

AS-IF DOMINANT STRATEGY MECHANISMS

Lea Nagel*

Roberto Saitto*

December 30, 2025

[\[Click here for latest version\]](#)

Abstract

We show that achieving dominant strategy incentive compatibility often requires to choose a mechanism which severely limits what agents can observe about others' previous moves.

However, experiments and theoretical arguments suggest increasing the transparency of a mechanism's extensive form can improve reliability of its predictions—even if it breaks the dominant strategy property.

To help resolve this dilemma, we define *as-if dominant strategy mechanisms*: (i) Each agent has at least one strategy that becomes dominant if the others were restricted to behave *as if* the mechanism was static, and (ii) all combinations of such strategies are ex-post equilibria. These mechanisms look like a dominant strategy one to cognitively limited agents who neglect others may condition their behavior in sophisticated ways, and can help them avoid dominated behaviors. Moreover, they ensure sophisticated agents never have an incentive to deviate.

Our framework rationalizes the auction format chosen by prominent online platforms, such as eBay. It also provides a unified explanation for experimental evidence in various settings. Further, we provide sufficient conditions for as-if dominant strategy mechanisms to also be weak dominance solvable.

*Nagel: Stanford University, Department of Economics, email: lnagel@stanford.edu. Saitto: Stanford University, Department of Economics, email: rsaitto@stanford.edu. We are especially grateful to Mohammad Akbarpour, Doug Bernheim, Matt Jackson, Paul Milgrom, Muriel Niederle and Ravi Jagadeesan for their advice. We would also like to thank Itai Ashlagi, Eric Budish, Martino Banchio, Yeon-Koo Che, Ben Golub, Shengwu Li, Andrew Mackenzie, Antonio Penta, Al Roth, Ilya Segal, Zhiwen Zheng, the anonymous referees at EC'24, as well as the audiences at LUISS, EC'24, EIEF, SITE 2024, EEA-ESEM 2024, the DIMACS Workshop on Simplicity in Mechanism Design and Preference Elicitation, and at various seminars and Stanford workshops for their helpful comments.

1 Introduction

For the practical efficacy of mechanism design, it is important to use mechanisms that allow for accurate predictions of what real-world agents will do. Traditionally, the reliability of the predicted outcome is evaluated purely on the basis of the assumptions required by the solution concept it relies upon. This explains the popularity of *dominant strategy mechanisms*—which offer each agent a strategy that is optimal no matter what others do, and hence have predictions that rely only on the assumption that agents are rational.

However, there is increasing evidence that other aspects of a mechanism’s extensive form can also matter. In particular, several papers report surprising evidence that increasing the *transparency* of a mechanism—by allowing agents to observe more about the other agents’ past moves—can significantly improve its performance, even when this comes at the cost of breaking dominant strategy incentive compatibility (Kagel and Levin, 2009; Bó and Hakimov, 2020; Klijn et al., 2019; Hakimov and Raghavan, 2022). Indeed, Kagel and Levin (2009, p. 234) even posit that

“implementation by a mechanism that has a weaker solution concept but that is more transparent may result in closer conformity to the planner’s desired outcome,”

and stress the relevance of this dilemma when agents are not fully rational. Since the performance of dominant strategy mechanisms can be low,¹ it would be useful to systematically predict how this dilemma should be resolved. In this paper, we propose a theoretical framework to make a step in this direction.

A MOTIVATING EXAMPLE. Suppose an online platform wants to design a single-item auction with the aim of raising revenues and ensuring a satisfactory consumer experience.² To attract bidders from different locations and time zones, they want to allow bidders to access the platform and bid at their convenience within some time window.³ We refer to such auction formats as asynchronous. Classical economic theory offers the sealed-bid second-price auction (Vickrey, 1961): Sell the item to the highest bidder at a price equal to the second-highest bid. This format has the appealing theoretical property that it is a *dominant strategy* for each bidder to bid her true value

¹See, e.g., Kagel and Levin (1993); Cason et al. (2006); Chen and Sönmez (2006); Echenique et al. (2016); Rees-Jones (2018); Masuda et al. (2022); Kagel and Levin (2009); Breitmoser and Schweighofer-Kodritsch (2022); Klijn et al. (2019); Hakimov and Raghavan (2022); Bó and Hakimov (2020).

²We consider the standard quasi-linear environment: Each bidder privately knows her value θ for the item; If she receives the item with probability x and pays t , her payoff is $x\theta - t$.

³In particular, this prevents the use of clock auctions.

for the item. Hence, rationality alone suffices to predict bidders will bid truthfully. Then, (i) no bidder ever regrets her actions ex-post and (ii) under certain conditions, properly setting a reserve price maximizes expected revenues (Myerson, 1981).

Importantly, to preserve the dominant strategy property, when a bidder enters the platform to submit her bid, it is necessary that *no information* about others' bids is disclosed to her.⁴ This, however, contrasts with design choices made in practice by prominent online platforms such as eBay, which all employ significantly more *transparent* auction formats: They disclose all available information about the already-submitted bids—except for the value of the current highest bid. Should the platform choose the same disclosure policy as eBay? Should we expect higher transparency to improve the efficacy of the auction, as was the case in different settings studied in the experiments cited above? Is there a limit after which additional information may be in fact detrimental for the strategic incentives? The theoretical framework we develop in this paper allows us to provide formal answers to such questions.

First, we formalize the dilemma transparency poses to mechanism design. We say that one mechanism is *more transparent* than another if they share the same game tree and the first has a strictly finer information partition than the second. This definition lets us formalize benefits of transparency: We argue it simplifies the contingent reasoning required to spot dominated behaviors, and that it can enhance trust.⁵ However, we show that achieving dominant strategy incentive compatibility severely restricts the transparency of a mechanism's extensive-form. Further, we show that, except in special cases, increasing a mechanism's transparency eventually forces a designer to break dominant strategy incentive compatibility, and only ex-post incentive compatibility is retained.⁶ This represents a dilemma since the predictions of ex-post equilibria do not arise from rationality alone, but instead rely on equilibrium arguments and assumptions about agents ability to engage in higher-order thinking. As we discuss, simply adopting ex-post equilibria in favor of higher transparency would suggest the use of mechanisms with intuitively implausible theoretical predictions,

⁴Otherwise, some bidder no longer has a strategy that is optimal for *any* behavior of the others. In particular, truthful bidding is then only an ex-post equilibrium and hence requires additional assumptions on what she believes others do to predict her behavior. Intuitively, this is because another bidder could "adversarily" condition her bid on the information disclosed to her.

⁵As we discuss in detail in Section 3, these benefits of transparency have previously been argued informally, e.g., in Glazer and Rubinstein (1996); Zhang and Levin (2021); Cramton (1998); Kagel and Levin (2009); Bó and Hakimov (2020); Breitmoser and Schweighofer-Kodritsch (2022); Pathak (2017).

⁶Ex-post equilibria extend Nash equilibria to incomplete information games: They are strategy profiles such that, for all type profile realizations, agents always best respond to each other. Even in private value environments, ex-post and dominant strategy incentive compatibility are equivalent only for direct revelation mechanisms.

which are indeed never observed in practice.⁷

To help resolve this dilemma, we propose *as-if dominant strategy mechanisms*, which have two properties: (i) Each agent has at least one strategy that becomes dominant if the others are restricted to behave *as if* the mechanism was static—that is, they can condition their behavior only on their own private information—and (ii) all combinations of such strategies are ex-post equilibria.

From a classical perspective, these mechanisms have a particularly robust set of ex-post equilibria: As long as each agent plays *any* of their best responses to a certain *set* of strategies for the others—which includes all those that only condition behavior on their own private information—they end up in one of these equilibria.

From a behavioral perspective, condition (i) says that, to a cognitively limited agent who neglects that others may condition their behavior on hypothetical contingencies, the mechanism looks as-if it were dominant strategy. Condition (ii) says it is ex-post incentive compatible for sophisticated agents to adhere to the as-if dominant strategy prediction—even if they think some other agent may be cognitively limited.

As-if dominant strategy mechanisms are a superset of the dominant strategy ones and, importantly, allow for greater transparency. Hence, they can help cognitively limited agents recognize dominated behaviors, while otherwise being indistinguishable from dominant strategy mechanisms for them. Moreover, they can enhance agents’ trust in the mechanism, another crucial requirement for the reliability of its predictions.

While similar in scope, we use a mechanism’s extensive form to relax dominance, which makes our approach differ substantially from the recent stream of literature that uses it to refine dominance, starting from Li (2017b)’s obviously dominant strategy mechanisms.⁸ When applicable, such refinements can be very helpful to designers, but they often exclude all available dominant strategy mechanisms for important application settings.⁹ Further, their use may be prevented by practical concerns. For

⁷A vast theoretical literature supports the idea that prediction based on dominant strategies are stronger than those relying on Nash equilibria (e.g., Arrow, 1951; Luce and Raiffa, 1957; Kohlberg and Mertens, 1986; Brandenburger et al., 2008). Moreover, there is experimental evidence that play often does not correspond to equilibrium (see, e.g., Crawford et al., 2013, for a survey), and that transparent mechanisms with ex-post (Nash) equilibria perform significantly worse than similar games with a dominant strategy—e.g., comparing the ultimatum game with the dictator game (Forsythe et al., 1994).

⁸See, e.g., Zhang and Levin (2021); Pycia and Troyan (2023b); Nagel and Saitto (2024); Ferraioli and Ventre (2022); Chew and Wang (2022).

⁹For example, consider obviously dominant strategy mechanisms, defined by Li (2017b) in his seminal work on simplicity. In matching settings, stable outcomes can be implemented by an obviously dominant strategy mechanism only under strong restrictions on agents’ preferences. See also Bade and Gonczarowski (2016); Li (2017b); Ashlagi and Gonczarowski (2018); Troyan (2019); Thomas (2020); Arribillaga et al. (2023); Mandal and Roy (2022) for characterizations of social choice rules that are implementable in obviously dominant strategies and the recent general characterization of obviously

example, [Li \(2017b\)](#) shows that clock auctions are the only obviously dominant strategy mechanisms in their setting: However, they require the designer to gather all bidders together so that they can bid against the clock at the same time.

In [Section 5](#), we study one setting in which the existing refinement cannot be used due to practical concerns. In particular, our theoretical framework rationalizes the choice of auction format made by online auction houses such as eBay, Copart and Auctionet. To study this problem, we introduce a new model of second-price auctions with asynchronous bids. We show that, among those, the Online English Auction achieves maximal transparency subject to being an as-if dominant strategy mechanism. This format allows agents to observe all information about previously submitted losing bids but hides the value of the currently winning bid. By having chosen this information disclosure policy, these platforms ensure even agents who cannot profitably win—but whose value may affect the final price—have a strict incentive to bid truthfully. Any additional information would break as-if dominance.¹⁰

We also show that our approach provides a unified explanation of the experimental evidence discussed above, which shows that breaking dominant strategy incentive compatibility in favor of greater transparency sometimes improves the performance and predictability of a mechanism.¹¹ In particular, we show that the [Ausubel \(2004\)](#) auction with public dropouts, and sequential versions of the Deferred Acceptance and Top Trading Cycle algorithms are all as-if dominant strategy mechanisms that offer greater transparency to participants than their dominant-strategy counterparts, which can help account for their better predictability and performance.

Finally, in [Section 7](#), we propose perfect as-if dominant strategy mechanisms, a refinement of the as-if dominant strategy ones. They are weak dominance solvable mechanisms whose predictions are robust to the possibility that some agents are cognitively limited. We provide easy-to-check sufficient conditions to verify an as-if dominant strategy mechanism is perfect, and show they are satisfied by most of our applications. Consequently, we can apply our results to conclude that the [Ausubel \(2004\)](#) auction with public dropout information, and sequential versions of the Deferred Acceptance and Top Trading Cycles Algorithms are all weak dominance solvable. While a direct proof of the former result is in [Ausubel \(2004\)](#), to the best of our knowledge, the results for the latter mechanisms are novel.

While our approach can help designers identify effective transparent mechanisms, dominant strategy mechanisms provided by [Pycia and Troyan \(2023a\)](#).

¹⁰Other examples of mechanisms that are not as-if dominant include, e.g., the well-known beauty contest, and tâtonnement auctions for heterogeneous goods ([Ausubel, 2006a](#); [Gul and Stacchetti, 2000](#)).

¹¹See [Kagel and Levin \(2009\)](#); [Bó and Hakimov \(2020\)](#); [Klijn et al. \(2019\)](#); [Hakimov and Raghavan \(2022\)](#); [Bó and Hakimov \(2023\)](#).

our analysis abstracts away from some considerations that can be important for practical design choices. First, the form in which information is disclosed is important for agent to be able to process it: For example, eBay highlights the value of the current highest losing bid, but gives agents the opportunity to access more information if they want. Second, there may be privacy concerns limiting the amount of transparency a designer wants to provide. For example, in school choice settings, a designer may prefer to only reveal aggregate statistics, such as school-specific cutoffs.¹² Third, time constraints may prevent frequent communication between the agents and the mechanism. Finally, if agents are sophisticated and interact repeatedly—e.g., in government procurement auctions—too much transparency may help support collusive behavior. Thus, as-if dominance is best-suited for settings in which agents interact sporadically.

1.1 Other Related Literature

We relate to the theory of ex-post implementation [Bergemann and Morris \(2008, 2005\)](#). [Ausubel \(2004\)](#) defines the refinement of ex-post perfect equilibrium. While all as-if dominant strategy mechanisms are ex-post incentive compatible, even an ex-post perfect incentive compatible mechanism may not be an as-if dominant strategy mechanism. Our solution concept also does not nest with the concepts of obvious ex-post equilibrium, partition obvious Nash equilibrium and everywhere dominant equilibrium—proposed, respectively, by [Li \(2017a\)](#), [Zhang and Levin \(2017\)](#) and [Mackenzie and Zhou \(2022\)](#).

We also relate to the literature studying mechanisms whose normal form allows agents to have a bounded depth of rationality ([Saran, 2016](#); [de Clippel et al., 2019](#); [Börgers and Li, 2019](#); [Crawford, 2021](#); [Jackson et al., 1994](#); [Gorelkina, 2018](#); [Kneeland, 2022](#)). Our approach is substantially different, since as-if dominance is an extensive-form concept: It builds on the idea that the information partition of a mechanism can lead agents who are less than fully rational to perceive others' strategy space as simpler, restoring the incentive properties of a dominant strategy mechanisms. Hence, even agents with limited sophistication are guided towards equilibrium play.

Finally, our characterization of the Online English Auction relates us to a literature on online auctions more generally, and eBay specifically, as reviewed in [Hasker and Sickles \(2010\)](#). Unlike existing models of decentralized online platforms, such as [Bodoh-Creed et al. \(2020\)](#), we do not focus on the problem of matching buyers and sellers, but on the choice of information disclosure policy for a single auction.

¹²The role of cutoffs in school choice is considered in, e.g., [Azevedo and Leshno \(2016\)](#); [Hakimov and Raghavan \(2022\)](#); [Leshno and Lo \(2020\)](#).

2 Preliminary Definitions

We consider environments with an at most countable set N of agents n , each with a type space Θ_n . Each type $\theta_n \in \Theta_n$ defines a complete preference ordering \leq_{θ_n} over some outcome space Ω . Let $\Theta := \times_n \Theta_n$ and define Θ_{-n} in the standard way.

2.1 Trees, Information Partitions and Mechanisms

For what follows, it will be convenient to distinguish a mechanism's game tree from the chosen information partition. Let c denote the chance player and $N_c := N \cup \{c\}$. A *tree* is a triple $X = (H, P, \omega)$ composed of:

- (i) A set H of histories h from a set of actions, equipped with the standard precedence partial order \preceq . Let \emptyset be the root history and Z the set of terminal histories z .¹³ For all $h \in H \setminus Z$, let $A(h)$ be the the set of actions a with $(h, a) \in H$.
- (ii) A player function $P: H \setminus Z \rightarrow N_c$. Agent n is called to play at h if $P(h) = n$.
- (iii) An outcome function $\omega: Z \rightarrow \Omega$.

Given a tree X and a partition \mathcal{I} of a subset of the set $H \setminus Z$ of non-terminal histories, let $\mathcal{I}(h)$ be the element of \mathcal{I} that contains h .¹⁴ An *information partition* \mathcal{I} for a tree X is a partition of $H \setminus Z$, whose elements are called *information sets* and satisfy the following: For all agents n , there is $\mathcal{I}_n \subset \mathcal{I}$ such that \mathcal{I}_n is a partition of the set of histories at which n is called to play. If $P(h) = c$, then $\mathcal{I}(h)$ is a singleton. If $\mathcal{I}(h) = \mathcal{I}(h')$, then $A(h) = A(h')$. All agents have perfect recall.¹⁵

A *mechanism* is a pair (X, \mathcal{I}) , consisting of a tree X and an information partition \mathcal{I} for X . Given a mechanism (X, \mathcal{I}) , a *behavior* σ_n for agent n maps each history h at which n is called to play into some action $a \in A(h)$ and is \mathcal{I} -measurable, that is: If $\mathcal{I}_n(h) = \mathcal{I}_n(h')$, then $\sigma_n(h) = \sigma_n(h')$. Let Σ_n be the set of behaviors for agent n . If n is never called to play, let $\Sigma_n := \{\emptyset\}$. For the chance player, a behavior σ_c and Σ_c are defined analogously. Let $\Sigma := (\times_n \Sigma_n) \times \Sigma_c$, with generic element σ , and define Σ_{-n} in the standard way. For all σ , let $\omega(\sigma)$ be the outcome that is implemented at the

¹³More formally, H is a set of sequences h of actions. If $h = (h', h'') \in H$, then $h' \in H$. All non-terminal histories are finite sequences. If an infinite sequence z of actions is such that there is a sequence of histories $h^k \in H$ such that $h^k \rightarrow z$ as k goes to infinity, then $z \in H$. The precedence order is defined such that $h' \preceq h''$ if and only if $h'' = (h', h)$.

¹⁴Throughout, we treat two trees X_1 and X_2 as the same if there exists an order isomorphism $f: H_1 \rightarrow H_2$ such that (i) $P_2 \circ f = P_1$, (ii) $\omega_2 \circ f = \omega_1$, and (iii) the ℓ -th component of $f(h')$ and $f(h'')$ is the same if and only if the ℓ -th component of h' and h'' is the same.

¹⁵Agent n 's *experience* at h is the sequence of all information sets I_n and messages m that—along h —agent n respectively encounters and sends, in the order these events occur. There is *perfect recall* if n 's experience at any two given histories is the same whenever they lie in the same I_n , for all n .

terminal history z that is reached under σ . Furthermore, for all non-terminal histories $h \in H$ and behavior profiles $\sigma \in \Sigma$, let $\omega(\sigma | h)$ be the outcome that is implemented if h is reached and starting from h , all agents and the chance player follow σ .¹⁶ We extend behaviors from histories to information sets in the obvious way.

