F(\theta^Y) - F(\theta^{\rho^*})\right)^K\right)\left(\rho^*_{\mathcal{F}} + \overline{\Delta}\right).$$

The left-hand side of the preceding equation corresponds to Equation (14), the lower bound on the expected loss from misreporting as a shill bidder, and the right-hand side corresponds to Equation (13), the upper bound on the expected gain from misreporting and thus we have shown that it is equilibrium not to shill when θ^Y is sufficiently high. □

Proof of Proposition 4.12.

Let F_m be the discrete approximation of the exponential distribution with rate $\lambda = 1$ and atoms at $\{0, 2, \ldots, 2m\}$. Then, using the argument from the proof of Proposition 4.6, F_m is regular and MHR. For all $m > 2$, an optimal reserve is $\theta^{\rho^*} = 4$. Observe that $Y^* = 3$ satisfies Equation (12) because

$$\frac{\theta^{\rho^*}}{\theta^{\rho^*} + 2\overline{\Delta}}\left(\frac{f(\theta^{\rho^*})}{1 - F(\theta^{\rho^*})}\right)^n = \frac{1}{2}(e^2 - 1)^n > 1 \text{ for all } n \geqslant 1.$$

We apply Lemma C.10 to conclude that the ascending, screening auction with screen level θ^{Y^*} is weakly shill-proof, ex-post incentive compatible, and optimal for all F_m. Then, $Q^{AS,Y^*}(F_m) = 2$ and $Q^E(F_m) = m - 2$. Thus, $Q^{AS}(F_m, Y^*)/Q^E(F_m) \to 0$ as $m \to \infty$, concluding the proof. □

D Affiliation and Interdependence (Section 5) Appendix

Lemma D.1 (Karlin and Rinott (1980)). *For any non-decreasing function g, if $F' \geq_{\text{Aff}} F$, then $\mathbb{E}_{\theta\sim F'}[g(\theta)] \geqslant \mathbb{E}_{\theta\sim F}[g(\theta)]$.*

Lemma D.2. *For any function $g : \mathbb{R}^N \to \mathbb{R}$ with decreasing differences, if $F' \geq_{\text{Aff}} F$, then for any index i and for all $\theta_i^1 > \theta_i^2 > \theta_i^3$,*

$$\begin{aligned}&\mathbb{E}_{\theta\sim F'}\left[\left(g(\theta_i^1,\theta_{-i}) - g(\theta_i^2,\theta_{-i})\right) - \left(g(\theta_i^2,\theta_{-i}) - g(\theta_i^3,\theta_{-i})\right)\right]\\ &\quad\leqslant \mathbb{E}_{\theta\sim F}\left[\left(g(\theta_i^1,\theta_{-i}) - g(\theta_i^2,\theta_{-i})\right) - \left(g(\theta_i^2,\theta_{-i}) - g(\theta_i^3,\theta_{-i})\right)\right].\end{aligned}$$

Proof. Consider any i, Δ^1, and Δ^2. Then, define

$$\tilde{g}(\theta_i,\theta_{-i}) = \left(g(\theta_i+\Delta^1+\Delta^2,\theta_{-i}) - g(\theta_i+\Delta^1,\theta_{-i})\right) - \left(g(\theta_i+\Delta^1,\theta_{-i}) - g(\theta_i,\theta_{-i})\right).$$

By the assumption that g has decreasing differences, $\tilde{g}$ is non-increasing in θ_i. We can then immediately apply Lemma D.1 to complete the proof. □

Proof of Lemma 5.5

To prove that T has the claimed form, we will consider a specific non-truthful reporting: if a bidder has value θ^m, she commits to mis-reporting (selecting partitions) $\theta^{m'}$ for the rest of the game. Since our direct mechanism is an equilibrium for real bidders, we must have

$$\begin{aligned}&\mathbb{E}\Big[\Big(X_i(\theta^{j_m};\Theta)v(\theta^{j_m},\theta_{-i}) - T_i(\theta^{j_m};\Theta)\Big) -\\ &\qquad\qquad \left(X_i(\theta^{j_{m-1}};\Theta)v(\theta^{j_m},\theta_{-i}) - T_i(\theta^{j_{m-1}};\Theta)\right) \mid \theta_i = \theta^{j_m}, \theta_{-i}\in\Theta_{-i}\Big] \geqslant 0, \text{ and}\\ &\mathbb{E}\Big[\Big(X_i(\theta^{j_{m-1}};\Theta)v(\theta^{j_{m-1}},\theta_{-i}) - T_i(\theta^{j_{m-1}};\Theta)\Big) -\\ &\qquad\qquad \left(X_i(\theta^{j_m};\Theta)v(\theta^{j_{m-1}},\theta_{-i}) - T_i(\theta^{j_m};\Theta)\right) \mid \theta_i = \theta^{j_{m-1}}, \theta_{-i}\in\Theta_{-i}\Big] \geqslant 0.\end{aligned}$$

Defining U_i to be the ex-interim utility for bidder i, the preceding expressions become:

$$\begin{aligned}U_i(\theta^{j_m};\Theta) &\geqslant U_i(\theta^{j_{m-1}};\Theta) + (v_i^*(\theta^{j_{m-1}},\theta^{j_m};\Theta) - v_i^*(\theta^{j_{m-1}},\theta^{j_{m-1}};\Theta))X_i(\theta^{j_{m-1}};\Theta), \text{ and}\\ U_i(\theta^{j_{m-1}};\Theta) &\geqslant U_i(\theta^{j_m};\Theta) + (v_i^*(\theta^{j_m},\theta^{j_m};\Theta) - v_i^*(\theta^{j_m},\theta^{j_{m-1}};\Theta))X_i(\theta^{j_m};\Theta).\end{aligned}$$