We assume that all mechanisms (X, \mathcal{I}) are such that, for all chance moves $\sigma_c \in \Sigma_c$, there is a finite number $B(\sigma_c)$ such that all terminal histories $z(\sigma_{-c}, \sigma_c)$ that can be reached for some agents' behavior profile σ_{-c} have length at most $B(\sigma_c)$. All finite-length mechanisms satisfy this assumption.

In what follows, a key role is played by the mechanism which has the coarsest information partition for a given tree. To this end, we state the following result.¹⁷

PROPOSITION 1. *Every tree has a coarsest information partition.*

2.2 Dominant Strategies, Rationality and Ex-post Equilibria

We now define the solution concepts that in the paper we use as a benchmark for the strength of the strategic incentives a mechanism provides.¹⁸

Fix a mechanism (X, \mathcal{I}) . Recall that a strategy is a fully contingent plan which specifies the agent's behavior in the mechanism conditional on their private information. Formally, a *strategy* for agent n is a map $s_n: \Theta_n \rightarrow \Sigma_n$. Let S_n denote the set of strategies s_n for agent n . A *strategy profile* s specifies a strategy for each agent n . Let $s(\theta) := (s_n(\theta_n))_n$ and $S := \prod_n S_n$. Define $s_{-n}(\theta_{-n})$ and S_{-n} in the obvious way.

We now define the benchmark solution concepts. A strategy profile s_n is a *dominant strategy equilibrium* if each s_n is a dominant strategy, i.e. a best response for agent n no matter how others behave. Formally, a strategy s_n is *dominant* if $\omega(\sigma_n, \sigma_{-n}) \leq_{\theta_n} \omega(s_n(\theta_n), \sigma_{-n})$ for all types θ_n , behaviors σ_n , and others' behavior profiles σ_{-n} .

Mechanisms providing each agent with a dominant strategy are theoretically attractive because their predictions rely on agents' rationality alone: Following a stream of foundational literature—see, e.g. [Arrow \(1951\)](#); [Luce and Raiffa \(1957\)](#); [Kohlberg and Mertens \(1986\)](#); [Brandenburger et al. \(2008\)](#)—an agent is *rational* if she only plays admissible strategies, that is strategies for which there is no other strategy that always leads to a better outcome, sometimes strictly so. Formally, a strategy s_n for agent n is *admissible* or *undominated* if there is no type θ_n and behavior σ_n such that

¹⁶Formally, $\omega(\sigma | h) := \omega(z)$ where z is the unique terminal history with $h \preceq z$ and $(h', \sigma(h')) \preceq z$ for all h' such that $h \preceq h' \prec z$. Then, $\omega(\sigma) := \omega(\sigma | \emptyset)$.

¹⁷For mechanisms of finite length, [Cohen and Li \(2023\)](#) provide a constructive proof.

¹⁸For these definitions, we treat chance as any other agent. That is, agents do not need to take expectations over chance moves, as in [Li and Dworzak \(2022\)](#). This allows us to simplify notation and exposition—in particular, we can avoid measure-theoretic technicalities. Our results would not be significantly affected by adopting the more traditional view.

$\omega(s_n(\theta_n), \sigma_{-n}) \leq_{\theta_n} \omega(\sigma_n, \sigma_{-n})$ for all σ_{-n} , with at least one strict inequality.¹⁹ In a dominant strategy mechanism, only the outcomes associated with a dominant strategy profile can result from the play of rational agents. Importantly, to determine their own admissible strategies, an agent does not need to know the others' payoffs nor that the others are rational.

The second benchmark solution concept, *ex-post equilibrium*, relies on stronger assumptions than rationality. Intuitively, it is a strategy profile such that agents play a Nash equilibrium for each possible type realization. Formally, an *ex-post equilibrium* is a profile s such that $\omega(\sigma_n, s_{-n}(\theta_{-n}), \sigma_c) \leq_{\theta_n} \omega(s_n(\theta_n), s_{-n}(\theta_{-n}), \sigma_c)$ for all agents n , types θ_n , behaviors σ_n , others' type profiles θ_{-n} and chance behaviors σ_c . An *ex-post equilibrium* is *strict* if the inequality always holds strict. All dominant strategy equilibria are *ex-post equilibria*, but the converse is not true. To see that rationality alone does not suffice to support *ex-post equilibria*, observe that, in order for rational agents to play an *ex-post equilibrium*, they need to additionally (i) mutually specify an exact belief about others' strategies and (ii) break any indifference in a manner consistent with others' beliefs. Assuming that agents adopt a specific mutually consistent system of strategic beliefs is much stronger than the assumption that agents are rational. Whenever there are multiple undominated *ex-post equilibria*, even rational agents may be uncertain about how the others will play. Moreover, even if there is a unique *ex-post equilibrium*, knowledge of the others' payoffs and rationality, and the ability to engage in higher-order thinking are necessary (and not always sufficient) for agents to converge to the equilibrium belief.

We refer to a mechanism with a dominant strategy (ex-post) equilibrium as a dominant strategy (ex-post incentive compatible) mechanism. Consistently with the above discussion, we say that dominant strategy mechanisms provide stronger strategic incentives than ex-post incentive compatible ones.²⁰

3 The Dilemma of Mechanism Transparency

In this section we formally define our notion of transparency, and argue that increasing transparency can increase a mechanism's performance. We then show that this leads to a dilemma for mechanism design, since, often, a designer must choose between increasing the transparency of a mechanism's extensive form and maintaining

¹⁹The notion of rationality adopted in this paper is stronger than the one which only requires agents to play strategies that are not *strictly dominated*, i.e. for which there is another strategy that always leads to a strictly better outcome (see, e.g., [Bernheim, 1984](#); [Pearce, 1984](#)).

²⁰Note, however, that in a direct revelation mechanism, the truthful revelation strategy profile is a dominant strategy equilibrium if and only if it is an *ex-post* one.

dominant strategy incentive compatibility.

3.1 The Benefits of Transparency

We will say that a mechanism is more transparent than another if they share the same tree and the former has a strictly finer information partition than the latter. This defines a straightforward partial order which allows us to compare the transparency of mechanisms that share the same game tree.

DEFINITION 1. Given a tree X , a mechanism (X, \mathcal{I}') is more transparent than a mechanism (X, \mathcal{I}'') if \mathcal{I}' is strictly finer than \mathcal{I}'' .

Intuitively, a mechanism is more transparent if, whenever an agent moves, she is provided with more information about other agents' past moves. If agents' preferences are interdependent, the value of more transparent mechanisms is well known. However, even in private value environments, there are several reasons why increasing transparency may improve a mechanism's performance.

To formalize this, fix a strategy profile s for a mechanism (X, \mathcal{I}) .²¹ Given an agent n with type θ_n and a history h at which n moves, an action $a \in A(h)$ is *deviating* if $s_n(\theta_n)(h) \neq a$. A deviating action a is *obviously dominated* for θ_n if, for all $h' \in \mathcal{I}(h)$, terminal histories z such that $(h'', a) \prec z$ for some $h'' \in \mathcal{I}(h)$, and $\sigma_{-n} \in \Sigma_{-n}$, we have

$$\omega(z) \leq_{\theta_n} \omega(s_n(\theta_n), \sigma_{-n} \mid h').$$

Intuitively, obviously dominated deviating actions can be seen to be dominated without having to condition on the other agents' behavior.

PROPOSITION 2. *Let s be a dominant strategy equilibrium of some mechanism (X, \mathcal{I}) . Then, given a more transparent mechanism (X, \mathcal{I}') : For all agents n , types θ_n and histories h at which n moves, the set of obviously dominated deviating messages increases.*

The above proposition says that for a mechanism with a dominant strategy profile s , increasing transparency makes more deviations from the strategy profile s look obviously dominated—despite the fact that now agents have more strategies. This can make it easier for agents with limited ability to reason contingently to play optimally.

Moreover, even if the set of obviously dominated deviating actions stays the same, the following fact always holds.

²¹The following definitions are the natural counterpart to the definition of an obviously dominant strategy provided by Li (2017b).

FACT 1. *Suppose we increase a mechanism’s transparency. Then, whenever agent n moves at some h' , comparing her different continuation behaviors requires less contingent reasoning.*

Formally, for any belief σ_{-n} about others’ continuation behavior, the information set containing h' contains a smaller set of histories h agent n needs to condition upon to compare the outcomes $\omega(\sigma'_n, \sigma_{-n} \mid h)$ and $\omega(\sigma''_n, \sigma_{-n} \mid h)$ of two continuation behaviors σ'_n and σ''_n .

A large experimental literature provides evidence that this benefit of more transparent mechanisms—reducing uncertainty and making the consequences of an agent’s actions more apparent—can help less than fully rational agents to play optimally (e.g., Kagel and Levin, 2009; Bó and Hakimov, 2020; Hakimov and Raghavan, 2022).²² Strikingly, Breitmoser and Schweighofer-Kodritsch (2022) show that transparency can be helpful even within the set of relatively simple mechanisms that satisfy obvious strategy-proofness: In their experiment, a more transparent OSP clock auction in which bidders see the number of remaining competitors outperforms the OSP clock auction in which this information is not provided to bidders. Similar evidence is also provided by the experiment in Esponda and Vespa (2014) in the context of obviously strategy-proof voting games.²³

FACT 2. *Suppose we increase a mechanism’s transparency. Then, agents can more easily detect whether the designer deviates from the mechanism’s rules.*

Formally, suppose agent n moves at h . For any given continuation behavior σ_n , the set of sequences of future information sets and outcomes consistent with the mechanism’s rules for some behavior profile σ_{-n} is non-increasing in the size of the information set containing h .

The above fact shows that transparency as we define it makes it easier to detect deviations by the designer, mitigating the concern that agents may deviate from the predicted behavior because they do not trust the mechanism. In particular, increasing the transparency of a mechanism preserves properties such as auditability, verifiability, trustworthiness, and credibility (see, e.g., Grigoryan and Möller, 2023; Akbarpour and Li, 2020; Woodward, 2020; Hakimov and Raghavan, 2022; Pycia and Ünver, 2020). In contrast, reducing transparency may break these properties.

3.2 A Dilemma for Mechanism Design

The benefits of transparency may lead a designer prefer information partitions that are as transparent as possible, while satisfying any constraints imposed by the setting, such as anonymity of the players or time constraints. However, the amount of

²²This also informally argued in, e.g., Cramton (1998); Zhang and Levin (2021). Similarly, Rubinstein (1986) formalizes the idea that the extensive form of a game can help agents play optimally.

²³See also Niederle and Vespa (2023) for a review of the evidence on failures of contingent thinking.

information provided also affects the agents' strategic incentives. The next proposition shows that increasing the transparency of a mechanism can preserve dominant strategy incentive compatibility only if agents' preferences satisfy a rapidly increasing set of inequalities. Intuitively, if agent m is revealed information about what agent n does, agent n must prefer the worst outcome that may realize for any one of agent m 's continuation behaviors if she follows her dominant strategy to the best outcome that may realize for any one of agent m 's continuation behaviors if she deviates in a way that agent m can detect.

PROPOSITION 3. *Let s be a dominant strategy equilibrium of some mechanism (X, \mathcal{I}) . Suppose there is a behavior σ_n for agent n which first deviates from $s_n(\theta_n)$ at I_n , and let σ_{-n} be such that some agent $m \neq n$ can detect the deviation. That is, there are $I'_m \neq I''_m$ such that*

- (i) *under $(s_n(\theta_n), \sigma_{-n})$, I_n is reached and agent m is first called to play again at $I'_m \succ I_n$,*
- (ii) *under (σ_n, σ_{-n}) , I_n is reached and agent m is first called to play again at $I''_m \succ I_n$.*

Then, $\omega(\sigma_n, \sigma'_m, \sigma_{-n,m}) \leq_{\theta_n} \omega(s_n(\theta_n), \sigma''_m, \sigma_{-n,m})$ for all σ'_m, σ''_m that agree with σ_m up to I_n .

As an illustration of just how severely transparency is limited if the designer wants to maintain dominant strategy incentive compatibility, consider any dominant strategy direct revelation mechanism. Each agent is asked to report her type and she has no information about what the others do. By Proposition 3, increasing transparency by revealing whether agent n reported to be of a certain type to all the other agents can preserve dominant strategy incentive compatibility only if that type prefers the worst outcome that may realize if she reports truthfully to the best outcome that may realize if she misreports. Thus, this type must be able to verify that truthful reporting is optimal for her without conditioning on the others' behavior.²⁴

This requirement, which is much stronger than dominance, was introduced by Li (2017b) as obvious dominance. Formally, a behavior σ'_n is *obviously dominant* for type θ_n if, for all information sets I_n that are reached under (σ'_n, σ_{-n}) for some σ_{-n} ,

$$\omega(\sigma''_n, \sigma'_{-n}) \leq_{\theta_n} \omega(\sigma'_n, \sigma''_{-n})$$

for all $\sigma''_n, \sigma'_{-n}, \sigma''_{-n}$ such that $\sigma'_n(I_n) \neq \sigma''_n(I_n)$ and I_n is reached under both $(\sigma''_n, \sigma'_{-n})$ and $(\sigma'_n, \sigma''_{-n})$. A strategy profile s is an *obviously dominant strategy equilibrium* if $s_n(\theta_n)$ is obviously dominant for θ_n , for all agents n and types θ_n .

²⁴As a concrete example, in a sealed-bid second price auction, revealing any information about a bidder's bid breaks the dominant strategy property.

Proposition 3 above speaks to how severely transparency restricts the payoffs in given mechanism, i.e., on an intensive margin, if we want to maintain dominant strategy incentive compatibility. Using obvious dominance, we can now also derive an extensive margin result, which shows for which mechanisms increasing transparency will eventually lead to a violation of dominant strategy incentive compatibility.²⁵

PROPOSITION 4. *Let s be a dominant strategy equilibrium of some mechanism (X, \mathcal{I}) . Unless s is obviously dominant when all information sets are singletons, there is a mechanism more transparent than (X, \mathcal{I}) for which s is still an ex-post equilibrium, but is no longer dominant.*

That is, unless the equilibrium of a dominant strategy can become obviously dominant, increasing transparency eventually leads to a violation of dominant strategy incentive compatibility. Hence, the designer necessarily faces a dilemma between increasing the transparency of a mechanism’s extensive form and the strength of the solution concept used for the implementation.

Given the existing literature studying obviously dominant mechanisms, we conclude this dilemma is present in many practical applications: First, for several relevant dominant strategy-implementable choice rules, obviously dominant mechanisms do not exist—e.g., for stable outcomes in matching settings and for the VCG outcome in multi-unit settings.²⁶ Second, even when such mechanisms do exist—e.g., the single-unit Vickrey outcome can be implemented by a clock auction—they may be infeasible for practical reasons. In particular, obviously dominant mechanisms are generally dynamic and thus require a high level of coordination.²⁷

While existing experimental evidence suggests that at least in some applications, increasing a mechanism’s transparency can improve its performance even at the cost of abandoning dominant strategy incentive compatibility (Kagel and Levin, 2009; Bó and Hakimov, 2020; Hakimov and Raghavan, 2022; Bó and Hakimov, 2023), only a theoretical approach can help designers systematically predict when this is the case. Indeed, it is well-documented across different settings that (ex-post) Nash equilibria are generally not a good prediction of agents’ behavior.²⁸

²⁵This result relates to Mackenzie (2020), who shows that in a mechanism of perfect information—i.e., in which all information sets are singletons—a strategy is dominant if and only if it is obviously dominant. Ashlagi and Gonczarowski (2018) show obvious dominance is preserved under perfect information.

²⁶Ashlagi and Gonczarowski (2018) show that stable mechanisms are not obviously dominant, for the multi-unit auction setting, see Section 6.1.

²⁷Pycia and Troyan (2023a) show that obviously dominant strategy mechanisms are equivalent to generalized millipede games, in which agents move sequentially.

²⁸See, e.g., Crawford et al. (2013) for a survey of some of the experimental evidence on non-equilibrium behavior, as well as some of the theoretical literature developing non-equilibrium models in response to this evidence.

4 As-If Dominant Strategy Mechanisms

In Section 3.1, we established a dilemma between dominant strategy incentive compatibility and the transparency of a mechanism’s extensive form. Next, we propose *as-if dominant strategy mechanisms* to help designers resolve this dilemma.

Fix a mechanism (X, \mathcal{I}) , and let \mathcal{S} be the coarsest information partition that can be defined on X . A behavior $\sigma_n \in \Sigma_n$ is *coarse* if it is \mathcal{S} -measurable. Let $\Sigma_n(\mathcal{S}) \subset \Sigma_n$ be the set of coarse behaviors for agent n , and let $\Sigma_{-n}(\mathcal{S})$ be the set of coarse behavior profiles for the others.²⁹

DEFINITION 2. An as-if dominant strategy mechanism is a mechanism such that:

- (i) Each agent has at least one coarsely optimal strategy, i.e., a strategy s_n such that $\omega(\sigma_n, \sigma_{-n}) \leq_{\theta_n} \omega(s_n(\theta_n), \sigma_{-n})$ for all types $\theta_n \in \Theta_n$, behaviors $\sigma_n \in \Sigma_n$, and coarse behavior profiles $\sigma_{-n} \in \Sigma_{-n}(\mathcal{S})$.
- (ii) All strategy profiles s such that each s_n is coarsely optimal are ex-post equilibria.

Given an as-if dominant strategy mechanism, an *as-if dominant strategy* is a strategy that is coarsely optimal.³⁰

From a classical perspective, as-if dominant strategy mechanisms have a robust set of ex-post equilibria, such that (i) each agent is playing optimally against set of behaviors for the others—including all the coarse ones—and (ii) they do not require agents to coordinate on one of them in order to best respond to each other. This allows us to predict agents’ behavior even if we relax the assumption that all agents hold a specific common strategic belief.

From a behavioral perspective, as-if dominant strategy mechanisms are ex-post incentive compatible mechanisms whose predictions are robust to the possibility that some agents, while rational, have cognitive limitations. Specifically, we say that an agent is *cognitively limited* if:

²⁹The definition of coarse behaviors can depend on the labeling of actions, and we take the labels as given for the application settings. All our results would stand if we allowed to relabel the mechanism to minimize the set of coarse actions.

³⁰The definition could be stated in slightly greater generality. In particular, we could say that a mechanism (X, \mathcal{I}) is as-if dominant strategy mechanism if there exists a mechanism (X', \mathcal{I}') in which conditions (i) and (ii) hold true that satisfies the following: It has the same reduced normal form as (X, \mathcal{I}) , and is such that $H \subset H'$ and the maps P , ω , and \mathcal{I} respectively coincide with the restrictions of P' , ω' , and \mathcal{I}' on H . This would allow a designer who provides information that is known to make certain strategies outcome-equivalent to eliminate such duplicate strategies—which would increase the number of coarse behaviors—without making condition (i) of our definition harder to satisfy. For ease of exposition, throughout the paper, when considering applications we simply assume designers do not eliminate such actions. Our results would be unchanged under this more general definition.

- (a) She has a simplified view of what others can do, neglecting that they may condition their behavior on more than just their private information.
- (b) She cannot engage in higher-order thinking, and thus may play any behavior which is undominated given her simplified view of what others can do.

Intuitively, this means we assume agents who cannot engage in higher-order thinking are also limited in their contingent reasoning ability, as they fail to realize that others may condition their behavior on hypothetical contingencies. Now suppose if no agent was cognitively limited, they would play certain ex-post equilibrium,³¹ but there is a chance that some agents are cognitively limited. In general, this would make the prediction of a mechanism via ex-post equilibrium unreliable. However, if the equilibrium is a coarsely optimal profile of an as-if dominant strategy mechanism, this does not constitute a problem. First, condition (i) of Definition 2 ensures that, as in a dominant strategy mechanism, the play of cognitively limited agents can be predicted by rationality alone—they have a strategy that looks dominant to them. Second, condition (ii) ensures that all agents—cognitively limited or not—are best responding to each other: As in an ex-post incentive compatible mechanism, sophisticated agents have no incentive to deviate—even if they think that some other agent may be cognitively limited.³² Moreover, cognitively limited agents cannot gain from unilaterally developing a more sophisticated view of what the others can do.

As we show next, the cognitively limited agents' view that others can only behave coarsely is always consistent with observed play. This implies every node in the tree can be reached by some coarse behavior, and thus supports the idea that, by an Occam's razor argument, a cognitively limited agent may simply neglect that others have access to behaviors that condition on hypothetical contingencies.³³

PROPOSITION 5. *For all mechanisms, if history h is reached under some behavior profile (σ_n, σ_{-n}) , then there is a coarse behavior profile σ'_{-n} such that h is reached under (σ_n, σ'_{-n}) .*

Importantly, as-if dominant strategy mechanisms are a proper superset of the dominant strategy ones.

³¹In Section 7, we refine as-if dominant strategy mechanisms such that their predictions are in fact their weak dominance solution and hence no longer rely on any equilibrium arguments.

³²If each coarse behavior is undominated for some type when all other agents behave coarsely, the coarsely optimal strategies are in fact the only strategies that are a best response to all profiles of coarsely optimal behaviors for the others. In this case, as long as sophisticated agents consider the presence of cognitively limited agents possible, they can only safely coordinate on coarsely optimal equilibria. This holds true in all the applications we consider in the paper.

³³The idea that agents hold a simplified view of how the others behave is also the basis of analogy-based expectation equilibrium (Jehiel, 2005), a solution concept for multi-stage games with perfect information in which players best respond to expectations about average behavior.

THEOREM 1. *If a mechanism is dominant strategy, it is as-if dominant strategy. Moreover, all as-if dominant strategy profiles are ex-post equilibria in undominated strategies.*

The main step of the proof of this result consists of showing that a dominated strategy can never be coarsely optimal—which is not immediate since the set of coarse behaviors for the others is, in general, a strict subset of all behaviors. Thus, as-if dominant strategy equilibria are also always undominated ex-post equilibria.

Given the above result, the usefulness of as-if dominant strategy mechanisms relies on the fact that they can allow for higher transparency than dominant strategy ones. Indeed, by Fact 1 and Proposition 2, increasing transparency can help agents with limited ability to reason contingently, helping them spot their dominated behaviors. Hence, a transparent as-if dominant strategy mechanism may be effective exactly when there are cognitively limited agents whose limited ability to reason contingently also prevents them from behaving optimally in a less transparent dominant strategy mechanism. Moreover, by Fact 2, increasing transparency can enhance agents' trust in the mechanism, a crucial condition for the reliability of its predictions.

We conclude the section observing that, to determine which mechanisms are as-if dominant strategy, we cannot consider solely their normal form. Indeed, an as-if dominant strategy can differ from a dominant strategy only if the set of coarse behaviors is a proper subset of the set of all behaviors.³⁴ The following two examples illustrate the crucial role that the information revealed by the extensive-form has in determining whether a mechanism is (as-if) dominant strategy. We also use them to show that there are as-if dominant strategy mechanisms which are not dominant strategy, and that there are ex-post incentive compatible mechanisms which are not as-if dominant strategy.

EXAMPLE 1 (Open Outcry English Auction). Several bidders are in a room bidding for an item, for which each has a private value $\theta_n \in \{0, 1, \dots, v\}$. Bidders have quasi-linear payoffs, and prefer winning the item at their value rather than losing. Starting from zero, the auctioneer asks for offers in $\{1, \dots, v\}$.³⁵ At the end of each round of offers, the item is tentatively assigned to one of the highest bidders—ties broken according to some predetermined order. Bidders can only make offers that exceed their previous ones. If at some point no new offer is made, the item is assigned to the tentative winner at the last offer she made. The fact that each bidder can observe everything the others do and the possibility of jump bidding implies this auction

³⁴This is a fundamental difference to approaches studying strategic simplicity beyond dominant strategy incentive compatibility that focus on a mechanism's normal form, as Börgers and Li (2019).

³⁵Bidders answer the auctioneer simultaneously.

is *not* a dominant strategy mechanism.³⁶ However, it is an as-if dominant strategy mechanism: A coarsely rational bidder should always make the minimum offer that can make her the tentative winner until this is below her valuation. This classification thus provides a formal justification for the long-standing habit—for private value environments—of modeling the English auction as a second-price sealed-bid auction (e.g., McAfee and McMillan, 1987; Milgrom, 1989)

EXAMPLE 2 (Ultimatum Game). Agent 1 makes a proposal to agent 2 on how to split a dollar.³⁷ Agent 2 observes the offer and either accepts or rejects. If he accepts, each agent gets their respective proposed share of the dollar. If he rejects, both get zero.

The unique ex-post equilibrium consistent with subgame perfection is for agent 1 to demand the entire dollar, and for agent 2 to accept all offers. However, this *is not* an as-if dominant strategy mechanism. To see why, notice that if player 2 is restricted to coarse behaviors, it is a dominant strategy for player 1 to demand the entire dollar. However, player 2 has no strict incentive to accept an offer that yields him zero payoffs. Indeed, it is also coarsely optimal for him to accept if and only if agent one offers him a positive amount. Hence, even if all agents play a coarsely optimal strategy, they may not end up in an ex-post equilibrium.³⁸

Consistently with this classification, the idea that the strategic incentives of agent 1 are substantially affected by the possibility that agent 2 can condition his decision on agent 1's offer is supported by experimental evidence (e.g., Forsythe et al., 1994).

4.1 A Microfoundation of As-if Dominant Strategy Mechanisms

Before moving on to applications, we provide one possible microfoundation of the behavioral interpretation of as-if dominant strategy mechanisms presented above.

Fix a mechanism (X, \mathcal{I}) in which each agent has a coarsely optimal strategy. Suppose agents have two selves, which play the following two-stage game.

The first self chooses her second self's view of the others' strategy spaces, by determining which hypothetical contingencies others can condition their behavior upon. Formally, each agent n 's first self picks an information partition \mathcal{J}_{-n} for the others, which is coarser than \mathcal{I}_{-n} . We say it is (*strictly*) *cognitively costly* for agents to consider more hypothetical contingencies other agents can condition their behavior upon if, conditional on the same outcome being reached, they (*strictly*) prefer to

³⁶Intuitively, this is because others could condition their behavior on the increment of the price.

³⁷The important assumption here is that agent 2 is indifferent between rejecting and accepting the lowest proposal agent 1 can make.

³⁸Another example of a mechanism with a (perfect) ex-post equilibrium that is not as-if dominant strategy is Ausubel (2006b)'s heterogeneous commodity ascending auction.

choose a strictly coarser partition.³⁹

The second self is cognitively limited: She forms a belief about others' strategies which only considers strategies measurable with respect to the view \mathcal{J}_{-n} chosen by her first self. Moreover, she does not engage in higher-order reasoning and simply plays one of the undominated best responses given her belief. The first self does not know which strategic belief her second self will hold, nor which way eventual indifference will be resolved. We model this by letting chance pick at random one strategic belief σ_{-n} that is measurable with respect to the chosen view \mathcal{J}_{-n} , and then pick at random one of the undominated best responses to σ_{-n} .

PROPOSITION 6. *Consider the two-stage game described above.*

- (i) *Suppose that, whenever considering more hypothetical contingencies is costly, choosing to consider the coarsest information partition is an ex-post equilibrium. Then, (X, \mathcal{I}) is an as-if dominant strategy mechanism.*
- (ii) *Suppose (X, \mathcal{I}) is an as-if dominant strategy mechanism and that considering more hypothetical contingencies is (strictly) costly. Then, choosing to consider the coarsest information partition is a (strict) ex-post equilibrium.*

Hence, as-if dominant strategy mechanisms are exactly those where each agent has a coarsely optimal strategy and, as long as contingent thinking is costly, agents would rationally choose to neglect behaviors for the others that are not coarse.

5 An Application to Online Auctions

We now apply the concepts of transparency and as-if dominance to characterize single-unit online auctions in which the as-if dominant strategy equilibrium implements either the profit maximizing or the efficient outcome. In particular, suppose an online platform wants to choose a auction format that provides as-if dominant strategy incentives and that is also highly transparent to benefit from the additional trustworthiness and simplicity that transparency can provide.

In principle, clock auctions are a possible way to achieve the platform's goals. Indeed, clock auctions are obviously dominant, and hence can be made fully transparent while maintaining dominant strategy—and hence also as-if dominant strategy—incentives. The strategic simplicity, credibility and practical efficacy of this auction format has been extensively discussed by a recent stream of literature (Li, 2017b;

³⁹For ease of exposition, we assume all agents incur these cognitive costs. An analogous argument holds if only some agents do.

Akbarpour and Li, 2020; Pycia and Troyan, 2023b; Nagel and Saitto, 2024; Zhang and Levin, 2021; Milgrom and Segal, 2020). Moreover, recent experimental evidence supports that the most transparent version of the clock auction—in which bidders can observe at which price the other bidders drop out—is particularly well-suited to implement the Vickrey outcome (Breitmoser and Schweighofer-Kodritsch, 2022).⁴⁰

However, practical considerations can prevent clock auctions from being feasible. Indeed, a live clock auction requires the platform to gather all bidders together for the entire duration of the auction. While such a high level of coordination is sometimes employed for high-stakes auctions,⁴¹ it may be infeasible for online platforms which run hundreds of thousands of smaller auctions each year. Indeed, in practice, these platforms almost exclusively choose asynchronous formats, i.e., let bidders submit their bid at any time within a given time frame.

This raises the question of which asynchronous auction formats are as-if dominant strategy and allow for a high level of transparency. We propose an asynchronous format—the Online English Auction—as an answer to the platform’s design problem. We argue that this format satisfies the discussed desirable properties, as (i) it achieves maximal transparency among the set of as-if dominant strategy second-price auctions with asynchronous bids, and (ii) it implements the dominant strategy outcome of a standard second-price auction, and hence can implement both the efficient and the profit maximizing outcome. While this format is novel to the literature, it is employed by prominent online auction platforms such as eBay.

5.1 Second-Price Auctions with Asynchronous Bids

We start by defining second-price auctions with asynchronous bids. Throughout, we consider a quasi-linear single-unit auction setting. In particular, let \mathcal{D} be a finite subset of the non-negative real line containing zero, and define $\Theta_n := \mathcal{D}$ and $\Omega_n := \{0, 1\} \times \mathcal{D}$. The set of outcomes Ω consists of all pairs (x, b) specifying an allocation $x = (x_n)_n$ and a payment vector $b = (b_n)_n$ such that $(x_n, b_n) \in \Omega_n$ for all n , and $\sum_n x_n \leq 1$. Agents’ utility is quasi-linear, that is, for each agent n and type $\theta_n \in \Theta_n$, we have $(x, b) \leq_{\theta_n} (x', b')$ if and only if $\theta_n x_n - b_n \leq \theta_n x'_n - b'_n$. It is convenient to let the set of agents N be the set of all positive integers.

Informally, a second-price auction with asynchronous bids is an auction that allows agents to submit bids over some fixed time interval. Each agent privately decides when to access an online bidding platform to submit her bid. She can access it

⁴⁰In particular, Breitmoser and Schweighofer-Kodritsch (2022) find subjects underbid less in clock auctions that reveal when other bidders drop out than in ones that do not reveal this information.

⁴¹One example is the FCC incentive auction (Milgrom and Segal, 2017). A high level of coordination is also required for the auctions run by auction houses such as Sotheby’s and Christie’s.

several times and revise her bid upwards if she likes. The platform can decide which information is revealed about previously submitted bids when an agent accesses the platform. At the end of the auction, the item is sold to the agent who first submitted the highest bid at a price equal to the highest bid submitted by another agent.

DEFINITION 3. A mechanism (X, \mathcal{I}) is a *second-price auction with asynchronous bids* if it satisfies the following. The tree X is such that an agent can submit at most L bids, where $L \geq 1$, within some time interval $[0, T]$ or $[0, T)$, where $T > 0$. Along each terminal history $z \in Z$, at the root, chance picks $\bar{n} \in N$ which determines the set $\bar{N} := \{n \in N \mid n \leq \bar{n}\}$ of agents who take part in the auction. Then, letting $b_n^0 := 0$, we have the following.

- (i) First, chance moves and selects a reserve price r , from some set $R \subset [0, \infty)$.
- (ii) Second, in some order, all agents $n \in \bar{N}$ choose a first access date $\tau_n^1 \in [0, T]$.
- (iii) Then, for all integers $k > 0$, iteratively defining $N^k \subset \bar{N}$ as the set of agents who submitted less than L bids and for whom $\tau_n^k \leq T$:
 - (a) If N^k is nonempty, let \underline{N}^k be the set of all agents $n \in N^k$ with $\tau_n^k \leq \tau_m^k$ for all $m \in N^k$. All $n \in \underline{N}^k$ are called to play in some order. Each submits a bid b_n^k from the set $\{b \in \mathcal{D} \mid b_n^{k-1} \leq b\}$ and selects a next access date $\tau_n^k + 1$ from the set $\{\tau \in [0, T] \cup \{\infty\} \mid \tau_n^k < \tau\}$. For all other agents $n \in \bar{N} \setminus \underline{N}^k$, let $\tau_n^{k+1} := \tau_n^k$ and $b_n^k := b_n^{k-1}$.
If $k = 1$, chance selects a *winning bid* at random among the set of maximum bids b_n^1 submitted by some $n \in \underline{N}^1$. If $k > 1$, let m be the agent who submitted the previous winning bid. If $\max\{b_n^k \mid n \in \underline{N}^k\} < b_m^k$, the new winning bid is b_m^k . Otherwise, chance selects the new winning bid at random from the set of maximum bids b_n^k submitted by some agent $n \in \underline{N}^k \cup \{m\}$. All b_n^k that are not winning are said to be *losing bids*.
 - (c) If N^k is empty, the auction ends. The agent n who submitted the winning bid b_n^{k-1} is assigned the item if and only if $b_n^{k-1} \geq r$. If agent n is assigned the object, she pays $\max\{r, \max_{m \neq n} b_m^{k-1}\}$. All other bidders pay zero.

For all $h \in H$ such that all agents have selected at least one access date, define agent n 's next access date $\tau_n(h)$ as the most recent access date chosen by agent n . If agent n is called to play at h , we refer to $\tau_n(h)$ as agent n 's current access date. The information partition \mathcal{I} is such that no agent ever receives information about the number of agents $\bar{n} = |\bar{N}|$ selected by chance, and:

- (i) Agents select their first access dates τ_n^1 with no information about others' moves.

- (ii) If agent n is called to bid at h , for all $m \neq n$ with $\tau_n(h) \leq \tau_m(h)$, the information set $\mathcal{I}_n(h)$ is such that the only information about agent m 's next access date that agent n observes is that it satisfies $\tau_n(h) \leq \tau_m(h)$.
- (iii) If two agents n have the same next access date $\tau_n(h)$, they are called to play simultaneously whenever they are first called to play along any successor of h .
- (iv) If agent n enters the platform for the k -th time and is revealed the final selling price will be at least β , a bid b_n^k is *admissible* if it is strictly larger than both b_n^{k-1} and β . If she does not submit an admissible bid, others do not observe she entered and receive no information about her bid.⁴²

An asynchronous second-price auction in which *no* information about previous bids is disclosed to bidders corresponds to the standard Vickrey auction, and thus has a dominant strategy equilibrium. In fact, the only information that can be revealed while maintaining dominant strategy incentive compatibility is the reserve price. At the other end, in an asynchronous second-price auction in which *all* information about previous bids is disclosed, truthful bidding is only an ex-post equilibrium. In between these two extremes there are many other disclosure policies that could be chosen by the online platform hosting the auction.

Our next result characterizes the set of second-price auctions with asynchronous bids that are as-if dominant strategy mechanisms. Consider a second-price auction with asynchronous bids, and suppose that n is called to submit a bid at h , on date $\tau_n(h)$: A *previously submitted* bid is a bid submitted along h on some date $\tau < \tau_n(h)$.

THEOREM 2. *A second-price auction with asynchronous bids is an as-if dominant strategy mechanism if and only if the following holds for all β : Until it is certain the final selling price is at least β , no agent observes whether someone else previously submitted a bid of at least β .*

From Theorem 2, we conclude that the restriction imposed by as-if dominance on the information partition of the second-price auction with asynchronous bids concerns only the information that is disclosed about the currently winning bid.⁴³ The

⁴²That is, agents cannot distinguish actions that are equivalent to submitting no bid. Using the alternate definition of as-if dominance from Footnote 30, we could directly assume bidders cannot submit such bids. We reiterate that the present formulation is for expositional convenience only. All results would stay the same under the alternate definition.

⁴³Interestingly—in this setting—the same bound on transparency is necessary and sufficient to ensure the outcomes in undominated strategies of an asynchronous second-price auction are the same as those of a sealed-bid second-price auction (possibly with a reserve). In principle, implementation in undominated strategies requires agents to find the appropriate alternative to exclude each dominated behavior (see, e.g., Jackson et al. (1994)): Conveniently, the as-if dominance property ensures that cognitively limited agents can do so *with a single strategy*. This relationship between as-if dominance and implementation in undominated strategies does not hold for any of our other application settings.

intuition is as follows. Suppose that, at some point in the mechanism, agent n is not currently winning and observes that another agent submitted a bid of at least β . Then, agent n could condition her behavior on the value of β in a way that changes the price at which another bidder wins the object. If it is possible that both the highest losing bid and the reserve price are lower than β , she can do so in a way that breaks as-if dominance. Indeed, for agent n , bidding any value below β when β is above her value is compatible with playing a coarsely optimal strategy, as she is indifferent between submitting any losing bid. This implies that—even if all agents play rationally conditional on their belief that others behave coarsely—the agent who submitted the currently winning bid may not be best responding to agent n .