Thus,

$$\begin{aligned}&(v_i^*(\theta^{j_{m-1}},\theta^{j_m};\Theta) - v_i^*(\theta^{j_{m-1}},\theta^{j_{m-1}};\Theta))X_i(\theta^{j_{m-1}};\Theta)\\ &\qquad\leqslant U_i(\theta^{j_m};\Theta) - U_i(\theta^{j_{m-1}};\Theta)\\ &\qquad\leqslant (v_i^*(\theta^{j_m},\theta^{j_m};\Theta) - v_i^*(\theta^{j_m},\theta^{j_{m-1}};\Theta))X_i(\theta^{j_m};\Theta).\end{aligned}$$

Then, by the same logic as in Lemma B.2, a mechanism with an optimal transfer rule is such that

$$U_i(\theta^{j_m};\Theta) = \sum_{k=2}^{m}(v_i^*(\theta^{j_{k-1}},\theta^{j_k};\Theta) - v_i^*(\theta^{j_{k-1}},\theta^{j_{k-1}};\Theta))X_i(\theta^{j_{k-1}};\Theta).$$

Therefore, we have that

$$T_i(\theta^{j_m};\Theta) = X_i(\theta^{j_m};\Theta)v_i^*(\theta^{j_m},\theta^{j_m};\Theta) - \sum_{k=2}^{m}(v_i^*(\theta^{j_{k-1}},\theta^{j_k};\Theta) - v_i^*(\theta^{j_{k-1}},\theta^{j_{k-1}};\Theta))X_i(\theta^{j_{k-1}};\Theta). \tag{15}$$

Re-arranging concludes the proof. □

Proof of Lemma 5.6

There are two cases to consider. In the first case, suppose $\max_{\theta\in S}\{\theta\} > \max_{\theta\in S'}\{\theta\}$. In this case, we can form subsets $S_1,\dots,S_K$ and $S'_1,\dots,S'_K$ such that

(i) $S = \bigcup_k S_k$ and $S' = \bigcup_k S'_k$;

(ii) For all k, $x \in S_k$, and $y \in S'_k$: $x > y$; and

(iii) For all $k' > k$, $x \in S_k$, and $y \in S'_{k'}$: $x < y$.

Such partitions can be formed inductively. First, define the base case as

$$S_K = \left\{\theta \in S : \theta > \max_{\tilde\theta\in S'}\left\{\tilde\theta\right\}\right\} \text{ and } S'_K = \left\{\theta \in S' : \theta > \max_{\tilde\theta\in S\setminus S_K}\left\{\tilde\theta\right\}\right\}.$$

Then, define the inductive case as

$$S_k = \left\{\theta \in S \setminus \left(\bigcup_{k'>k} S_{k'}\right) : \theta > \max_{\tilde\theta\in S'\setminus\left(\bigcup_{k'>k} S'_{k'}\right)}\left\{\tilde\theta\right\}\right\} \text{ and}$$

$$S'_k = \left\{\theta \in S' \setminus \left(\bigcup_{k'>k} S_{k'}\right) : \theta > \max_{\tilde\theta\in S\setminus\left(\bigcup_{k'\geqslant k} S_{k'}\right)}\left\{\tilde\theta\right\}\right\}.$$

Note that these two partitions have the same number of elements because $\min_{\theta\in S}\{\theta\} > \min_{\theta\in S'}\{\theta\}$. Then, we can re-write the right-hand side of Equation (1) as

$$\sum_{k=1}^{K}\left(\mathbb{P}_F\left[\theta_i \in S_k \mid S\right]\mathbb{E}_{\theta\sim F}\left[g(\theta) \mid \theta_i \in S_k\right] - \mathbb{P}_F\left[\theta_i \in S'_k \mid S'\right]\mathbb{E}_{\theta\sim F}\left[g(\theta) \mid \theta_i \in S'_k\right]\right).$$

Note that by assumption, the marginals are equal: $\mathbb{P}_{F'}\left[\theta_1 \in S_k\right] = \mathbb{P}_F\left[\theta_1 \in S_k\right]$ for all k. We can then recall that g is non-decreasing, weakly super-modular, and that affiliation is equivalent to log-supermodularity of the type distribution to conclude that for all k,

$$\begin{aligned}&\mathbb{P}_{F'}\left[\theta_i \in S_k \mid S\right]\mathbb{E}_{\theta\sim F'}\left[g(\theta) \mid \theta_i \in S_k\right] - \mathbb{P}_{F'}\left[\theta_i \in S'_k \mid S'\right]\mathbb{E}_{\theta\sim F'}\left[g(\theta) \mid \theta_i \in S'_k\right]\\ &\quad\geqslant \mathbb{P}_F\left[\theta_i \in S_k \mid S\right]\mathbb{E}_{\theta\sim F}\left[g(\theta) \mid \theta_i \in S_k\right] - \mathbb{P}_F\left[\theta_i \in S'_k \mid S'\right]\mathbb{E}_{\theta\sim F}\left[g(\theta) \mid \theta_i \in S'_k\right]\end{aligned}$$

Then, we can sum over all k to conclude the proof in this case.