In contrast, if agent n knows for sure that either the reserve price or the highest losing bid is at least β , then she can only affect the final outcome by submitting a bid above β . Doing so may lead her to win the auction, and hence is never compatible with playing a coarsely optimal strategy if she has a value below β , since she would risk making a negative payoff. Hence, as-if dominance is preserved in this case.

We now show that *any* as-if dominant strategy profile of a second-price auction with asynchronous bids implements the same outcome as the dominant strategy outcome of a standard sealed-bid second-price auction (possibly with a reserve price).

PROPOSITION 7. *All as-if dominant strategy profiles of a second-price auction with asynchronous bids are such that, for all $\theta \in \Theta$, after chance selects a reserve price r :*

- (i) *If $\max_n \theta_n > r$, the item is sold to one of the agents with the highest value, for a price equal to the maximum between r and the highest value among the losing agents.*
- (ii) *If $\max_n \theta_n = r$, the item is either unsold or sold at a price of r to an agent with value r .*
- (iii) *Otherwise, the item is unsold.*

Thus, we can use as-if dominant strategy mechanisms in this setting to implement either the profit maximizing outcome (by setting the optimal reserve price) or the efficient outcome (by not setting a reserve price).

5.2 The Online English Auction

In Theorem 2 above, we characterized all as-if dominant second-price auctions with asynchronous bids. This showed that although as-if dominance restricts the information that can be provided to bidders, it still allows the platform to choose among various information disclosure policies. Given the experimental evidence as well as theoretical arguments that higher transparency may lead to closer conformity to the

predicted outcomes (see Section 3.1), we now introduce the *Online English Auction*, and show it achieves maximal transparency among the set of as-if dominant strategy asynchronous second-price auctions. As we discuss, this format coincides with the default format used by prominent online platforms, such as eBay, providing a rationalization of their choice.

DEFINITION 4. A second-price auction with asynchronous bids is an Online English Auction if its information partition is such that, whenever an agent is called to submit a bid, she observes:

- (i) The reserve price selected by chance.
- (ii) The order of submission and submission dates of all previous admissible bids, and who submitted them.
- (iii) The value of all previously submitted admissible losing bids .
- (iv) Nothing else (but all her previous moves).

COROLLARY 1. *The Online English Auction is an as-if dominant strategy mechanism.*

PROOF. Immediate consequence of Definition 4 and Theorem 2. □

Next, we show that this format achieves maximal transparency within the set of as-if dominant strategy asynchronous second-price auctions.

THEOREM 3. *There is no as-if dominant strategy second-price auction with asynchronous bids that is more transparent than the Online English Auction.*

The Online English Auction is the default auction format used by prominent online platforms, such as eBay.⁴⁴ Thus, Theorem 3 provides a theoretical justification for the popularity of this format: The information provided is such that bidders still have strong strategic incentives to bid truthfully. Subject to the as-if dominant strategy incentive constraint, the chosen information policy then maximizes transparency, which can help bidders understand their optimal strategy and increase trust in the online platform, as we argued in Section 3.

⁴⁴Other platforms using Online English Auctions include Copart USA for cars (<https://www.copart.com/>), auctionet for art and design (<https://auctionet.com/>), and the U.S. General Services Administration auctions (<https://www.gsauctions.gov/>), all accessed on February 10, 2024. For exposition, we abstracted away from the possibility that bidders can revise their maximum bid downwards as long as they remain above the currently displayed price. This is also an as-if dominant strategy mechanism but prevents reliably revealing anything about highest bid by design, only making our characterization easier to prove.

5.3 Shill-Proofness

The asynchronous second-price auctions considered above are designed to resolve the problem of coordinating bidding times. With this in mind, a particularly interesting class of ex-post equilibria are those in which each bidder enters the platform only a single time, and not all bidders enter at the same time.⁴⁵ We call such ex-post equilibria *fully asynchronous*. Fully asynchronous truthful equilibria exist in all second-price auctions. However, in practice, they are only realistic if bidders can trust that their behavior will not be exploited by the platform or by the private sellers of the items.

As argued in Section 3—see Fact 2—increasing the transparency of an asynchronous second-price auction can increase bidders’ confidence that the platform is adhering to the auction rules it specified. Indeed, in the Online English Auction defined above, the winning bidder can always verify that the price they paid indeed corresponds to the bid submitted by another bidder.

A distinct property, shill-proofness, is that the private sellers do not have an incentive to submit bids of their own in an attempt to raise their item’s selling price.

To model this, suppose Chance may (or may not) pick one of the bidders at random and change her preferences such that she now only wants to maximize revenues raised in the auction. Only such a *shill bidder* knows her preferences have changed. A truthful⁴⁶ and fully asynchronous equilibrium is *shill-proof* if a shill bidder cannot raise expected revenues by entering and submitting a positive bid.⁴⁷

PROPOSITION 8. *Let bidders’ types be independent draws from a regular distribution.*

(i) *All truthful and fully asynchronous equilibria of an Online English Auction with optimal reserve price are shill-proof.*

(ii) *Consider a fully asynchronous truthful equilibrium of an asynchronous second-price auction. If there is an on-path history at which a bidder receives more information than she would in an Online English Auction, the equilibrium is not shill-proof.*

⁴⁵In fact, truthfully revealing one’s reservation value once is the strategy recommended by eBay, see <https://www.ebay.com/help/buying/bidding/automatic-bidding?id=4014>, accessed Aug 02, 2024.

⁴⁶We say an equilibrium is truthful if each bidder bids her true value when she enters the auction, unless she knows another bidder already bid weakly more, in which case she may bid in any way that guarantees her to lose the auction.

⁴⁷A probability distribution $\pi \in \Delta(\mathcal{D})$ is *regular* if it has full support and, letting \mathcal{D} be ordered such that $d_k \leq d_\ell$ if and only if $k \leq \ell$, the map

$$k \mapsto d_k - (d_{k+1} - d_k) \frac{1 - \pi(\{d_\ell \mid k \leq \ell\})}{\pi(\{d_k\})}$$

is increasing. If bidders types are independent draws from a Myerson-regular distribution, a reserve price is *optimal* if it maximizes expected revenue in a sealed-bid second-price auction.

That is, in an Online English Auction, bidders can enter the auction at their most convenient time, bid their reservation value and still sellers do not have an incentive to submit shill bids. Revealing any additional information—which would break as-if dominance—would sometimes give the seller a strict incentive to submit shill bids.

While truthful shill-proof equilibria can exist in more transparent asynchronous second-price auctions that are not as-if dominant strategy, they are not fully asynchronous: Bidders would have to be present throughout the entire auction—similarly to a clock auction—defeating the very purpose of the asynchronous format.

6 Other Applications

In this section, we consider two complex application settings—multi-unit auctions and matching problems—for which obviously dominant strategy mechanisms either do not exist, or exist only under strong restrictions on agents’ preferences. Thus, Proposition 4 implies there is necessarily a dilemma between transparency and dominant strategy implementation.

In both settings, experimental evidence shows that sometimes using a more transparent mechanism can achieve closer conformity to its theoretical prediction even if the additional transparency breaks the dominant strategy property. We apply as-if dominance to provide a unified explanation for these findings: We show that the better-performing transparent mechanisms are all as-if dominant strategy. This implies that, to cognitively limited agents, their incentives are indistinguishable from those of a dominant strategy one. The additional transparency, however, can actually increase the rate at which these cognitively limited agents are able to identify their dominated behaviors and hence behave optimally. Indeed, for all applications we consider, the more transparent mechanisms *strictly* increase the number of deviating actions that look obviously dominated for such agents, c.f. Proposition 2. Thus, we propose that these mechanisms can actually lead to better adherence to the theoretical predictions than their dominant strategy counterparts in complex settings with cognitively limited agents.

6.1 Multi-Unit Auctions

Consider the multi-unit analogue of the standard single-unit auction setting. Let $L \geq 2$ denote the maximum number of units of the good that may be allocated. The set of possible outcomes for agent n in this setting is given by $\Omega_n := X_n \times T_n$, where $X_n := \{0, 1, \dots, L\}$ is the allocation space for agent n and T_n is the set of possible transfers for agent n , given by a finite subset of the non-negative real line including

zero. An outcome $(x, t) \in \Omega$ specifies a pair $(x_n, t_n) \in \Omega_n$ for all agents n and satisfies $\sum_n x_n \leq L$. An agent's type $\theta_n \in \Theta_n$ is a vector $(\theta_n^1, \dots, \theta_n^L)$ of non-increasing marginal values for each unit of the good the agent may be allocated. We assume symmetry, i.e., Θ_n is such that there is some $\Delta \subset \mathbb{R}_+^L$ with at least two elements such that, for all n , $\theta_n \in \Theta_n$ if and only if $\theta_n^\ell \in \Delta$ for all ℓ and $\theta_n^\ell \leq \theta_n^{\ell-1}$ for all $\ell \geq 2$. Finally, agent's utility is quasi-linear, i.e., $(x', t') \leq_{\theta_n} (x'', t'')$ if and only if $\sum_{\ell \leq x'_n} \theta_n^\ell - t'_n \leq \sum_{\ell \leq x''_n} \theta_n^\ell - t''_n$.

We say a multi-unit allocation problem is *proper* if there are at least two agents, two marginal values and two units.⁴⁸ We will be interested in mechanisms which implement the Vickrey outcome.⁴⁹ Our first result in this section shows that if the setting is proper, then it is not possible to implement the Vickrey outcome in obviously dominant strategies (c.f. [Gkatzelis et al. \(2017\)](#)).⁵⁰

PROPOSITION 9. *For proper multi-unit allocation problems with transfers, there is no mechanism that implements the Vickrey outcome in obviously dominant strategies.*

Together, [Theorem 9](#) and [Proposition 4](#) imply that a designer who wants to implement the Vickrey outcome in this setting will necessarily face a trade-off between the strength of the incentives and transparency. The canonical dominant strategy mechanism in this setting is the standard multi-unit sealed-bid Vickrey auction, which is static and hence provides minimal transparency. However, [Ausubel \(2004\)](#) proposed a class of dynamic auctions—multi-unit analogues of the standard ascending clock auction—which offer higher transparency than a static mechanism.

6.1.1 Ausubel Tree

We now give a brief description of the game tree.⁵¹ First, chance selects a rationing priority order among all possible permutations of agents. Then, the discrete-time auction proceeds in $r \geq 1$ rounds. Each agent n initially demands all available units $x_n^0 = L$. Each round r then works as follows: First, the auctioneer announces the round's unit price $p^r := \Delta_r$ to all active agents,⁵² where Δ_r denotes the r -th smallest element of Δ . Then, all active bidders communicate to the auctioneer how many

⁴⁸With one good, it reduces to the standard single-unit setting considered in [Section 5](#), in which a obviously dominant strategy clock auction can implement the Vickrey outcome—see [Li \(2017b\)](#).

⁴⁹We remark that the Vickrey outcome not only allocates units efficiently, but also specify a payment vector where each agent payment is equal to the externality she imposes on others.

⁵⁰We provide a direct proof. The result also follows from the fact that OSP mechanisms are weakly group-strategyproof ([Li, 2017b](#)), combined with [Proposition 7.3.](#) in [Gkatzelis et al. \(2017\)](#) which shows no weakly group-strategyproof mechanism can guarantee better than a $\sqrt{2}$ -approximation to optimal social welfare in a setting with ≥ 2 agents who have non-increasing marginal values over ≥ 2 units.

⁵¹A formal definition can be found in [Appendix D](#) of [Nagel and Saitto \(2024\)](#).

⁵²An agent is active in round $r \geq 1$ if they demanded at least one unit in round $r - 1$.

(additional) units they want to drop from their demand. Next, based on each agent n 's round r demand, x_n^r , the auctioneer computes the total demand at the end of the round, $\sum_n x_n^r$. Whenever there is excess demand, i.e. $\sum_n x_n^r > L$, the auctioneer then assigns $\chi_n^r := \max\{0, L - \sum_{n' \neq n} x_{n'}^r\} - \max\{0, L - \sum_{n' \neq n} x_{n'}^{r-1}\}$ units to agent n at price p^r . If additionally $r < |\Delta| - 1$, the auction proceeds to the next round. If $r = |\Delta| - 1$, the the auction ends and unassigned units are rationed according to agents' unfulfilled demands and the priority order, at price equal to the highest possible marginal value.⁵³ Whenever there is no excess demand at the end of round $r \leq |\Delta| - 1$, the auction also ends. All agents are assigned the units which exactly fulfill their demand x_n^r at price p^r . Then, unassigned units are rationed according to the priority order at price p^r , under the condition agent n cannot be assigned more than $x_n^{r-1} - x_n^r$ additional units. We call a bidder's strategy *sincere bidding* if it prescribes to drop the ℓ -th unit in round r if and only if $\theta_n^\ell \leq p^r$.

6.1.2 Information Partitions

We call mechanism a *Ausubel auction with private dropouts* if the tree is as described above and the information partition is such that bidders communicate how many units they would like to drop privately to the auctioneer within a round and, at the end of a round, a bidder is not informed how many units the other bidders dropped. We call a mechanism a *Ausubel auction with public dropouts* if the information partition is such that bidders communicate how many units they would like to drop privately to the auctioneer within a round but, at the end of a round, it is revealed to bidders how many units the others dropped. In the Ausubel auction with private dropouts, bidding sincerely is a dominant strategy for a bidder. In the Ausubel auction with public dropouts, it is only an ex-post equilibrium, and the unique outcome of iterated elimination of dominated strategies.⁵⁴ However, our next result shows that the Ausubel auction with public dropouts maintains the incentive as straightforward as in its version with private dropouts, in the sense that it is an as-if dominant strategy mechanism.

THEOREM 4. *In the Ausubel auction with public dropouts, bidding sincerely is an as-if dominant strategy.*

⁵³That is, unassigned units are rationed following the priority order until supply is exhausted, subject to assigning each agent n at most $x_n^{|\Delta|-1} - \sum_{r \leq |\Delta|-1} \chi_n^r$ additional units. Intuitively, if reached, the last round is played by a computer that knows all units left should be dropped at $p = |\Delta|$. This is a minor adaptation of the definition by Ausubel (2004), who instead assumes there is a round $|\Delta| + 1$ which is never reached in equilibrium.

⁵⁴See Theorem 1 and 2 in Ausubel (2004) and Okamoto (2018) for proofs of these statements.

It is straightforward to see that the Ausubel auction with public dropouts is strictly more transparent than the one with private dropouts. Together with Theorem 4, this provides a theoretical explanation of existing experimental evidence: [Kagel and Levin \(2009\)](#) compare the performance of the standard sealed-bid multi-unit Vickrey auction, the Ausubel auction with private dropouts, and the one with public dropouts. The authors find that—against the theory of dominant strategy implementation—the Ausubel auction with public dropouts achieves significantly better performance than the Ausubel auction with private dropouts and the sealed-bid Vickrey auction. For example, in the first 12 auctions with $L = 2$ units, truthful bidding on the higher-valued unit is 18% for the sealed-bid Vickrey auction, 43% for the Ausubel auction with private dropouts, but 83% when dropouts are public. The results are similar for the lower-valued unit, although the gap between the auction format becomes slightly smaller as bidders become more experienced. [Kagel and Levin \(2009\)](#) suggest the additional transparency by the dropout information simplifies the bidding task even at the expense of the solution concept. We provide a theoretical justification: The Ausubel auction with public dropouts is more transparent and an as-if dominant strategy mechanism. Thus, cognitively limited agents have what looks like a dominant strategy to them, and can benefit from the additional transparency, which may help them spot their dominated behaviors more easily. Moreover, it is ex-post incentive compatible for sophisticated agents to adhere to the as-if dominant strategy prediction—even if they are worried that some of the other agents is cognitively limited. This does not extend to the more complex ascending auction formats for heterogeneous goods, such as those considered by [Ausubel \(2006b\)](#) and [Gul and Stacchetti \(2000\)](#), which are not as-if dominant strategy mechanisms.

6.2 Two-Sided Matching Problems

Our next application is to two-sided matching problems. Throughout, we refer to agents n as students, and let $J := \{1, \dots, L\}$ denote the set of schools a student may be assigned to. A student may also be unassigned, denoted 0. Thus, the set of possible outcomes for an agent n is $\Omega_n = L \cup \{0\}$. Schools have capacities $K_j \geq 1$. Hence, a social outcome $\omega \in \times_n \Omega_n$ is feasible if the matching ω is such that at most K_j students are assigned to school j . Each school $j \in J$ has a strict ranking over the set of students and 0, where ranking a agent n below 0 means that student is unacceptable to the school. Finally, we model schools as non-strategic and always assume their preference ranking is drawn by nature at the beginning of the mechanism.⁵⁵

⁵⁵This is reasonable in context where the preferences of one side of the market are determined, e.g., as a function of zip codes, siblings at the school, and a random tie breaker.

It is well-known the (student-proposing) Deferred Acceptance Algorithm (DAA) can be used to find the student-optimal stable outcome in this setting (Gale and Shapley, 1962).⁵⁶ While the stable outcome can be reached with a dominant strategy mechanism, Ashlagi and Gonczarowski (2018) show that no obviously dominant strategy mechanism can do so. Hence, as in the previous setting, there is an inevitable dilemma between dominant strategy incentive compatibility and transparency.

6.2.1 The DA Algorithm

The DA algorithm can be described as follows. All students start unassigned. In each round, students that are not tentatively assigned to a school apply to one school among those they have not applied to yet, or choose to leave the mechanism and thus remain unassigned. At the end of each round, given both sides' preferences, the clearinghouse tentatively matches students to schools: Students that exceed a school's capacity or are unacceptable are rejected. A round is the final round if at the end of it, each student is either assigned to a school or has left the mechanism.