Now, we consider the case where $\max_{\theta\in S}\{\theta\} < \max_{\theta\in S'}\{\theta\}$. In this case we can form subsets $S_1,\dots,S_K$ and $S'_1,\dots,S'_{K+1}$ such that

(i) $S = \bigcup_k S_k$ and $S' = \bigcup_k S'_k$;

(ii) For all k, $x \in S_k$, and $y \in S'_k$: $x > y$; and

(iii) For all $k' > k$, $x \in S_k$, and $y \in S'_{k'}$: $x < y$.

The inductive construction is the same as the first case with the base case changed to

$$S'_{K+1} = \left\{\theta \in S' : \theta > \max_{\tilde{\theta} \in S} \left\{\tilde{\theta}\right\}\right\} \text{ and } S_K = \left\{\theta \in S : \theta > \max_{\tilde{\theta} \in S' \setminus S'_{K+1}} \left\{\tilde{\theta}\right\}\right\}.$$

We then re-write the right-hand side of Equation (1) as

$$\begin{aligned}
&\frac{1}{2}\mathbb{P}_F\left[\theta_i \in S_1\right]\mathbb{E}_{\theta\sim F}\left[g(\theta) \mid \theta_i \in S_1\right] - \mathbb{P}_F\left[\theta_i \in S'_1\right]\mathbb{E}_{\theta\sim F}\left[g(\theta) \mid \theta_i \in S'_1\right] + \\
&\frac{1}{2}\mathbb{P}_F\left[\theta_i \in S_K\right]\mathbb{E}_{\theta\sim F}\left[g(\theta) \mid \theta_i \in S_K\right] - \mathbb{P}_F\left[\theta_i \in S'_{K+1}\right]\mathbb{E}_{\theta\sim F}\left[g(\theta) \mid \theta_i \in S'_{K+1}\right] + \\
&\quad\frac{1}{2}\sum_{k=2}^{K}\Bigg(\left(\mathbb{P}_F\left[\theta_i \in S_k\right]\mathbb{E}_{\theta\sim F}\left[g(\theta) \mid \theta_i \in S_k\right] - \mathbb{P}_F\left[\theta_i \in S'_k\right]\mathbb{E}_{\theta\sim F}\left[g(\theta) \mid \theta_i \in S'_k\right]\right) \\
&\qquad\qquad - \left(\mathbb{P}_F\left[\theta_i \in S'_k\right]\mathbb{E}_{\theta\sim F}\left[g(\theta) \mid \theta_i \in S'_k\right] - \mathbb{P}_F\left[\theta_i \in S_{k-1}\right]\mathbb{E}_{\theta\sim F}\left[g(\theta) \mid \theta_i \in S_{k-1}\right]\right)\Bigg)
\end{aligned}$$

Now, observe that the first two differences must decrease under F' because g is non-decreasing and weakly super-modular. We can then recall that g has decreasing differences, apply Lemma D.2, and then sum over all k to conclude that the sum also must decrease and therefore obtain the desired conclusion. □

Proof of Theorem 5.4

Affiliation. Towards contrapositive, suppose that the auction is not shill-proof (strong or weak, we will highlight where the proofs diverge) under F and consider some $F' \succeq_{\text{Aff}} F$. There exists some realization of θ under F such that a shill bidder i deviates at Θ. Since $\tilde{x}$, μ, and ξ_0 are all fixed, we know from Lemma 5.5 that under weak shill-proofness that the only way for expected revenue under weak shill-proofness to change is for the ex-interim expected value of some bidder $j \neq i$ to change. We can also apply Lemma B.3 to see that the same is true under strong shill-proofness. Thus, in order to prove that the auction is not shill-proof for either weak or strong shill-proofness under F', it is sufficient to show that for all $j \neq i$, if the expected transfers from T_j from manipulation $\hat{\Theta}$ compared to truthful play Θ^* under F, then the same manipulation is profitable under F'. Fix j and consider the profitably manipulated possible values $\hat{\Theta}$. We can re-express the expected transfers from

bidder j from Lemma 5.5 as

$$\begin{aligned}
&\mathbb{E}_{\theta\sim F}\left[T_j(\theta_j;\Theta^*) \mid \theta\in\Theta^*\right]\\
&\quad=\mathbb{E}_{\theta\sim F}\Bigg[\mathbb{E}\left[\tilde{x}_j(\theta_j,\theta_{-j})v(\theta_j,\theta_{-j}) \mid \theta_j=\theta_j,\tilde{x}_j(\theta_j,\theta_{-j})=1\right]\\
&\qquad-\sum_{m:\theta^{k_m}<\theta_j}\mathbb{P}\left[\tilde{x}_j(\theta^{k_m},\theta_{-j})=1\right]\Big(\mathbb{E}\left[v(\theta^{k_m},\theta_{-j}) \mid \theta_j=\theta^{k_{m+1}},\tilde{x}_j(\theta^{k_m},\theta_{-j})=1\right]\\
&\qquad\qquad-\mathbb{E}\left[v(\theta^{k_m},\theta_{-j}) \mid \theta_j=\theta^{k_m},\tilde{x}_j(\theta^{k_m},\theta_{-j})=1\right]\Big) \mid \theta_{-j}\in\Theta^*_{-j},\theta_j\in\Theta^*_j\Bigg].
\end{aligned}$$

Now observe that with manipulation only changes that the expectation is taken with respect to $\theta_{-j}\in\hat{\Theta}_{-j}$. Note that we are not worrying about the true distribution from bidder $j'\neq j$ because when considering revenue from bidder j, we consider her expected transfers, and then integrate over j's true expected distribution, $\theta_j\in\Theta^*_j$. Given that v satisfies Assumption 5.1 and $\tilde{x}$ is orderly, we know that

$$\begin{aligned}
&\mathbb{E}\left[\tilde{x}_j(\theta_j,\theta_{-j})v(\theta_j,\theta_{-j}) \mid \theta_j=\theta_j,\tilde{x}_j(\theta_j,\theta_{-j})=1\right]\\
&\quad-C_m\cdot\Big(\mathbb{E}\left[v(\theta^{k_m},\theta_{-j}) \mid \theta_j=\theta^{k_{m+1}},\tilde{x}_j(\theta^{k_m},\theta_{-j})=1\right]-\mathbb{E}\left[v(\theta^{k_m},\theta_{-j}) \mid \theta_j=\theta^{k_m},\tilde{x}_j(\theta^{k_m},\theta_{-j})=1\right]\Big)
\end{aligned}\tag{16}$$

also has the same properties (with respect to θ_j) for any m such that $\theta^{k_m}<\theta_j$ and constant C_m. Therefore, we can conclude that the necessary conditions for Lemma 5.6 are satisfied and apply it to conclude that

$$\mathbb{E}_{\theta\sim F'}\left[T_j(\theta_j;\hat{\Theta})-T_j(\theta_j;\Theta^*) \mid \theta\in\Theta^*\right]\geqslant\mathbb{E}_{\theta\sim F}\left[T_j(\theta_j;\hat{\Theta})-T_j(\theta_j;\Theta^*) \mid \theta\in\Theta^*\right].$$

Thus, the auction is not shill-proof under F'.

Common Values. Towards contrapositive, suppose that the auction is not shill-proof (strong or weak, we will highlight where the proofs diverge) under v and consider some $v'\geq_{\text{Com}}v$. There exists some realization of v under v such that a shill bidder i deviates at Θ. Since $\tilde{x}$, μ, and ξ_0 are all fixed, we know from Lemma 5.5 that under weak shill-proofness that the only way for expected revenue under weak shill-proofness to change is for the ex-interim expected value of some bidder $j\neq i$ to change. We can also apply Lemma B.3 to see that the same is true under strong shill-proofness. Thus, in order to prove that the auction is not shill-proof for either weak or strong shill-proofness under F', it is sufficient to show that for all $j\neq i$, if the expected transfers from T_j from manipulation $\hat{\Theta}$ compared to truthful play Θ^* under v, then the same manipulation is profitable under v'. This is immediate: since v' is point-wise higher than v and more super-modular, then we can see that Equation (16) is larger under v' than v no matter the constant and therefore the auction is not shill-proof.

Proof of Proposition 5.7

The proof that a strongly shill-proof auction must be pay-as-bid is the same as in Lemma 3.1 because that argument is ex-post and bidders' values are weakly increasing in all types. To

prove that the conjectured efficient allocation rule is efficient, observe that by the assumption that a bidder values her own signal more than that of others, the efficient ex-post allocation is such that $\tilde{x}_i = 0$ for $i \notin \text{argmax}_j \{\theta_j\}$ and $\sum_j \tilde{x}_j(\theta) = 1$ for all θ. The assumption that the seller has 0 value for the good completes the proof that $\tilde{x}^E$ is the efficient allocation rule.

Recall that a shill-proof auction must be pay-as-bid, and observe that bids must be strictly increasing as a function of own signal in order to be incentive compatible. Next, because the bidder who has the highest ex-ante signal will have the highest ex-post valuation, under any allocation rule, the bidder with the highest signal who is allocated will have the highest bid. Therefore, to maximize transfers the optimal allocation rule is such that $\tilde{x}_i = 0$ for $i \notin \text{argmax}_j \{\theta_j\}$. To conclude the proof, observe that by Lemma 5.5, the fact that changing the allocation rule conditional on only the maximum value will not change that bidder's value, and by the standard arguments, the optimal allocation rule uses a reserve type.

E Supplemental Appendix

The following definition for an extensive-form auction is taken[31] from Li (2017):

Definition E.1. An **extensive form auction** G is defined as the tuple $(H, \prec, A, \mathcal{A}, P, \{\mathcal{I}_i\}_{i \in B}, (x, t))$ such that:

(i) H is a set of histories, along with a binary relation $\prec$ on H that represents precedence. In addition:

  (a) $\prec$ forms a partial order and $(H, \prec)$ forms an arborescence.

  (b) There exists an initial history $h_\varnothing$ such that there does not exist h' where $h' \prec h_\varnothing$.

  (c) The set of terminal histories is $Z \equiv \{h : \neg\exists h \text{ such that } h \prec h'\}$.

  (d) The set of immediate successors to h is $\text{succ}(h)$.

(ii) A is the set of possible actions.

(iii) $\mathcal{A} : H \backslash h_\varnothing \to A$ maps histories to the most recent action taken to reach it. In addition:

  (a) For all h, $\mathcal{A}(h)$ is one-to-one on $\text{succ}(h)$.

  (b) The set of actions available at h is

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

(iv) $P : H \backslash Z \to B$ is the player function for any given non-terminal history.

(v) $\mathcal{I}_i$ is a partition of $\{h : P(h) = i\}$ such that:

  (a) $A(h) = A(h')$ when h and h' are in the same cell of the partition, and

  (b) $A(h) \cap A(h') = \varnothing$ when h and h' are not in the same cell of the partition.