6.2.2 Static and Sequential DAA

In the *static DAA*, each student reports a ranking without any information about others' moves, and a clearinghouse implements the student-optimal stable outcome by running the above algorithm. We refer to the dynamic version of this mechanism, in which students actively participate by applying to schools one by one as the *sequential DAA* if the information partition is such that (i) schools' preferences are not revealed to students and (ii) students only observe their own tentative assignment at the end of each round.⁵⁷

Being truthful is a dominant strategy for each student in the static DAA (Dubins and Freedman, 1981; Roth, 1982), but only an ex-post equilibrium in the sequential DAA (e.g., Bó and Hakimov, 2022; Klijn et al., 2019).⁵⁸ Nevertheless, and against the theoretical prediction of dominant strategy implementation, several experiments find significantly higher fractions of truth-telling in the sequential version of the DAA than in the static version (Bó and Hakimov, 2020; Klijn et al., 2019; Hakimov and Raghavan, 2022). Once again, our approach can provide an explanation for this experimental

⁵⁶A *stable outcome* in this setting is a matching in which all students and schools that are assigned to each other are acceptable to each other, and there is no pair of a school and a student would both rather be matched to each other than to (one of) their partners in the matching. The student-optimal stable matching is the stable matching most preferred by all students.

⁵⁷To compare transparency, we represent the static DAA and sequential DAA using the same tree.

⁵⁸The sequential DAA only provides information on a student's own outcome each round, so the transparency that can be provided under dominant strategy incentives is very limited in this context.

evidence: The sequential DAA is an as-if dominant strategy mechanism, and hence provides transparency while maintaining strategic incentives that—to agents with limited ability to reason contingently—resemble those of a dominant strategy one.

THEOREM 5. *In the sequential Deferred Acceptance Algorithm, reporting truthfully is the unique as-if dominant strategy.*

Providing students with maximal transparency in this context may be not be feasible since it may reveal sensitive information. However, depending on the structure of schools’ preferences, it may be feasible to inform students (i) at which schools they will be admitted for sure, (ii) at which schools it is uncertain whether they will be admitted and (iii) at which schools they will not be admitted for sure. Relatedly, the experiments in [Bó and Hakimov \(2020\)](#) and [Hakimov and Raghavan \(2022\)](#) consider versions of the sequential DAA in which information about schools’ current cutoffs are provided at the end of each round, and show that these mechanisms, which are as-if dominant, outperform their dominant strategy counterparts.⁵⁹ Moreover, the experiment in the latter paper highlights the important role of this additional transparency for students’ ability to verify the outcome.

6.3 One-Sided Matching Problems

The last application for this section are housing allocation problems. Each agent n owns a house j_n . The type space of an agent n is the set of all strict preference orderings over the set of houses \mathcal{J} . An outcome $\omega \in \Omega := \mathcal{J}^N$ assigns a house to each agent, such that no two agents are assigned the same house. Given a type profile θ , an outcome ω is *core stable* if there is no coalition $\mathcal{C} \subset \mathcal{N}$ such that there is an outcome $\omega' \in \Omega$ with $\omega'_n \in \{j_k \mid k \in \mathcal{C}\}$ and $\omega_n <_{\theta_n} \omega'_n$ for all $n \in \mathcal{C}$.

6.3.1 Top Trading Cycles

For all type profiles, the *Top Trading Cycles (TTC)* algorithm always implements the unique core stable outcome, and works as follows. In round k , each not-yet-assigned agent points towards her k –th most preferred house. Then, all sets of unassigned agents—interpreted as vertices—together with their pointing decisions—interpreted as directed edges—that form a cycle are identified. If an agent is part of one of these sets, she is assigned the house she pointed at. The algorithm continues until all agents are assigned to an house.

⁵⁹If students can only apply to schools at which they meet the cutoff changes the game tree and reduces students’ strategy spaces, so an improvement in performance cannot solely be attributed to the additional transparency.

The *static TTC algorithm* is such that all agents are asked to report their preference order with no information about what others do. The *sequential TTC algorithm* is such that, after all (still unassigned) agents simultaneously report their k -th most preferred house, the k -th round is run and agents are informed of its outcome. Reporting truthfully is a dominant strategy equilibrium in the static TTC, while it is not in the more transparent sequential TTC. However, we have the following result.

THEOREM 6. *In the sequential Top Trading Cycles algorithm, reporting truthfully is the unique as-if dominant strategy.*

Hence, the sequential TTC algorithm provides higher transparency and maintains as-if dominant incentives. Consistently, [Hakimov and Raghavan \(2022\)](#) provide evidence that truthful reporting is higher in the sequential TTC algorithm is higher than in its static counterpart.

7 Perfection and Weak Dominance Solvability

In this section, we identify a class of as-if dominant strategy mechanisms—called *perfect as-if dominant strategy mechanisms*. They are weak dominance solvable mechanisms whose predictions are robust to the possibility that some of the agents are cognitively limited and hence, while able to recognize their own dominated behaviors, are not able to engage in the higher-order thinking necessary to operate iterated deletion of weakly dominated strategies.

To operationalize the refinement, we then provide a set of easy-to-check sufficient conditions which hold for most of our applications. This allows us to prove the novel results that the sequential Deferred Acceptance Algorithm and the sequential Top Trading Cycle are weak dominance solvable. Throughout this section, we always assume the type space is finite and restrict attention to finite mechanisms.

7.1 Weak Dominance Solvability

To introduce our refinement, we first briefly define weak dominance solvability.

Fix a mechanism (X, \mathcal{I}) . Two behaviors $\sigma'_n, \sigma''_n \in \Sigma_n$ are *equivalent* if $\omega(\sigma'_n, \sigma_{-n}) = \omega(\sigma''_n, \sigma_{-n})$ for all $\sigma_{-n} \in \Sigma_{-n}$. That is, these behaviors always lead to the same outcome regardless of the other's behavior. Next, for all $I_n \in \mathcal{I}_n$, we say that two behaviors $\sigma'_n, \sigma''_n \in \Sigma_n$ are *I_n -equivalent* if $\omega(\sigma'_n, \sigma_{-n} \mid h) = \omega(\sigma''_n, \sigma_{-n} \mid h)$ for all $\sigma_{-n} \in \Sigma_{-n}$ and $h \in I_n$. That is, these behaviors lead to the same outcome regardless of the other's behavior, conditional on having reached I_n . We treat two strategies $s'_n, s''_n \in S_n$ such

that $s'_n(\theta_n)$ is equivalent to $s''_n(\theta_n)$ for all $\theta_n \in \Theta_n$ as the same.⁶⁰

A set $S' \subseteq S$ is a *strategy rectangle* if $S' := \times_n S'_n$ with $S'_n \subseteq S_n$ for all agents n .

Given a strategy rectangle S' , a strategy $s_n \in S'_n$ is *dominated with respect to S'* if there is a strategy $s'_n \in S'_n$ such that

$$\omega(s_n(\theta_n), s_{-n}(\theta_{-n}), \sigma_c) \leq_{\theta_n} \omega(s'_n(\theta_n), s_{-n}(\theta_{-n}), \sigma_c)$$

for all $s_{-n} \in S'_{-n}$, $\sigma_c \in \Sigma_c$ and $\theta \in \Theta$, with at least one inequality holding strict. Observe that s_n is dominated with respect to S if and only if it is a dominated strategy.

A strategy rectangle S' can be obtained by iterated deletion of dominated strategies if there exists a (finite) sequence of strategy rectangles $\{S^k\}_k$ with $S^{k+1} \subseteq S^k$ for all k , such that: $S^1 = S$, $\bigcap_k S^k = S'$, and, whenever a strategy s_n is in $S^k_n \setminus S^{k+1}_n$, then s_n is dominated with respect to S^k .

A strategy rectangle $S^\infty \subseteq S$ is *admissible* if (i) no strategy $s_n \in S^\infty_n$ is dominated with respect to S^∞ , and (ii) the set S^∞ can be obtained by iterated deletion of dominated strategies.

DEFINITION 5. A mechanism is weak dominance solvable if there is a strategy profile s such that $\{s\}$ is the unique admissible strategy rectangle.

We refer to the unique admissible strategy profile of a weak dominance solvable mechanism as its (*weak dominance*) *solution*.

7.2 Perfect As-If Dominant Strategy Mechanisms

We now define perfect as-if dominant strategy mechanisms. Recall that a strategy s_n is *coarsely optimal* if $\omega(\sigma_n, \sigma_{-n}) \leq_{\theta_n} \omega(s_n(\theta_n), \sigma_{-n})$ for all types $\theta_n \in \Theta_n$, behaviors $\sigma_n \in \Sigma_n$, and coarse behavior profiles $\sigma_{-n} \in \Sigma_{-n}(\mathcal{I})$.

DEFINITION 6. A finite mechanism is a perfect as-if dominant strategy mechanism if:

- (i) Each agent has a unique coarsely optimal strategy s_n .
- (ii) The mechanism is weak dominance solvable, and $(s_n)_n$ is its solution.

Intuitively, perfect as-if dominant strategy mechanisms are weak dominance solvable mechanisms that are robust to the presence of agents who, while rational, are cognitively limited as defined in Section 4. This means that—as long as it is common knowledge among the fully sophisticated agents that all agents are either fully sophisticated themselves or cognitively limited—their predictions no longer rely on

⁶⁰Formally, we consider the quotient space of S_n with respect the equivalence relation defined above.

equilibrium arguments. Indeed, cognitively limited agents have a strategy that looks dominant to them, which implies their behavior can be predicted by rationality alone. Moreover, an agent with perfect ability to engage in higher order thinking can solve the mechanism by iteratively eliminating dominated strategies. Importantly, even if she is uncertain whether some of the other agents are cognitively limited, this does not pose a problem: She knows that the cognitively limited agents will play their unique coarsely optimal strategy.

The above definition constitutes a refinement of as-if dominant strategy mechanisms. As we show later, it is satisfied by most of the applications we consider.

PROPOSITION 10. *A perfect as-if dominant strategy mechanism is as-if dominant strategy.*

Observe that while all perfect as-if dominant strategy mechanisms are by definition weak dominance solvable, the converse is not true. For example, the well-known beauty contest game is weak dominance solvable, but not even an as-if dominant strategy mechanism (see, e.g., Nagel (1995)).

7.3 Sufficient Conditions for Perfection

Since it is usually difficult to prove directly that a given mechanism is weak dominance solvable, we provide sufficient conditions for a mechanism to be perfect as-if dominant strategy, and hence weak dominance solvable. This allows us to operationalize our refinement and to apply it to the multi-unit auction and matching settings considered above.

To this end, we introduce a regularity condition for mechanisms which we call a *round structure*, and which is satisfied by all the applications in the preceding sections. Intuitively, a mechanism has a round structure if the information revealed by the information partition is such that, whenever an agent is called to play, she knows how far into the mechanism she is called to play at. This is always the case for finite mechanisms in which information sets “do not cross”: If $h \in I$ precedes $h' \in I'$, then there is no $h'' \in I'$ that precedes some $h''' \in I$.

DEFINITION 7. A mechanism (X, \mathcal{I}) has a *round structure* if there exists a finite partition \mathcal{R} of the non-terminal histories $H \setminus Z$ that satisfies the following:

- (i) If a round $R \in \mathcal{R}$ and some $I \in \mathcal{I}$ are not disjoint, then $I \subseteq R$.
- (ii) Within each round, play is simultaneous: For all $R \in \mathcal{R}$ and $I \subseteq R$, for all $h \in I$: All terminal histories z with $h \prec z$ pass through the same information sets $I' \in R$, and encounter them in the same order.

- (iii) The binary relation $\preceq_{\mathcal{R}}$ over \mathcal{R} defined by $R' \preceq_{\mathcal{R}} R''$ if and only if $h' \preceq h''$ for some $h' \in R'$ and $h'' \in R''$ is acyclic.

The following results gives sufficient conditions for a mechanism to be perfect as-if dominant strategy which are easy to verify in any given application setting. Recall that given a mechanism (X, \mathcal{I}) , the symbol \mathcal{J} denotes the coarsest information partition that can be defined on the tree X .

THEOREM 7. *Suppose the type space Θ is finite and consider an as-if dominant strategy mechanism (X, \mathcal{I}) that satisfies the following conditions.*

(i) *The mechanism is finite and has a round structure.*

(ii) *Each agent has a coarse as-if dominant strategy s_n , such that:*

- (a) *For all information sets $I_n \in \mathcal{I}_n$, types $\theta_n \in \Theta_n$, and strategies $\sigma_n \in \Sigma_n$ that are not I_n -equivalent to $s_n(\theta_n)$,*

$$\omega(\sigma_n, \sigma_{-n} \mid h) \leq_{\theta_n} \omega(s_n(\theta_n), \sigma_{-n} \mid h)$$

for all $h \in I_n$ and coarse $\sigma_{-n} \in \Sigma_{-n}(\mathcal{J})$, with at least one strict inequality.

- (b) *For all coarse $\sigma_n \in \Sigma_n(\mathcal{J})$, there is $\theta_n \in \Theta_n$ such that $s_n(\theta_n)$ is equivalent to σ_n .*

Then, the mechanism is a perfect as-if dominant strategy mechanism.

Intuitively, condition (ii) requires that each agent has an as-if dominant strategy that (a) remains the unique coarsely optimal continuation strategy even after the agent (mistakenly) deviates from it, and (b) is such that all coarse behaviors are coarsely optimal for some type.

The assumptions of the above theorem are satisfied by the Ausubel Auction with public dropouts, the sequential Deferred Acceptance algorithm, and the sequential Top Trading cycle algorithm.⁶¹ As a corollary, we have the following result, which demonstrates perfect as-if dominant strategy mechanisms are interesting also from a technical perspective, as they allow us to quickly identify some mechanisms that are weak dominance solvable.

COROLLARY 2. *The as-if dominant strategy profiles of the Ausubel Auction with public dropouts, the sequential Deferred Acceptance Algorithm and the sequential Top Trading Cycle algorithm are the weak dominance solutions of their respective mechanisms.*

⁶¹The only condition the Online English Auction fails to satisfy in a non-remediable way is the uniqueness of the as-if dominant strategy profile.

The first part of this corollary is proven directly as Theorem 2 in [Ausubel \(2004\)](#). Our Theorem 7 reveals that Ausubel’s result is effectively an immediate consequence of his Theorem 1, which proves that in the Ausubel Auction with private dropouts—the coarsest information partition for its tree—sincere bidding is a dominant strategy, even after a deviation. To the best of our knowledge, the statements concerning the sequential DA and TTC algorithms are novel.

This implies that even from the traditional theoretical perspective the three mechanisms from Corollary 2—which have been shown to outperform their dominant strategy counterparts (see Section 6)—were not only ex-post incentive compatible but in fact weak dominance solvable mechanisms. However, many experiments have found that subjects are typically not able to engage in several steps of elimination of even strictly dominated strategies (e.g., [McKelvey and Palfrey \(1992\)](#); [Nagel \(1995\)](#)). We propose that instead condition (i) of perfect as-if dominant strategy mechanisms may help account for their superior performance: These mechanisms all offer cognitively limited agents a strategy that looks dominant to them, and provide transparency that can help them recognize their dominated behaviors, helping to guide them toward the weak dominance solution.

8 Conclusion

In this paper, we show achieving dominant strategy incentive compatibility severely limits the possibility of a designer to increase a mechanism’s transparency. However, both theoretical arguments and existing experimental evidence show that increasing transparency of a mechanism’s extensive form can sometimes increase the reliability of its predictions—even if this breaks dominant strategy incentive compatibility.

To help predict when this is the case, we define as-if dominant strategy mechanisms. We argue they allow designers to provide higher transparency that can help cognitively limited agents to identify their dominated behaviors, while being indistinguishable from dominant strategy mechanisms for them. Moreover, they make it ex-post incentive compatible for sophisticated agents to adhere to the as-if dominant strategy prediction—even if they think that some other agent is cognitively limited.

Our approach provides a unified explanation for experimental evidence in multiple settings, from matching problems to multi-unit auctions ([Kagel and Levin, 2009](#); [Bó and Hakimov, 2020](#); [Bó and Hakimov, 2023](#); [Hakimov and Raghavan, 2022](#); [Klijn et al., 2019](#)): We show the sequential DA and TTC algorithms and the Ausubel Auction with public dropouts are all (perfect) as-if dominant strategy mechanisms.

Applying our approach to the setting of online auctions, we characterize the class

of as-if dominant strategy second-price auctions with decentralized bids, and show that the Online English Auction—the format used by prominent online platforms such as eBay—achieves maximal transparency within this class: Any additional transparency would break as-if dominance.

Finally, we define a refinement, perfect as-if dominant strategy mechanisms. These are weak dominance solvable mechanisms whose predictions are once again robust to the possible presence of cognitively limited. We provide easy-to-verify sufficient conditions for as-if dominant strategy mechanisms to be perfect, and show that they are satisfied by most applications we consider. This lets us apply our results to conclude that the sequential versions of the DA and TTC algorithms are both weak dominance solvable mechanisms.

References

- Akbarpour, M. and Li, S. (2020). Credible auctions: A trilemma. *Econometrica*, 88(2):425–467.
- Arribillaga, R. P., Massó, J., and Neme, A. (2023). All sequential allotment rules are obviously strategy-proof. *Theoretical Economics*, 18(3):1023–1061.
- Arrow, K. J. (1951). Alternative approaches to the theory of choice in risk-taking situations. *Econometrica*, 19(4):404–437.
- Ashlagi, I. and Gonczarowski, Y. A. (2018). Stable matching mechanisms are not obviously strategy-proof. *Journal of Economic Theory*, 177:405–425.
- Ausubel, L. M. (2004). An efficient ascending-bid auction for multiple objects. *American Economic Review*, 94(5):1452–1475.
- Ausubel, L. M. (2006a). An efficient dynamic auction for heterogeneous commodities. *American Economic Review*, 96(3):602–629.
- Ausubel, L. M. (2006b). An efficient dynamic auction for heterogeneous commodities. *The American Economic Review*, 96:602–629.
- Azevedo, E. M. and Leshno, J. D. (2016). A supply and demand framework for two-sided matching markets. *Journal of Political Economy*, 124(5):1235–1268.
- Bade, S. and Gonczarowski, Y. A. (2016). Gibbard-satterthwaite success stories and obvious strategyproofness. *CoRR*.
- Bergemann, D. and Morris, S. (2005). Robust mechanism design. *Econometrica*, 73(6):1771–1813.
- Bergemann, D. and Morris, S. (2008). Ex post implementation. *Games and Economic Behavior*, 63(2):527–566.
- Bernheim, B. D. (1984). Rationalizable strategic behavior. *Econometrica*, 52(4):1007–1028.
- Bó, I. and Hakimov, R. (2020). Iterative Versus Standard Deferred Acceptance: Experimental Evidence. *The Economic Journal*, 130(626):356–392.
- Bó, I. and Hakimov, R. (2022). The iterative deferred acceptance mechanism. *Games and Economic Behavior*, 135:411–433.

- Bodoh-Creed, A. L., Boehnke, J., and Hickman, B. (2020). How Efficient are Decentralized Auction Platforms? *The Review of Economic Studies*, 88(1):91–125.
- Brandenburger, A., Friedenberg, A., and Keisler, H. J. (2008). Admissibility in games. *Econometrica*, 76(2):307–352.
- Breitmoser, Y. and Schweighofer-Kodritsch, S. (2022). Obviousness around the clock. *Experimental Economics*, 25(2):483–513.
- Bó, I. and Hakimov, R. (2023). Pick-an-object mechanisms.
- Börgers, T. and Li, J. (2019). Strategically simple mechanisms. *Econometrica*, 87(6):2003–2035.
- Cason, T. N., Saijo, T., Sjöström, T., and Yamato, T. (2006). Secure implementation experiments: Do strategy-proof mechanisms really work? *Games and Economic Behavior*, 57(2):206–235.
- Chen, Y. and Sönmez, T. (2006). School choice: an experimental study. *Journal of Economic Theory*, 127(1):202–231.
- Chew, S. H. and Wang, W. (2022). Generalizing obvious dominance using the sure-thing principle. Working paper.
- Cohen, S. and Li, S. (2023). Sequential cursed equilibrium.
- Cramton, P. (1998). Ascending auctions. *European Economic Review*, 42(3):745–756.
- Crawford, V. P. (2021). Efficient mechanisms for level-k bilateral trading. *Games and Economic Behavior*, 127:80–101.
- Crawford, V. P., Costa-Gomes, M. A., and Iriberry, N. (2013). Structural models of nonequilibrium strategic thinking: Theory, evidence, and applications. *Journal of Economic Literature*, 51(1):5–62.
- de Clippel, G., Saran, R., and Serrano, R. (2019). Level-k Mechanism Design. *Review of Economic Studies*, 86(3):1207–1227.
- Dubins, L. E. and Freedman, D. A. (1981). Machiavelli and the gale-shapley algorithm. *The American Mathematical Monthly*, 88(7):485–494.
- Echenique, F., Wilson, A. J., and Yariv, L. (2016). Clearinghouses for two-sided matching: An experimental study. *Quantitative Economics*, 7(2):449–482.

- Elkind, E. (2007). Designing and learning optimal finite support auctions. In *Proceedings of the Eighteenth Annual ACM-SIAM Symposium on Discrete Algorithms, SODA '07*, page 736–745, USA. Society for Industrial and Applied Mathematics.
- Esponda, I. and Vespa, E. (2014). Hypothetical thinking and information extraction in the laboratory. *American Economic Journal: Microeconomics*, 6(4):180–202.
- Ferraioli, D. and Ventre, C. (2022). Obvious strategyproofness, bounded rationality and approximation. *Theory of Computing Systems*, 66(3):696–720.
- Forsythe, R., Horowitz, J. L., Savin, N., and Sefton, M. (1994). Fairness in simple bargaining experiments. *Games and Economic Behavior*, 6(3):347–369.
- Gale, D. and Shapley, L. S. (1962). College admissions and the stability of marriage. *The American Mathematical Monthly*, 69(1):9–15.
- Gkatzelis, V., Markakis, E., and Roughgarden, T. (2017). Deferred-acceptance auctions for multiple levels of service. In *Proceedings of the 2017 ACM Conference on Economics and Computation, EC '17*, page 21–38, New York, NY, USA. Association for Computing Machinery.
- Glazer, J. and Rubinstein, A. (1996). An extensive game as a guide for solving a normal game. *Journal of Economic Theory*, 70(1):32–42.
- Gorelkina, O. (2018). The expected externality mechanism in a level-k environment. *International Journal of Game Theory*, 47(1):103–131.
- Grigoryan, A. and Möller, M. (2023). A theory of auditability for allocation and social choice mechanisms.
- Gul, F. and Stacchetti, E. (2000). The english auction with differentiated commodities. *Journal of Economic Theory*, 92(1):66–95.
- Hakimov, R. and Raghavan, M. (2022). Improving transparency in school admissions: Theory and experiment. Working paper. available at ssrn: <https://ssrn.com/abstract=3572020>.
- Hasker, K. and Sickles, R. (2010). ebay in the economic literature: Analysis of an auction marketplace. *Review of Industrial Organization*, 37(1):3–42.
- Jackson, M. O., Palfrey, T. R., and Srivastava, S. (1994). Undominated nash implementation in bounded mechanisms. *Games and Economic Behavior*, 6(3):474–501.

- Jehiel, P. (2005). Analogy-based expectation equilibrium. *Journal of Economic Theory*, 123(2):81–104.
- Kagel, J. H. and Levin, D. (1993). Independent private value auctions: Bidder behaviour in first-, second- and third-price auctions with varying numbers of bidders. *The Economic Journal*, 103(419):868–879.
- Kagel, J. H. and Levin, D. (2009). Implementing efficient multi-object auction institutions: An experimental study of the performance of boundedly rational agents. *Games and Economic Behavior*, 66(1):221–237.
- Klijn, F., Pais, J., and Vorsatz, M. (2019). Static versus dynamic deferred acceptance in school choice: Theory and experiment. *Games and Economic Behavior*, 113:147–163.
- Kneeland, T. (2022). Mechanism design with level-k types: Theory and an application to bilateral trade. *Journal of Economic Theory*, 201:105421.
- Kohlberg, E. and Mertens, J.-F. (1986). On the strategic stability of equilibria. *Econometrica*, 54(5):1003–1037.
- Leshno, J. D. and Lo, I. (2020). The Cutoff Structure of Top Trading Cycles in School Choice. *The Review of Economic Studies*, 88(4):1582–1623.
- Li, J. and Dworczak, P. (2022). Are simple mechanisms optimal when agents are unsophisticated? Working paper.
- Li, S. (2017a). Obvious ex post equilibrium. *The American Economic Review*, 107(5):230–234.
- Li, S. (2017b). Obviously strategy-proof mechanisms. *American Economic Review*, 107(11):3257–87.
- Luce, R. and Raiffa, H. (1957). *Games and Decisions: Introduction and Critical Survey*. Dover Books on Mathematics. Dover Publications.
- Mackenzie, A. (2020). A revelation principle for obviously strategy-proof implementation. *Games and Economic Behavior*, 124:512–533.
- Mackenzie, A. and Zhou, Y. (2022). Menu mechanisms. *Journal of Economic Theory*, 204:105511.
- Mandal, P. and Roy, S. (2022). Obviously strategy-proof implementation of assignment rules: A new characterization. *International Economic Review*, 63(1):261–290.

- Masuda, T., Mikami, R., Sakai, T., Serizawa, S., and Wakayama, T. (2022). The net effect of advice on strategy-proof mechanisms: an experiment for the vickrey auction. *Experimental Economics*, 25(3):902–941.
- McAfee, R. P. and McMillan, J. (1987). Auctions and bidding. *Journal of Economic Literature*, 25(2):699–738.
- McKelvey, R. D. and Palfrey, T. R. (1992). An experimental study of the centipede game. *Econometrica*, 60(4):803–836.
- Milgrom, P. (1989). Auctions and bidding: A primer. *Journal of Economic Perspectives*, 3(3):3–22.
- Milgrom, P. and Segal, I. (2020). Clock auctions and radio spectrum reallocation. *Journal of Political Economy*, 128(1):1–31.
- Milgrom, P. R. and Segal, I. R. (2017). Designing the US Incentive Auction. In Bichler, M. and Goeree, J. K., editors, *Handbook of Spectrum Auction Design*, chapter 37, pages 803–812. Cambridge University Press, Cambridge.
- Myerson, R. B. (1981). Optimal auction design. *Mathematics of Operations Research*, 6(1):58–73.
- Nagel, L. and Saitto, R. (2024). A measure of complexity for strategy-proof mechanisms. Working paper.
- Nagel, R. (1995). Unraveling in guessing games: An experimental study. *The American Economic Review*, 85(5):1313–1326.
- Niederle, M. and Vespa, E. (2023). Cognitive limitations: Failures of contingent thinking. *Annual Review of Economics*, 15(1):307–328.
- Okamoto, N. (2018). An efficient ascending-bid auction for multiple objects: Comment. *The American Economic Review*, 108(2):555–560.
- Pathak, P. A. (2017). *What Really Matters in Designing School Choice Mechanisms*, page 176–214. Econometric Society Monographs. Cambridge University Press.
- Pearce, D. G. (1984). Rationalizable strategic behavior and the problem of perfection. *Econometrica*, 52(4):1029–1050.
- Pycia, M. and Troyan, P. (2023a). Obviously strategyproof mechanisms in general environments. Working paper.

- Pycia, M. and Troyan, P. (2023b). A theory of simplicity in games and mechanism design. *Econometrica*, 91(4):1495–1526.
- Pycia, M. and Ünver, M. U. (2020). Arrowian efficiency and auditability in the allocation of discrete resources. Cepr discussion paper no. dp15377, available at ssrn: <https://ssrn.com/abstract=3723566>.
- Rees-Jones, A. (2018). Suboptimal behavior in strategy-proof mechanisms: Evidence from the residency match. *Games and Economic Behavior*, 108:317–330.
- Roth, A. E. (1982). The economics of matching: Stability and incentives. *Mathematics of Operations Research*, 7(4):617–628.
- Rubinstein, A. (1986). Finite automata play the repeated prisoner’s dilemma. *Journal of Economic Theory*, 39(1):83–96.
- Saran, R. (2016). Bounded depths of rationality and implementation with complete information. *Journal of Economic Theory*, 165:517–564.
- Thomas, C. (2020). Classification of priorities such that deferred acceptance is obviously strategyproof. Working paper.
- Troyan, P. (2019). Obviously strategy-proof implementation of top trading cycles. *International Economic Review*, 60(3):1249–1261.
- Vickrey, W. (1961). Counterspeculation, auctions, and competitive sealed tenders. *The Journal of Finance*, 16(1):8–37.
- Woodward, K. (2020). Self-auditable auctions. Working paper.
- Zhang, L. and Levin, D. (2017). Bounded rationality and robust mechanism design: An axiomatic approach. *American Economic Review*, 107(5):235–39.
- Zhang, L. and Levin, D. (2021). Partition obvious preference and mechanism design: Theory and experiment. Working paper.

A Appendix

LEMMA A.1. *Let \mathcal{C} be a chain of information partitions for some tree X , ordered by coarseness. Let $\mathcal{I}^{\mathcal{C}}$ be defined such that $\mathcal{I}^{\mathcal{C}}(h) := \{h' \in H \mid \text{there is } \mathcal{I} \in \mathcal{C} \text{ such that } \mathcal{I}(h) = \mathcal{I}(h')\}$. Then, $\mathcal{I}^{\mathcal{C}}$ is an information partition.*

PROOF. It is immediate to see that $\mathcal{I}(h') = \mathcal{I}(h'')$ implies $A(h') = A(h'')$ and $P(h') = P(h'')$. Next, we prove that $\mathcal{I}^{\mathcal{C}}$ is indeed a partition. Suppose $h \in \mathcal{I}^{\mathcal{C}}(h') \cap \mathcal{I}^{\mathcal{C}}(h'')$: Then, there are $\mathcal{I}' \in \mathcal{C}$ and $\mathcal{I}'' \in \mathcal{C}$ such that $\mathcal{I}'(h') = \mathcal{I}'(h)$ and $\mathcal{I}''(h'') = \mathcal{I}''(h)$. Without loss, we can assume that \mathcal{I}'' is coarser than \mathcal{I}' . This implies that $\mathcal{I}''(h') = \mathcal{I}''(h'') = \mathcal{I}''(h)$. Now let $h''' \in \mathcal{I}^{\mathcal{C}}(h')$. Then, there is $\mathcal{I}''' \in \mathcal{C}$ such that $\mathcal{I}'''(h''') = \mathcal{I}'''(h')$. Again without loss, we can assume that \mathcal{I}''' is coarser than \mathcal{I}'' . But then, $\mathcal{I}'''(h''') = \mathcal{I}'''(h'')$ and thus $h''' \in \mathcal{I}^{\mathcal{C}}(h'')$: This implies $\mathcal{I}^{\mathcal{C}}(h') \subseteq \mathcal{I}^{\mathcal{C}}(h'')$. A symmetric argument proves the other inclusion, yielding $\mathcal{I}^{\mathcal{C}}(h') = \mathcal{I}^{\mathcal{C}}(h'')$.

Finally, we show that $\mathcal{I}^{\mathcal{C}}$ satisfies perfect recall. Fix an agent n , some $I_n^{\mathcal{C}} \in \mathcal{I}_n^{\mathcal{C}}$ and two histories such that $h', h'' \in I_n^{\mathcal{C}}$. Then, there is $\mathcal{I} \in \mathcal{C}$ such that $\mathcal{I}(h') = \mathcal{I}(h'')$. Since \mathcal{I} satisfies perfect recall, along both h' and h'' , agent n called to play the same number of times. If agent n is never called to play before h' and h'' , then the two histories have the same experience and we are done, so suppose not. Let h'_k and h''_k be the k -th predecessor of h' and h'' at which n is called to play, respectively. Similarly, let a'_k and a''_k be the actions such that $(h'_k, a'_k) \preceq h'$ and $(h''_k, a''_k) \preceq h''$. Since \mathcal{I} has perfect recall, it must be $\mathcal{I}(h'_k) = \mathcal{I}(h''_k)$ and $a'_k = a''_k$ for all k . But $\mathcal{I}(h'_k) = \mathcal{I}(h''_k)$ implies $\mathcal{I}^{\mathcal{C}}(h'_k) = \mathcal{I}^{\mathcal{C}}(h''_k)$, and thus h' and h'' have the same experience under $\mathcal{I}^{\mathcal{C}}$. \square

LEMMA A.2. *The finest common coarsening of two information partitions is an information partition itself.*

PROOF. Let \mathcal{I}^1 and \mathcal{I}^2 be two information partitions for a tree X , and let \mathcal{I} be their finest common coarsening. By construction of \mathcal{I} , for each n , there exists $\mathcal{I}_n \subset \mathcal{I}$ that partitions the set of histories at which n is called to play, and that $A(h') = A(h'')$ whenever $\mathcal{I}(h') = \mathcal{I}(h'')$. Suppose by contradiction \mathcal{I} does not have perfect recall. Then, there exists a terminal history z' such that, for some $h' \prec z'$, there is $h'' \in \mathcal{I}_n(h')$ such that agent n 's experience at h' and h'' is not the same. Since the set of predecessors of z' is well-ordered, we can assume that h' is the earliest history preceding z' at which this is true. By construction of \mathcal{I} , there must be a finite sequence $\{I^k\}_{k=1}^K$ such that $I^k \in \mathcal{I}_n^1$ for all odd k and $I^k \in \mathcal{I}_n^2$ for all even k , the intersection $I^k \cap I^{k+1}$ is nonempty for all $k < K$, and $h' \in I^1$ and $h'' \in I^K$. Let $h^1 = h'$ and $h^K = h''$, and take $h^k \in I^k \cap I^{k-1}$ for all $1 < k < K$.

Let $\zeta(h^1)$ be the last history agent n is called to play before h^1 . (If she was never called to play before, let $\zeta(h^1)$ be the root of the game. Since \mathcal{I}^1 has perfect recall, it

follows that, along h^2 , agent n is called to play as many times as along h^1 , and takes the same actions. Then, we can define $\xi(h^2)$ analogously to $\xi(h^1)$. Furthermore—if agent n was ever called to play before—again by perfect recall, it has to be $\mathcal{I}^1(\xi(h^2)) = \mathcal{I}^1(\xi(h^1))$. This implies that $\mathcal{I}(\xi(h^2)) = \mathcal{I}(\xi(h^1))$. By the way we selected h' , it follows that the experience of player n at $\xi(h^2)$ and $\xi(h^1)$ is the same. But since they lie in the same information set and agent n took the same action at both, it follows that her experience at h^1 and h^2 is the same as well. Proceeding this way with an induction argument, we conclude that agent n 's experience at h' is the same as at h'' , reaching a contradiction. \square

PROOF OF PROPOSITION 1. Fix a tree X and equip the set of information partitions for X with the partial order defined by coarseness. Let C be a chain of information partitions. Then, \mathcal{I}^C such that $\mathcal{I}^C(h) := \{h' \in H \mid \text{there is } \mathcal{I}' \in C \text{ such that } h' \in \mathcal{I}'(h)\}$ for all non-terminal h is an information partition that is an upper bound for C . Hence, by Zorn's Lemma, there is a maximal information partition \mathcal{I} .

Next, we show that \mathcal{I} is the (unique) coarsest information partition for X . Suppose not: Then, there is an information partition \mathcal{I}' and an non-terminal h such that $\mathcal{I}'(h)$ is not a subset of $\mathcal{I}(h)$. For all non-terminal h , define $\mathcal{I}' \vee \mathcal{I}$ as the finest common coarsening of \mathcal{I}' and \mathcal{I} . Then, $\mathcal{I}' \vee \mathcal{I}$ is an information partition for X that is strictly coarser than \mathcal{I} , contradicting its maximality. The fact that \mathcal{I}^C and $\mathcal{I}' \vee \mathcal{I}$ are indeed information partitions for X follows from Lemmas A.1 and A.2 above. \square

PROOF OF PROPOSITION 2. Define (X, \mathcal{I}') and (X, \mathcal{I}'') such that \mathcal{I}' is finer than \mathcal{I}'' , and let Σ' and Σ'' denote their respective behavior spaces (with their projections defined in the standard way). Suppose agent n moves at h and that, under the coarser \mathcal{I}'' , the action $a \in A(h)$ is obviously dominated for some $\theta_n \in \Theta_n$. We want to show the same is true under the finer \mathcal{I}' . To this end, suppose $a \in A(h)$ is not obviously dominated for θ_n under \mathcal{I}' . Then, it must be that, for each \mathcal{I}'' -measurable σ''_n with $\sigma''_n(h) \neq a$, there is some $\sigma'_n \in \Sigma'_n$ with $\sigma'_n(h) = a$, some $h^1, h^2 \in \mathcal{I}'(h)$, and some $\sigma_{-n}^1, \sigma_{-n}^2 \in \Sigma'_{-n}$ such that

$$\omega(\sigma''_n, \sigma_{-n}^1 \mid h^1) <_{\theta_n} \omega(\sigma'_n, \sigma_{-n}^2 \mid h^2).$$

Since \mathcal{I}'' is coarser than \mathcal{I}' , we have $h^1, h^2 \in \mathcal{I}''(h)$. Moreover, again by Proposition 5, it is without loss to assume that all behaviors above are \mathcal{I}'' -measurable.⁶² But this implies that a is not obviously dominated under \mathcal{I}'' , completing the proof. \square

⁶²Indeed, there is a coarse behavior that leads to the terminal history reached from h^2 under $(\sigma''_n, \sigma_{-n}^2)$. Similarly, there is a coarse σ_{-n} such that, starting from h^1 , the behavior profile $(\sigma''_n, \sigma_{-n})$ reaches the same terminal history as $(\sigma''_n, \sigma_{-n}^1)$.