(vi) For every $z \in Z$, $z = (x, t)$, such that $\sum_{i=1}^{N} x_i \leqslant 1$, $x_i \in [0, 1]$, and $t_i \in \mathbb{R}$.

In order to define an information order, we will use the notation that when the current set of possible values is Θ, a player i's knowledge of what values are possible is $\mathcal{I}^i_\Theta$.

Definition E.2. A menu rule μ' is **more informative** than μ, $\mu' \succeq_{\text{Info}} \mu$, if $\mu(\cdot, \cdot) = \mu'(\cdot, \cdot)$ for all game states Θ and all possible information sets $\mathcal{I}, \mathcal{I}'$, when $\mathcal{I} \subseteq \mathcal{I}'$, it is the case that

$$\mathbb{P}\left[\mathcal{I}_\Theta = \mathcal{I} \mid \mu'\right] - \mathbb{P}\left[\mathcal{I}_\Theta = \mathcal{I}' \mid \mu'\right] \geqslant \mathbb{P}\left[\mathcal{I}_\Theta = \mathcal{I} \mid \mu\right] - \mathbb{P}\left[\mathcal{I}_\Theta = \mathcal{I}' \mid \mu\right].$$

Proposition E.3. *Consider any affiliated type distribution F and value function v satisfying Assumption 5.1. Suppose $(\tilde{x}, \tilde{t}^*(\tilde{x}, \mu, \xi_0, v, F), \mu, \xi_0)$ is orderly and strongly shill-proof. Then, for any μ' such that $\mu \succeq_{\text{Info}} \mu'$, it is the case that $(\tilde{x}, \tilde{t}^*(\tilde{x}, \mu', \xi_0, v, F), \mu', \xi_0)$ is strongly shill-proof.*

[31] We modify the definition to remove notation we do not use and to make it specific to auctions.

Proof. Towards contrapositive, suppose that the auction is not strongly shill-proof under μ and consider some $\mu' \geq_{\text{Info}} \mu$. There exists some realization of θ under F such that a shill bidder i deviates at Θ. Since F is affiliated and $\tilde{x}$, ξ_0, and μ are all fixed, we can apply Lemmatta 5.5, B.3 and D.1 to conclude that the only way for expected revenue under strong shill-proofness to change is for ex-interim expected values for the winner $j \neq i$ to change. There can be no change in allocation since j has the highest value among the real bidders and so since the allocation is orderly, any change in the allocation would result in a shill bidder winning the item. Recall that the winning bidder conditions her valuation on the realization that she wins, i.e., she conditions on the assumption that $(\theta_j, j) \rhd (\theta_k, k)$ for all $k \neq j$. Thus, it is without loss with regards to how the winning bidder j will estimate her valuation to assume that $\mathcal{I}_\Theta$ is such that for all $y \in \mathcal{I}_\Theta$ and $k \neq j$, $(\theta_j, j) \rhd (y_k, k)$. There are now two cases to consider. In the first case, consider that under μ', the information sets in the support are a super-set of those in the support under μ. We then condition on the event of the same realization of the information set that leads to a profitable deviation to find that the auction is not strongly shill-proof under μ'. In the second case, we do not assume that the support under μ contains elements that μ' does not. It may be that the realization of the information set for the winner's last move is not in the support of μ. However, because $\mu' \geq_{\text{Info}} \mu$, we can know that there is some realization of the information structure such that 0 is not contained in it. This is because we know that there exists an information set that is strictly smaller than that realization and by Bayes plausibility, the complement must also be in it. Further, a shill bidding deviation means selecting a partition which does not include 0 and by Lemma A.4 Condition 2a, we know that said partition does not intersect the partition containing 0. Since the allocation is fixed, one of the partitions must increase transfers from the winner in order to have the average transfer under every realization of the information set average out to the transfer under the coarser information set. We then select the highest transfer among these realizations to find a profitable deviation and complete the proof. □

Example E.4. Consider the sealed-bid, first-price (pay-as-bid) optimal auction with equilibrium bids b^1 in the regular, IPV environment. The naïve implementation of allocation and transfer rule $(\tilde{x}^*, \tilde{t}^1)$ in a public auction would be to query each bidder sequentially on what her value is and then have the payment rule be $\tilde{t}^1$, but $\tilde{t}^1$ is not the direct transfer rule of any equilibrium of this game. Indeed, consider the last bidder who takes a move, and label that bidder N. If $\theta_N > \max_{i<N}\{b_i\}$, then the only possible equilibrium bid—and therefore the transfer function—is $\max_{i<N}\{b_i\}$.

If we modify the direct transfer rule to represent the bid that each bidder submits in equilibrium in this sequential form (as one could solve for inductively), the auction would be weakly shill-proof by regularity. In particular, while a shill bid can force later bidders to pay a higher price, the probability that no one will want to pay that higher price outweighs the benefit by regularity. However, such an auction is not strongly shill-proof because, given full knowledge of real bidders' valuations, a shill bidder will be incentivized to bid just below the highest valuation of a subsequent bidder.

Example E.5. Let $\mathcal{F}(x) = 1 - e^{-0.1x}$ and let F_1, F_2 be discrete approximations of $\mathcal{F}$ with atoms $\{0, 5, 9, 14, 20\}$ and $\{0, 3, 7, 14, 20\}$, respectively. It can be verified that both these distributions are regular and have optimal reserve $\theta^{\rho^*} = 14$. Consider a variant of the

efficient Dutch auction, with the modification that, if all bidders have indicated values less than 20, then the auction queries bidders from lowest-to highest-priority as to whether their value is at least 9. If no one indicates that their value is at least 9, then the Dutch auction continues. If at least one person does indicate that their value is at least 9, then bidders are queried from lowest- to highest-priority as to whether their value is 14, and the transfer is 14 if at least two people have value 14, and 9 if only one person does. It can be verified that if the value distribution is F_1, the auction just described is weakly shill-proof, but if the value distribution is F_2, then the auction is not weakly shill-proof. When the value distribution is F_2, in expectation, a shill bidder will want to report that her value is 9. In fact, Lemma C.7 (see appendix) implies that if the value distribution is F_2, then the auction in this example must be a semi-Dutch auction with cutoff at least 14.

E.1 Credibility

The following definitions are adapted from Akbarpour and Li (2020) to match our notation and specialized to the auction setting:

For any extensive form game G, we can define a messaging game as follows:

1. The auctioneer chooses to:
   (a) Select an outcome and end the game; or
   (b) Go to step 2.
2. The auctioneer chooses some bidder $i \in B$ and sends a message $m = I_i \in \mathcal{I}_i$.
3. Bidder i privately observes message $m = I_i$ and chooses reply $r \in A(I_i)$.
4. The auctioneer privately observes r.
5. Go to step 1.

We can now write bidder i's observations in the game as $((m_i^k, r_i^k)_{k=1}^{\tau_i}, \omega_i)$ where τ_i is the number of observations i has and ω_i is the information partition over outcomes that i observes. Let $o_i(\sigma_0, \sigma, \theta)$ be i's observation when the auctioneer plays σ_0, the bidders play σ, and the type profile is θ.

Definition E.6 (Akbarpour and Li (2020))**.** Let σ_0^G be the **rule-following auctioneer strategy**. Formally, σ_0^G is defined by the following algorithm: Initialize $\hat{h} := h_\varnothing$. At each step, if $\hat{h} \in z$, end the game and choose outcome $(x, t)(\hat{h})$. Else, contact agent $P(\hat{h})$ and send message $m = \mathcal{I}_{P(\hat{h})}$ such that $P(\hat{h}) \in \mathcal{I}_{P(\hat{h})}$. Upon receiving reply r, update $\hat{h}$ to h such that $h \in \text{succ}(\hat{h})$ and $\mathcal{A}(h) = r$, then iterate.

Definition E.7 (Akbarpour and Li (2020), Definition 3)**.** Given some promised strategy profile (σ_0, σ), auctioneer strategy $\hat{\sigma}_0$ is safe if, for all agents $i \in B$ and all type profiles θ, there exists θ'_{-i} such that $o_i(\hat{\sigma}_0, \sigma, \theta) = o_i(\sigma_0, \sigma, (\theta_i, \theta'_{-i}))$. We denote by $\Sigma_0^*(\sigma_0, \sigma)$ the set of **safe strategies**.

Definition E.8 (Akbarpour and Li (2020), Definition 4)**.** (G, σ) is **credible** if

$$\sigma_0^G \in \underset{\sigma_0 \in \Sigma_0^*(\sigma_0^G, \sigma)}{\operatorname{argmax}} \left\{ \mathbb{E}_\theta \left[\sum_{i \in B} t_i(\sigma_0, \sigma, \theta) \right] \right\}$$

Proof of Proposition 6.1

Strong Shill-Proofness → Credibility. We prove the contrapositive: Suppose (G, σ) is not credible. Let $\hat{\sigma}_0 \in \Sigma_0^*(\sigma_0^G, \sigma)$ be a profitable and safe deviation by the auctioneer. By Definition E.7, there exists θ and $\{\theta'_{-i}\}_i$ such that $o_i(\hat{\sigma}_0, \sigma, \theta) = o_i(\sigma_0, \sigma, (\theta_i, \theta'_{-i}))$ for all i. By Lemma B.1, only one bidder i^* has $t_{i^*}(\sigma(\theta_{i^*}, \theta'_{-i^*})) \neq 0$ and so $t_{i^*}(\sigma(\theta_{i^*}, \theta'_{-i^*})) > t_{i^*}(\sigma(\theta))$ because $\hat{\sigma}_0$ is profitable. Then, let $R = \{i^*\}$, and by ex-post monotonicity,

$$t_{i^*}(\sigma(\theta_{i^*}, \theta'_{-i^*})) > t_{i^*}(\sigma(\theta)) \geqslant t_{i^*}(\sigma(\theta_{i^*}, 0)),$$

and so the auction is not strongly shill-proof.

Credibility → Weak Shill-Proofness. We prove the contrapositive: Suppose (G, σ) is not weakly shill-proof. Let $\hat{\sigma} \in \Sigma_S$ be a profitable shilling strategy. Then, by Definition A.1, for all R, there exists θ, θ' such that $\hat{\sigma}(\theta; R) = \sigma(\theta; B)$. Consider the following reporting strategy $\hat{\sigma}_0$: for all $i \in R$, report play as if $i \in R$ is following $\hat{\sigma}$; and for all $i \notin R$, report in the rule-following manner. This strategy is safe because $\hat{\sigma} \in \Sigma_S$. To see that it is profitable compared to σ_0^G, consider what happens when the winning bidder i is or is not in R.[32] Conditional on $i \in R$ winning, $\hat{\sigma}_0$ increases expected revenue because $\hat{\sigma}$ is a profitable shill bidding strategy. Conditional on $i \notin R$ winning, shill bidding would have led 0 revenue for the seller and, by Lemma B.1, $\hat{\sigma}_0$ must have non-negative revenue. Thus, our described reporting strategy is a profitable, safe strategy and therefore the auction is not credible. □