PROOF OF PROPOSITION 3. By perfect recall, for all σ'_m and σ''_m that satisfy the stated conditions there exists a unique behavior σ'''_m such that we have $\omega(s_n(\theta_n), \sigma'''_m, \sigma_{-n,m}) = \omega(s_n(\theta_n), \sigma'_m, \sigma_{-n,m})$ and $\omega(s_n(\theta_n), \sigma'''_m, \sigma_{-n,m}) = \omega(\sigma_n, \sigma''_m, \sigma_{-n,m})$. \square

LEMMA A.3. Consider two mechanisms (X, \mathcal{I}') and (X, \mathcal{I}'') such that \mathcal{I}' is finer than \mathcal{I}'' . If s is an ex-post equilibrium for (X, \mathcal{I}'') , then s is also an ex-post equilibrium for (X, \mathcal{I}') .

PROOF. We start by observing that measurability of behaviors is always preserved by an increase in transparency. Now, consider two mechanisms (X, \mathcal{I}') and (X, \mathcal{I}'') such that \mathcal{I}' is finer than \mathcal{I}'' and suppose that s is an ex-post equilibrium of the latter but not of the former. Then, for some agent n , there is $\theta \in \Theta$ and an \mathcal{I}' -measurable σ'_n such that

$$\omega(s_n(\theta_n), s_{-n}(\theta_{-n})) <_{\theta_n} \omega(\sigma'_n, s_{-n}(\theta_{-n})).$$

Then, there is a unique h' reached under both $(s_n(\theta_n), s_{-n}(\theta_{-n}))$ and $(\sigma'_n, s_{-n}(\theta_{-n}))$ at which n is called to play and $s_n(\theta_n)(h) \neq \sigma'_n(h)$. Let $I''_n = \mathcal{I}''(h')$. By perfect recall, I''_n is the last information set of agent n in \mathcal{I}'' crossed by both $(s_n(\theta_n), s_{-n}(\theta_{-n}))$ and $(\sigma'_n, s_{-n}(\theta_{-n}))$. Hence, there is a \mathcal{I}'' -measurable σ''_n such that, at all h at which n is called to play under $(\sigma'_n, s_{-n}(\theta_{-n}))$, we have $\sigma''_n(h) = \sigma'_n(h)$. But then, $\omega(\sigma''_n, s_{-n}(\theta_{-n})) = \omega(\sigma'_n, s_{-n}(\theta_{-n}))$, contradicting that s is an equilibrium of (X, \mathcal{I}'') . \square

PROOF OF PROPOSITION 4. Let s be a dominant strategy equilibrium for (X, \mathcal{I}) . By Lemma A.3, s is an ex-post equilibrium when X is equipped with its finest information partition, under which the mechanism is of perfect information. Hence, to complete the proof, it suffices to show that if \mathcal{I} is such that all information sets are singletons, then s is obviously dominant. This follows since, by perfect information, for all h at which n is called to play that can be reached under $(s_n(\theta_n), \sigma_{-n})$ for some σ_{-n} , we have that h is reached under $(s_n(\theta_n), \sigma'_{-n})$ and $(s_n(\theta_n), \sigma''_{-n})$ if and only if σ'_{-n} and σ''_{-n} specify the same action at all histories that precede h at which agent n is not called to play. Then, since all information sets are singletons, for all σ_n that specify the same action as $s_n(\theta_n)$ at all histories that precede h at which n is called to play, there exists a single σ'''_{-n} such that $\omega(\sigma_n, \sigma'''_{-n}) = \omega(\sigma_n, \sigma'_{-n})$ and $\omega(s_n(\theta_n), \sigma'''_{-n}) = \omega(s_n(\theta_n), \sigma''_{-n})$. \square

PROOF OF PROPOSITION 5. Suppose h is reached under (σ_n, σ_{-n}) . Let \mathcal{I}_m be the set of information sets of agent m under the coarsest information partition. By perfect recall, for all distinct $h', h'' \preceq h$ at which m is called to play, we have $\mathcal{I}_m(h') \neq \mathcal{I}_m(h'')$. Then, there is a coarse behavior σ'_m such that $\sigma'_m(h') = \sigma_m(h')$ for all $h' \preceq h$ at which

agent m is called to play. Hence, there is a coarse profile σ_{-n} such that h is reached under (σ_n, σ'_{-n}) . \square

LEMMA A.4. *If agent n has coarsely optimal strategy s_n with $s_n(\theta_n) = \sigma_n$, then σ_n is an undominated behavior for θ_n .*

PROOF. Suppose the statement is false for some mechanism (X, \mathcal{I}) . Then, there is an agent n for which there are a type θ_n and behaviors σ'_n and σ''_n such that σ'_n dominates σ''_n for θ_n , but σ''_n is a best response for θ_n to all coarse behavior profiles $\sigma_{-n} \in \Sigma_{-n}(\mathcal{I})$.

Let D_{-n} be the set of behaviors for the others, σ_{-n} , such that $\omega(\sigma''_n, \sigma_{-n}) <_{\theta_n} \omega(\sigma'_n, \sigma_{-n})$, i.e., such that type θ_n strictly prefers to follow σ'_n over σ''_n . Since σ'_n dominates σ''_n , the set D_{-n} is not empty. Moreover, there must be some σ_c for which D_{-n} is not empty even when we restrict attention to behavior profiles in which chance's behavior is given by σ_c . With a slight abuse of notation, let D_{-n} be as defined above but fixing chance's behavior to σ_c .

Next, for all σ_{-n} , let $d(\sigma_{-n})$ be the length of the longest history $h \preceq z(\sigma'_n, \sigma_{-n})$ such that, for all $h' \prec h$ at which n is not called to play, the following holds true: If there is $h'' \prec z(\sigma''_n, \sigma_{-n})$ such that $h' \in \mathcal{I}(h'')$, then $\sigma_{-n}(h') = \sigma_{-n}(h'')$. Thus, $d(\sigma_{-n})$ is the length of the longest history $h \preceq z(\sigma'_n, \sigma_{-n})$ for which σ_{-n} is consistent with reaching both (σ'_n, σ_{-n}) and $(\sigma''_n, \sigma_{-n})$ while behaving coarsely at all histories preceding h . Observe that, for all $\sigma_{-n} \in D_{-n}$, it must be that $d(\sigma_{-n})$ is strictly less than the length of $z(\sigma'_n, \sigma_{-n})$. Otherwise, we can construct a coarse $\sigma^c_{-n} \in \Sigma_{-n}(\mathcal{I})$ with $\sigma^c_{-n}(h) = \sigma_{-n}(h)$ for all h at which n is not called to play such that $h \prec z(\sigma_n, \sigma_{-n})$ for some $\sigma_n \in \{\sigma'_n, \sigma''_n\}$: But then, we would have $z(\sigma_n, \sigma^c_{-n}) = z(\sigma_n, \sigma_{-n})$ for all $\sigma_n \in \{\sigma'_n, \sigma''_n\}$, contradicting coarse optimality.

Let d be the largest integer such that $d(\sigma_{-n}) = d$ for some $\sigma_{-n} \in D_{-n}$. This is well defined since D_{-n} is nonempty and since the tree is of finite length for fixed σ_c . Let σ_{-n} be such that $d(\sigma_{-n}) = d$. By perfect recall, we can assume without loss that $\sigma_{-n}(\mathcal{I}(h)) = \sigma_{-n}(h)$ for all $h \prec z(\sigma'_{-n}, \sigma_{-n})$ at which n is not called to play and whose length is strictly less than d .

Next, let h' be the unique strict predecessor of $z(\sigma'_{-n}, \sigma_{-n})$ of length d . By construction, h' is well defined, neither chance nor agent n is not called to play at h' , and there is $h'' \prec z(\sigma''_{-n}, \sigma_{-n})$ such that $\mathcal{I}(h') = \mathcal{I}(h'')$ and $\sigma_{-n}(h') \neq \sigma_{-n}(h'')$. Let σ'_{-n} and σ''_{-n} be defined such that $\sigma'_{-n}(h) = \sigma''_{-n}(h) = \sigma_{-n}(h)$ at all $h \notin \mathcal{I}(h')$ at which n is not called to play, but $\sigma'_{-n}(\mathcal{I}(h')) = \sigma_{-n}(h')$ while $\sigma''_{-n}(\mathcal{I}(h')) = \sigma_{-n}(h'')$. Since \mathcal{I} satisfies perfect recall, we have $z(\sigma'_n, \sigma'_{-n}) = z(\sigma'_n, \sigma_{-n})$ and $z(\sigma''_n, \sigma''_{-n}) = z(\sigma''_n, \sigma_{-n})$. Moreover, by construction, we have $d < d(\sigma'_{-n})$ and $d < d(\sigma''_{-n})$, since these behaviors behave coarsely at least one more step than σ_{-n} . Then, by definition of d , it must be

that $\{\sigma'_{-n}, \sigma''_{-n}\}$ and D_{-n} are disjoint sets: Since σ'_n dominates σ''_n for type θ_n , it follows that

$$\begin{aligned}\omega(\sigma'_n, \sigma''_{-n}) &=_{\theta_n} \omega(\sigma''_n, \sigma''_{-n}) = \omega(\sigma''_n, \sigma_{-n}) \\ &<_{\theta_n} \omega(\sigma'_n, \sigma_{-n}) = \omega(\sigma'_n, \sigma'_{-n}) =_{\theta_n} \omega(\sigma''_n, \sigma'_{-n}).\end{aligned}$$

Finally, observe that $\sigma_{-n}(h') \neq \sigma_{-n}(h'')$ implies that $\mathcal{I}(h') \neq \mathcal{I}(h'')$. This implies that we can define σ'''_{-n} such that $\sigma'''_{-n}(h) = \sigma_{-n}(h)$ for all $h \notin \mathcal{I}(h') \cup \mathcal{I}(h'')$, $\sigma'''_{-n}(h') = \sigma_{-n}(h'')$ and $\sigma'''_{-n}(h'') = \sigma_{-n}(h')$. But then, by perfect recall, $z(\sigma'_n, \sigma'''_{-n}) = z(\sigma'_n, \sigma''_{-n})$ and $z(\sigma''_n, \sigma'''_{-n}) = z(\sigma''_n, \sigma'_{-n})$ which, by the above inequality, implies $\omega(\sigma'_n, \sigma'''_{-n}) <_{\theta_n} \omega(\sigma''_n, \sigma'''_{-n})$. This contradicts the assumption that, for type θ_n , behavior σ''_n is dominated by σ'_n . \square

PROOF OF THEOREM 1. By Lemma A.4, coarsely optimal strategies are undominated. This proves the second part of the statement. Moreover, it implies that in a dominant strategy mechanism, for each agent, only a dominant strategy can be coarsely optimal, i.e., satisfy condition (i) of Definition 2. It is easy to see that any strategy profile consisting of dominant strategies for each agent also satisfies condition (ii), proving the first part of the statement. \square

PROOF OF PROPOSITION 6. We first prove statement (i). The statement's assumptions can hold true only if, as long as chance picks a coarsely optimal strategy for each agent, the resulting strategy profile is an ex-post equilibrium. Indeed, otherwise, there would have to be some agent who has a non-coarse coarsely optimal strategy, and we can find (small enough) cognitive costs such that at least one agent is not ex-post best responding by only considering the coarsest partition. Thus, (X, \mathcal{I}) is an as-if dominant strategy mechanism.

We now prove statement (ii). By condition (ii) of Definition 2, as long as the others' first selves choose to consider the coarsest information partition, choosing the coarsest partition is enough make the second self play optimally in the second stage. Since considering finer partitions is (strictly) costly, the claim follows. \square

LEMMA A.5. *In a second-price auctions with asynchronous bids, a strategy s_n is coarsely optimal if and only if it satisfies the following at all I_n that may be reached for some σ_{-n} :*

- (i) *If at I_n it has been revealed either that $r \geq \theta_n$ or that a previous bid b_m by some agent $m \neq n$ satisfies $\theta_n \leq b_m$: Agent n bids truthfully or such that she either loses for sure, and chooses any next access date.*
- (ii) *Otherwise: Agent n bids weakly below θ_n . If she bids strictly below θ_n , then this is not her L -th bid and she selects a finite next access date.*

PROOF. It is easy to see that strategies which satisfy the above conditions are undominated. Next, observe that if others behave coarsely, from the perspective of agent n , each of their behaviors is equivalent to submitting a bid in a sealed-bid auction. Then, as in a sealed-bid second-price auction, it is in the strict interest of agent n to bid truthfully, unless at some point it is revealed to her there is no way to make a strictly positive payoff, in which case she can bid anything that ensures she loses the auction. The two conditions ensure agent n behaves consistently with these incentives—in particular, since she is not certain of the number of bidders, ending the auction with a bid that is not θ_n cannot be coarsely optimal unless bidder n knows she can never win at a profit. \square

PROOF OF THEOREM 2. We start by proving the if direction. Suppose (X, \mathcal{I}) is a second-price auction with asynchronous bids and that \mathcal{I} satisfies the condition from the statement. Note that the coarsest information partition \mathcal{S} is corresponds to a standard second-price sealed-bid auction with the only difference that agents choose when to submit their bids and can revise their bids upwards up to L times. Clearly, bidding truthfully once and at the earliest possible access date is an undominated strategy for each agent that satisfies condition (i) of Definition 2.

We now show that under the information condition from the statement of the theorem, condition (ii) of Definition 2 is also satisfied. Suppose not. This implies there is some agent n and a profile of coarsely optimal strategies of the others, s_{-n} , such that for some type profile $\theta \in \Theta_n$ and $\sigma_c \in \Sigma_c$ there exists a behavior $\sigma_n \in \Sigma_n$ such that $\omega(\sigma_n, s_{-n}(\theta_{-n}), \sigma_c) >_{\theta_n} \omega(s_n(\theta_n), s_{-n}(\theta_{-n}), \sigma_c)$ for some coarsely optimal strategy s_n . Recall that, for any agent n and type $\theta_n \in \Theta_n$, in any coarsely optimal strategy profile s , we must have that $s_n(\theta_n)$ satisfies the conditions stated in Lemma A.5. Let σ_c be such that the reserve price is r .

(Case 1.) Suppose that $\omega(s_n(\theta_n), s_{-n}(\theta_{-n}), \sigma_c)$ is such that agent n is not allocated the object, and $\omega(\sigma_n, s_{-n}(\theta_{-n}), \sigma_c)$ allocates the object to agent n at some price strictly below θ_n . If an agent m is allocated the object under $\omega(s_n(\theta_n), s_{-n}(\theta_{-n}), \sigma_c)$, it is straightforward to see that agent m must have bid at least θ_n at some point, and thus, by Lemma A.5, we must have $\theta_m \geq \theta_n \geq r$. But then, again by Lemma A.5, agent n cannot get allocated the object at a price strictly below θ_n under $\omega(\sigma_n, s_{-n}(\theta_{-n}), \sigma_c)$, contradiction. If no one is allocated the object, under $\omega(s_n(\theta_n), s_{-n}(\theta_{-n}), \sigma_c)$, then $r \geq \theta_n$, which once again contradicts that agent n is allocated the object at a price strictly below θ_n under $\omega(\sigma_n, s_{-n}(\theta_{-n}), \sigma_c)$.

(Case 2.) Suppose $\omega(s_n(\theta_n), s_{-n}(\theta_{-n}), \sigma_c)$ allocates agent n the object at some payment b , whereas $\omega(\sigma_n, s_{-n}(\theta_{-n}), \sigma_c)$ allocates agent n the object at some payment $b' < b$. (Observe that we know $b \leq \theta_n$ by Lemma A.5.) If $r = b$, we get an immediate

contradiction. If $b > r$, then $\max_{m \neq n} \theta_m > r$. Observe that, by the rules of the auction, no bids strictly above b and strictly below $s_n(\theta_n)$ were submitted. Then, the information partition is such that whenever an agent $m \neq n$ submitted a bid, she did not know whether a bid strictly above b was submitted (even if agent n did). Hence, Lemma A.5 implies that $b = \max_{m \neq n} \theta_m$. But then, by a similar argument as above, under the information partition from the statement, there must be a bid of at least θ_m submitted by some agent other than agent n even if she behaves according to σ_n , which contradicts that agent n is allocated the object at a price $b' < b$ under $\omega(\sigma_n, s_{-n}(\theta_{-n}), \sigma_c)$.

This proves every coarsely optimal strategy profile is an ex-post equilibrium, and hence concludes the proof of the if direction.

For the only if direction, suppose there is some information set I_n at which it is revealed to agent n that a bid of at least β has been made by some other agent, but she does not know for certain that the final selling price will be at least β . Note that this implies agent n cannot have bid above β up to I_n , and that there exists some history $h \in I_n$ such that $r < \beta$ and β is strictly higher than any other submitted bid. Next, take $\theta_{-n} \in \Theta_{-n}$ such that each $\theta_{m'}$ is equal to the last highest bid $\beta_{m'} < \beta$ submitted by agent m' along h (and to zero if she never submitted a bid). Take θ_n to be equal to the last submitted bid by agent n before I_n (and to zero if she has not submitted one yet) and recall that we must have $\theta_n < \beta$. Then, note that there is a coarsely optimal strategy s_n that prescribes the following behavior to θ_n : If I_n is reached, bid β . Else, never bid higher than θ_n . Let m denote the agent who submits the bid of at least β along h . Let (s_{-n}, σ_c) be such that h is reached under $(s(\theta), \sigma_c)$ and all agents other than n always bid exactly their value by the end of the auction. (Observe that these are coarsely optimal strategies by Lemma A.5.)

Let σ_m be the behavior in which agent m only bids β once, and on a date that comes (weakly) after the date at which agent n chooses to enter the platform next at the last information set she is called to play before I_n .