E.2 Generalizing Credibility in the Single-Action Case

Definition E.9. Fix a single-action auction with exogenous signal ψ and a set of real bidders R. The set of safe deviations to report to $i \in B$ is

$$\mathcal{A}_i^\psi(\theta_{j \leqslant i}) = \left\{ a : \exists \tilde{\theta}_{-i} \text{ such that } \left[j < i, \theta_j \neq 0 \implies \tilde{\theta}_j = \theta_j \right] \text{ and } a = \left(\sigma_i(\theta_i; \psi_i(\tilde{\theta}_{j<i})), \sigma_{-i}\left(\theta_i, \tilde{\theta}_{-i}; B\right) \right) \right\}.$$

The total set of **safe deviations** is

$$\mathcal{A}^\psi(\theta) = \left\{ \left\{ a^{\rightsquigarrow i} \right\} : a^{\rightsquigarrow i} \in \mathcal{A}_i^\psi(\theta_{j \leqslant i}) \text{ and } \sum_{i=1}^N x_i\left(a^{\rightsquigarrow i}\right) \leqslant 1 \right\}.$$

Definition E.9 allows the auctioneer to report any value she chooses when a bidder's declared valuation is 0. Note that if we take the canonical setting where there are no exogenous signals, the above assumption is without loss.

[32] Note that if no one has value above the optimal reserve, there will be no winner under any safe strategy, so let us only consider the case where the good is allocated.

Definition E.10. A single-action auction is ψ**-credible** if for all θ and $\{a^{\rightsquigarrow i}\} \in \mathcal{A}^{\psi}(\theta)$, we have

$$\sum_i t_i(a^{\rightsquigarrow i}) \leqslant \sum_i t_i(\sigma(\theta; B)).$$

Lemma E.11. *For a single-action auction, define the augmented (direct) inverse* $\breve{\psi}_i^{-1}$ *as* $\breve{\psi}_i^{-1}(\theta) = \{0\} \cup \psi_i^{-1}(\theta_{j<i})$. *Then for*

$$\vartheta^{\psi}(\theta) = \left\{ \left\{\theta^{\rightsquigarrow i}\right\} : \theta^{\rightsquigarrow i} \in \breve{\psi}_i^{-1}(\theta), \sum_i \tilde{x}_i(\theta^{\rightsquigarrow i}) \leqslant 1 \right\},$$

the auction is credible if and only if for all θ, *and* $\{\theta^{\rightsquigarrow i}\} \in \vartheta^{\psi}(\theta)$,

$$\sum_i \tilde{t}_i(\theta^{\rightsquigarrow i}) \leqslant \sum_i \tilde{t}_i(\theta).$$

Proof. Apply Lemma C.1, specifically the unique mapping between $(\tilde{x}, \tilde{t})$ and (x, t) to Definitions E.8 and E.9 to see that the lemma holds. □

Lemma E.12. *Suppose a single-action auction is weakly shill-proof, but not strongly shill-proof. Then, there exist* R, θ_R, *and* θ_{-R} *such that*

$$\sum_{k \in R} \tilde{t}_k(\theta_R, \theta_{-R}) > \sum_{k \in R} \tilde{t}_k(\theta_R, 0). \tag{17}$$

Proof. Suppose that (G, σ) is weakly shill-proof, but not strongly shill-proof. Because (G, σ) is weakly shill-proof, for all θ and R, R', we can define $\hat{\sigma}(\theta) \equiv \sigma(\theta; R) = \sigma(\theta; R')$. Since (G, σ) is not strongly shill-proof, $\hat{\sigma}$ must not be an ex-post strategy for the shill bidders. So, for some realization of R and θ_R there exists a profitable deviation for the shill bidders; examining the set of possible deviations Σ_S in Definition A.1, we see that any profitable deviating actions induces a profitable misreport θ_{-R} in the direct mechanism for some R and θ_R; proving Equation (17) can be satisfied. □

Lemma E.13. *If a single-action auction is strongly shill-proof, then for all* R, $i \notin R$, θ_i, *and* θ_{-i}, $\sum_{k \in R} \tilde{t}_k(\theta_i, \theta_{-i}) \leqslant \sum_{k \in R} \tilde{t}_k(0, \theta_{-i})$.

Proof. Towards contradiction, suppose that there exists R, $i \notin R$, θ_i, and θ_{-i} such that $\sum_{k \in R} \tilde{t}_k(\theta_i, \theta_{-i}) > \sum_{k \in R} \tilde{t}_k(0, \theta_{-i})$. So in the direct game reporting 0 is not a dominant strategy for shill bidders. This implies, from Condition i of Lemma C.1, that there exists a deviation in the auction such that for some value vectors, the seller raises more revenue. Therefore, the auction is not strongly shill-proof. □

Proof of Proposition 6.2

Weak Shill-Proofness $\rightarrow$ **(**$\psi = \mathrm{Id}$**)-Credibility.** Suppose the auction is not ($\psi = \mathrm{Id}$)-credible. Then, combining Lemma E.11 with the ex-post IR condition, there exist $\theta, \{\theta^{\rightsquigarrow i}\} \in \vartheta^{\psi}(\theta)$ and k^* such that $\tilde{t}_{k*}(\theta^{\rightsquigarrow k^*}) > \tilde{t}_{k*}(\theta)$. Applying the definition of orderly and the

winner-paying property, for all $j \neq k^*$, $\tilde{t}_j(\theta^{\rightsquigarrow k^*}) = 0$. Since $\psi = \mathrm{Id}$, for all $j \leqslant k$, it is the case that $\theta_j^{\rightsquigarrow k^*} = \theta_j$ or $\theta_j = 0$. Let $R = \{1, \ldots, k^*\}$. Then,

$$\begin{aligned}\sum_{i \in R} \tilde{t}_i(\theta^{\rightsquigarrow k^*}) &= \tilde{t}_{k^*}(\theta^{\rightsquigarrow k^*}) \\ &\geqslant \tilde{t}_{k^*}\left(\theta_1, \ldots, \theta_{k^*}, \theta_{k^*+1}^{\rightsquigarrow k^*}, \ldots, \theta_N^{\rightsquigarrow k^*}\right) \\ &\geqslant \tilde{t}_{k^*}\left(\theta_1, \ldots, \theta_{k^*}, 0, \ldots, 0\right) \\ &= \sum_{i \in R} \tilde{t}_i(\theta_R, 0).\end{aligned}$$

Thus, we can apply Lemma C.2 to conclude that the auction is not weakly shill-proof.

ψ-Credibility → Weak Shill-Proofness. Suppose the auction is ψ-credible. Towards contradiction, suppose the auction is not weakly shill-proof. So, there exists R and θ such that $\sigma(\theta; R) \neq \sigma(\theta; B)$. In particular, this means that shill bidders have, in expectation, a profitable deviation relative to acting as real bidders with valuation 0. If this is true in expectation, there must then exist $\theta = (\theta_R, 0)$ and $\tilde{\theta}_{-R}$ such that

$$\sum_{i \in R} \tilde{t}_i((\theta_R, \tilde{\theta}_{-R})) > \sum_{i \in R} \tilde{t}_i((\theta_R, 0)).$$

Now, let us consider the messaging deviation

$$\left\{\theta^{\rightsquigarrow i}\right\}_{i \in B} = \begin{cases} \left(\theta_R, \tilde{\theta}_{-R}\right) & i \in R \\ (\theta_R, 0) & \text{otherwise} \end{cases}.$$

By the definition of credibility, the auctioneer can report any value to other bidders when the value reported to him is 0 and bidders with value 0 are told the other bidders' true reports. Therefore, $\{\theta^{\rightsquigarrow i}\} \in \vartheta^{\psi}(\theta_R, 0)$ and

$$\sum_i \tilde{t}_i(\theta^{\rightsquigarrow i}) = \sum_{i \in R} \tilde{t}_i(\theta_R, \tilde{\theta}_{-R}) + \sum_{i \notin R} \tilde{t}_i(\theta_R, 0) > \sum_i \tilde{t}_i((\theta_R, 0)).$$

This contradicts Lemma E.11, and so the auction must be weakly shill-proof.

($\psi = \varnothing$)-Credibility → Strong Shill-Proofness. Suppose that the auction is not strongly shill-proof and $\psi = \varnothing$. There are two cases to consider: either the auction is not weakly shill-proof or it is. If the auction is not weakly shill-proof, then we can apply the previous case to conclude the auction is not ($\psi = \varnothing$)-credible. If the auction is weakly shill-proof, but not strongly shill-proof, then by Lemma E.12, there exists $R, k^* \in R$, and θ such that $\tilde{t}_{k^*}(\theta) > \tilde{t}_{k^*}(\theta_R, 0)$. Thus, we can construct the following profitable auctioneer reporting deviation:

$$\left\{\theta^{\rightsquigarrow i}\right\}_{i \in B} = \begin{cases} (\theta_R, \theta_{-R}) & i = k^* \\ (\theta_R, 0) & \text{otherwise} \end{cases}.$$

Since $\psi = \varnothing$, we know that $\left\{\tilde{\theta}^{\rightsquigarrow i}\right\} \in \vartheta^{\psi}(\theta_R, 0)$. The total transfers are then

$$\sum_i \tilde{t}_i\left(\theta^{\rightsquigarrow i}\right) = \tilde{t}_{k^*}(\theta) + \sum_{i \neq k^*} \tilde{t}_i\left(\theta_R, 0\right) > \sum_i \tilde{t}_i\left(\theta_R, 0\right).$$

Hence, by Lemma E.11, we see that the auction is not credible.

Strong Shill-Proofness → ψ-Credibility. Suppose that the auction is not ψ-credible. Then, combining Lemma E.11 with the ex-post IR condition, there exists $\theta, \{\theta^{\rightsquigarrow i}\} \in \vartheta^{\psi}(\theta)$ and k^* such that $\tilde{t}_{k^*}(\theta^{\rightsquigarrow k^*}) > \tilde{t}_{k^*}(\theta)$. Recall, by the definition of ψ^{-1}, that $\theta_{k^*}^{\rightsquigarrow k^*} = \theta_{k^*}$. Suppose $R = \{k^*\}$. Then,

$$\sum_{i \in R} \tilde{t}_i(\theta^{\rightsquigarrow k^*}) = \tilde{t}_{k^*}(\theta^{\rightsquigarrow k^*}) > \tilde{t}_{k^*}(\theta) \geqslant \tilde{t}_{k^*}(\theta_{k^*}, 0) = \sum_{i \in R} \tilde{t}_{k^*}(\theta_{k^*}, 0).$$

Therefore, by Lemma E.13, the auction is not strongly shill-proof. □