Then, I_n is not reached under $(\sigma_m, s_{-m}(\theta_{-m}), \sigma_c)$ and thus

$$\omega(s_m(\theta_m), s_{-m}(\theta_{-m}), \sigma_c) <_{\theta_m} \omega(\sigma_m, s_{-m}(\theta_{-m}), \sigma_c).$$

Indeed, the former outcome allocates agent m the object at a payment of $\max_{m' \neq m} \theta_{m'}$, whereas the latter allocates agent m the object at a payment of $\beta > \max_{m' \neq m} \theta_{m'}$ by construction of θ_{-m} . \square

PROOF OF PROPOSITION 7. By Lemma A.5, the information structure of an as-if dominant strategy second price auction with asynchronous bids is such that, whenever

playing an as-if dominant strategy, an agent ends the auction with a bid different from her value if and only if at some point she was revealed that either the reserve value or a losing bid of another agent was weakly above their value. This implies all the claims of the statement. \square

PROOF OF THEOREM 3. Consider a second-price auction (X, \mathcal{I}) with asynchronous bids that is more transparent than the online English auction (X, \mathcal{I}^E) . Then, there must be an agent n and a history h at which she is called to play such that $\mathcal{I}_n(h)$ is a proper subset of $\mathcal{I}_n^E(h)$. The information structure of the Online English Auction is such that—up to bids that are not admissible, about which neither \mathcal{I} nor \mathcal{I}^E can disclose information—each history in $\mathcal{I}_n^E(h)$ is uniquely identified with a possible value β for the currently winning bid. This implies that there must be at least two of these values, and that agent n did not submit the currently winning bid. Further, the lowest of such values, β' , is the value of the maximum between the reserve price and the highest admissible losing bid. Then, if $\beta' \notin \mathcal{I}_n(h)$, we immediately have a contradiction of Theorem 2. If $\beta' \in \mathcal{I}_n(h)$, then there is $\beta \in \mathcal{I}_n^E(h) \setminus \mathcal{I}_n(h)$ such that the lowest value for the currently winning bid in $\mathcal{I}_n(\beta)$ is strictly above β' . Once again, we have reached a contradiction. \square

PROOF OF PROPOSITION 8. Let $f(d_k) = \pi(\{d_k\})$ and $F(d_k) = \pi(\{d_\ell \mid k \leq \ell\})$.

For part (i), consider the Online English Auction with the optimal reserve price (which is independent of the number of bidders). Suppose first that a skill bidder, who without loss we can think of as an additional bidder 0, never submits a bid. Then, it an as-if dominant strategy—and hence, ex-post—equilibrium for all other bidders to be truthful whenever they access the platform to submit their bid. It remains to show that given this behavior for the others, bidder 0 indeed never wants to submit a bid $b > 0$.

Given the information partition of the Online English Auction, submitting a bid at any date t at which it is known that the highest bid is at least β is equivalent to setting a new reserve price among all bidders with values of at least β , since they are assumed to bid truthfully. Since the timing at which bidders submit bids is uninformative of their type, the optimal reserve price is independent of the number of bidders whenever the virtual values are increasing, the optimal bid that bidder 0 can submit solves

$$\max_{b \geq \beta} b \cdot \left(\frac{1 - F(b)}{1 - F(\beta)} \right).$$

However, comparing b_k to b_{k+1} for some $b_k \geq \beta$, by regularity, we have

$$b_{k+1} \left(\frac{1 - F(b_k)}{1 - F(\beta)} \right) - b_k \left(\frac{1 - F(b_{k-1})}{1 - F(\beta)} \right) \stackrel{\text{sgn}}{=} -\frac{1}{1 - F(\beta)} \left(b_k - (b_{k+1} - b_k) \frac{1 - F(b_k)}{f(b_k)} \right) < 0$$

whenever $\beta \geq r^*$, where r^* is the optimal reserve.⁶³ Thus, submitting some $b > \beta$ cannot be profitable in expectation, and setting $b = \beta$ does not affect the outcome, since at least one bidder other than the current winning bidder must have already submitted a bid of β . Hence, the Online English Auction has a truthful, fully asynchronous, and shill-proof equilibrium.

For part (ii), it suffices to observe that if a shill bidder observes that the highest bid is above β but the final price may be below β , submitting a bid of β strictly increases her expected revenues. \square

PROOF OF PROPOSITION 9. Without loss of generality, we can assume that the marginal values lie in some set $\{0, 1, \dots, u\}$ with $u \geq 1$. Suppose (X, \mathcal{I}) is a mechanism that implements the VCG outcome in obviously dominant strategies. Let s denote the obviously dominant strategy profile. With no loss we can assume that chance plays all its behaviors with positive probability and that \mathcal{I} is such that there is perfect information. Now, take a type profile $\theta' \in \Theta$ such that all but agents 1 and 2 have zero marginal value for all units. The other two agents have a marginal value of one for the first unit and zero for all others. Let z' be the terminal history that is reached under $(s(\theta'), \sigma'_c)$ for some $\sigma'_c \in \Sigma_c$. Then, let $h' \prec z'$ be the first on-path history at which an agent $n \in \{1, 2\}$ is called to play and $s_n(\theta'_n)(h') \neq s_n(\theta''_n)(h')$ for some $\theta''_n > \theta'_n$. Such a history must exist if the Vickrey outcome is implemented. Then, h' is reached under $(s(\theta), \sigma'_c)$ for all type profiles $\theta \in \Theta$ such that agents 1 and 2 have nonzero marginal values and $\theta_m = \theta'_m$ for all other agents m . In particular, h' is reached also under $(s(\theta''), \sigma'_c)$, where θ'' is such that agents 1 and 2 assign a marginal value of u to all units, and $\theta''_m = \theta'_m$ for all other agents m . Suppose $s_n(\theta'_n)(h') \neq s_n(\theta''_n)(h')$. Then, implementing the VCG outcome while preserving obvious dominance it must be that, for all z such that $(h', s_n(\theta'_n)(h')) \preceq z$, if outcome $\omega(z)$ assigns q units to agent n , it makes her pay at least qu . But this implies that $\omega(s(\theta'), \sigma_c)$ assigns no unit at price zero to agent n , contradicting the assumption that s implements the VCG outcome. The case $s_n(\theta'_n)(h') = s_n(\theta''_n)(h')$ is similar. \square

PROOF OF THEOREM 4. We prove the stronger claim that the Ausubel auction with public dropouts is a perfect as-if dominant strategy mechanism. For each agent n let

⁶³Recall that r^* is the lowest value $r \in \mathcal{D}$ such that $\left(b_k - (b_{k+1} - b_k) \frac{1 - F(b_k)}{f(b_k)} \right) > 0$ for all $b_k \geq r$ (see, e.g., Elkind (2007); Akbarpour and Li (2020)).

s_n denote sincere bidding. That is, for each type θ_n , let $s_n(\theta_n)$ denote the behavior that, whenever called to play, prescribes to drop the ℓ -th unit in the r -th round if and only if $\theta_n^\ell \leq p^r$.

We start by showing that the two conditions from Definition 2 are satisfied. First, note that, for each agent n , s_n is an undominated strategy that satisfies the condition (i) of the definition: If the others behave coarsely, this follows from the fact that sincere bidding is a weakly dominant strategy in the Ausubel auction without dropout information, see Theorem 1 in Ausubel (2004). It also implies that s_n is the only coarsely optimal strategy. Condition (ii) now follows since sincere bidding is an ex-post equilibrium of the Ausubel auction with public dropout information.

We now prove perfection by showing that the additional conditions from Theorem 7 are satisfied. Condition (i) holds by inspection. Further, sincere bidding is a coarse behavior, and due to the assumptions made on Θ_n , there is indeed a type dropping ℓ units at each round $r \leq \Delta - 1$, for all $\ell \in \{0, \dots, L\}$. Hence, condition (ii)(b) is satisfied.

To see that condition (ii)(a) is satisfied, start observing that, since the rationing priority order is predetermined, sincere bidding is never harmful, even after a deviation—see Okamoto (2018). Finally, fix some agent n , type θ_n and information set I_n . Then, for any behavior σ_n that prescribes taking a different action at I_n than $s_n(\theta_n)$, there is some $h \in I_n$ and coarse $\sigma_{-n} \in \Sigma_{-n}(\mathcal{S})$ such that $\omega(\sigma_n, \sigma_{-n} \mid h) <_{\theta_n} \omega(s_n(\theta_n), \sigma_{-n} \mid h)$: Suppose first that σ_n prescribes to drop more units than $s_n(\theta_n)$ at I_n . If the others behave such that, at the round containing I_n , agent n would have clinched all the units for which she had a larger marginal value if she did not drop them, and chance assigned her the lowest rationing priority, she is strictly worse off. Second, suppose that σ_n prescribes to drop fewer units than s_n . Then, if the others do not drop any unit in this round and all their (not yet clinched) units in the next, and chance assigned the highest rationing priority to agent n , she is strictly worse off. This implies that for all I_n and θ_n , any behavior that is not I_n -equivalent to $s_n(\theta_n)$ satisfies the inequality of condition (ii)(a). \square

PROOF OF THEOREM 5. We prove the stronger claim that the sequential DAA is a perfect as-if dominant mechanism. For each agent n and type θ_n let $s_n(\theta_n)$ be as follows: Whenever called to play, apply to the school ranked highest according to θ_n among the ones available (i.e., the ones she has not applied to yet).

We start by showing that the two conditions from Definition 2 are satisfied. Clearly, for each n , the strategy s_n is undominated and satisfies condition (i). Indeed, if the other agents behave coarsely, their behavior is equivalent to submitting a rank order list over schools. Since truthfully reporting one's ranking over schools is a dominant strategy in the static DAA, it follows that s_n undominated and satisfies condition

(i). Additionally, it is the only coarsely optimal strategy, as deviating at any point is always harmful for some coarse behavior of the others. Since all $s_n(\theta_n)$ are coarse behaviors, s is an ex-post equilibrium, and we are done.

We now prove perfection by showing the additional conditions from Theorem 7 hold: Condition (i) holds by inspection. Condition (ii)(b) is trivially satisfied since we assume that Θ_n contains all possible rankings over schools (and the outside option), and each coarse behavior corresponds to going down a list. To see that condition (ii)(a) is satisfied, fix some agent n , type θ_n and information I_n . Even if I_n is reached after a deviation—that is, θ_n did not always apply to her favorite available school in the past—continuing truthfully can never be harmful if the others behave coarsely, as it would not be for the type whose ranking differs from θ_n only to the extent that all previous moves are compliant with s_n . Finally, for all behaviors σ_n that prescribes sending a different message at I_n than $s_n(\theta_n)$, there is some $h \in I_n$ and $\sigma_{-n} \in \Sigma_{-n}(\mathcal{I})$ such that $\omega_n(\sigma_n, \sigma_{-n} \mid h) <_{\theta_n} \omega_n(s_n(\theta_n), \sigma_{-n} \mid h)$. Indeed, let σ_c be such that agent n is unacceptable to all schools she has not yet applied to at I_n except $s_n(\theta_n)(I_n)$ and $\sigma(I_n)$, where she is ranked first. This implies that for all I_n and θ_n , any behavior that is not I_n -equivalent to $s_n(\theta_n)$ satisfies the inequality of condition (ii)(a). \square

PROOF OF PROPOSITION 10. Condition (i) of Definition 6 immediately implies that for all agents there is exactly one strategy s_n which satisfies condition (i) of Definition 2. Thus, it suffices to show that $(s_n)_n$ is an ex-post equilibrium. All ex-post equilibria of the mechanism are Nash equilibria of a game at the start of which chance picks a type profile, each agent can only observe her realized type, and then the mechanism is played. (The chance player is indifferent among all outcomes.) The result then follows by an induction argument, recalling that, in a finite game, elimination of weakly dominated strategies cannot generate new Nash equilibria. \square

LEMMA A.6. *Let (X, \mathcal{I}) satisfy the conditions of Theorem 7. Then, there is a unique as-if dominant strategy profile.*

PROOF. Let s_n be a perfect as-if dominant strategy for agent n . Suppose that s'_n is another as-if dominant strategy for agent n and take any type θ_n . Then, for all $I_n \in \mathcal{I}_n$ at which n is first called to play, we have $(s_n(\theta_n), \sigma_{-n} \mid h) \leq_{\theta_n} (s'_n(\theta_n), \sigma_{-n} \mid h)$ for all $h \in I_n$ and coarse $\sigma_{-n} \in \Sigma_{-n}(\mathcal{I})$. Hence, it must be that $s'_n(\theta_n)$ is I_n -equivalent to $s_n(\theta_n)$ for all I_n at which n is first called to play, which implies that $s_n(\theta_n)$ and $s'_n(\theta_n)$ are equivalent behaviors. Since this holds for all types, $s'_n = s_n$. \square

LEMMA A.7. *Let (X, \mathcal{I}) satisfy the conditions of Theorem 7. Then, the as-if dominant strategy profile survives under all orders of iterated deletion of dominated strategies.*

PROOF. Let s be the as-if dominant strategy profile. Suppose that S^∞ is admissible but, for some agent n , the as-if dominant strategy s_n is not in S_n^∞ . Let $(S^k)_k$ be the sequence of eliminations such that $\bigcap_k S^k = S^\infty$. Let ℓ be the last round in which we have $s_m \in S_m^\ell$ for all $m \in N$. Then, we can take $n \in N$ such that $s_n \in S_n^\ell \setminus S_n^{\ell+1}$.

Let Σ_{-n}^ℓ be the set of behaviors $(\sigma_{-n,c}, \sigma_c) \in \Sigma_{-n}$ such that there is θ_{-n} and $s'_{-n} \in S_{-n}^\ell$ with $s'_{-n}(\theta_{-n}) = \sigma_{-n,c}$. Then, there is some type θ_n , behavior $\sigma_n \in \Sigma_n$ and information set $I_n \in \mathcal{I}_n$ with:

- (i) If she follows either $s_n(\theta_n)$ or σ_n , the information set I_n is reachable for some $\sigma_{-n} \in \Sigma_{-n}^\ell$.
- (ii) We have $\omega(s_n(\theta_n), \sigma_{-n}) <_{\theta_n} \omega(\sigma_n, \sigma_{-n})$ for some $\sigma_{-n} \in \Sigma_{-n}^\ell$ for which I_n is reached.
- (iii) We have $\omega(s_n(\theta_n), \sigma_{-n}) \leq_{\theta_n} \omega(\sigma_n, \sigma_{-n})$ for all $\sigma_{-n} \in \Sigma_{-n}^\ell$ for which I_n is reached.

Conditions (i) and (ii) imply that $s_n^{\theta_n}$ and σ_n are neither equivalent nor I_n -equivalent. Since s is a perfect as-if dominant strategy profile, there exists a coarse $\sigma_{-n} \in \Sigma_{-n}(\mathcal{I})$ such that I_n is reached and $(\sigma_n, \sigma_{-n}) <_{\theta_n} (s_n(\theta_n), \sigma_{-n})$. By the full support condition of perfect as-if dominant strategies, there is $\sigma'_{-n} \in \Sigma_{-n}$ such that I_n is reached, $(\sigma_n, \sigma'_{-n}) <_{\theta_n} (s_n(\theta_n), \sigma'_{-n})$ and $s_{-n}(\theta_{-n}) = \sigma'_{-n}$ for some θ_{-n} . But this implies $\sigma'_{-n} \in \Sigma_{-n}^\ell$, leading to a contradiction of condition (iii) above. \square

LEMMA A.8. *If (X, \mathcal{I}) has a round structure, then it has an elimination guide $(\mathcal{R}_k)_{k=1}^K$, i.e. a partition of \mathcal{R} such that, for all $k \in \{1, \dots, K\}$, $R' \in \mathcal{R}_k$, $h \in R'$ and $a \in A(h)$,*

- (i) either $(h, a) \in R'$
- (ii) or $(h, a) \in \mathcal{R}_{k-1} \cup Z$, where \mathcal{R}_0 is the empty set.

PROOF. Let $\mathcal{R}_1 \subseteq \mathcal{R}$ be the set of all $R \in \mathcal{R}$ such that, $R \preceq R'$ implies $R = R'$. Recursively, for all $k > 1$ such that $\mathcal{R} \neq \bigcup_{\ell < k} \mathcal{R}_\ell$, let \mathcal{R}_k be the set of all $R \in \mathcal{R}$ such that $R \preceq R'$ implies $R = R'$ or $R' \in \mathcal{R}_\ell$ for some $\ell < k$. Suppose the recursion is not finite. Then, at some step $k \geq 1$, the set \mathcal{R}_k is empty but $\mathcal{R} \neq \bigcup_{\ell < k} \mathcal{R}_\ell$. Take any $R_1 \in \mathcal{R}' := \mathcal{R} \setminus \bigcup_{\ell < k} \mathcal{R}_\ell$. Since $\mathcal{R}_k = \emptyset$ and by condition (iii) of Definition 7, for all $\ell > 1$, there is some R_ℓ such that $R_{\ell-1} \preceq R_\ell$ and $R_\ell \in \mathcal{R}' \setminus \{R_1, \dots, R_{\ell-1}\}$. But this implies \mathcal{R} is not finite, a contradiction. \square

LEMMA A.9. *Consider a mechanism satisfying the conditions of Theorem 7 with an elimination guide $(\mathcal{R}_k)_{k=1}^K$, and let \mathcal{R}_0 be the empty set. Consider a strategy rectangle S' such that there is a $k \geq 1$ satisfying the following condition:*

(C) For all agents n , $s'_n \in S'_n$ if and only if, for all θ_n , the behaviors $s_n(\theta_n)$ and $s'_n(\theta_n)$ are I_n -equivalent for all $I_n \in \mathcal{R}_{k-1}$ that can be reached under $s'_n(\theta_n)$, for some behavior of the others.

Then, for all agents n , all strategies $s'_n \in S'_n$ such that there are θ'_n and $I_n \in \mathcal{R}_k$ for which I_n can be reached and $s'_n(\theta'_n)$ is not I_n -equivalent to $s_n(\theta'_n)$ are dominated with respect to S' .

PROOF. Let s be the as-if dominant strategy profile. Suppose that $s'_n \in S'_n$ satisfies the condition we claim is sufficient to be dominated with respect S' . Pick $s''_n \in S'_n$ such that s''_n agrees with s'_n at all information sets that are in some $\mathcal{R}_{k'}$ with $k' < k$, and with s_n for all information sets that are in some \mathcal{R}'_k with $k' \geq k$. Then, since s is a perfect as-if dominant strategy profile and play within a round is simultaneous, by condition (C) and Proposition 5 we have $\omega(s''_n(\theta_n), s_{-n}(\theta_n), \sigma_c) \leq_{\theta_n} \omega(s'_n(\theta_n), s'_{-n}(\theta_{-n}), \sigma_c)$ for all $\theta \in \Theta$ and $s'_{-n} \in S'_{-n}$.⁶⁴ Furthermore, conditions (ii)(a) and (ii)(b) of Theorem 7 ensure that there are θ'_{-n} and $\sigma'_c \in \Sigma_c$ such that I_n is reached under both $s'_n(\theta'_n)$ and $s''_n(\theta'_n)$, and $\omega(s''_n(\theta'_n), s_{-n}(\theta'_n), \sigma'_c) <_{\theta'_n} \omega(s'_n(\theta'_n), s_{-n}(\theta'_{-n}), \sigma'_c)$. \square

LEMMA A.10. If a mechanism satisfying the conditions of Theorem 7 has a round structure, the singleton that contains the as-if dominant strategy profile is an admissible strategy rectangle.

PROOF. Let s be the as-if dominant strategy profile. Since the mechanism has a round structure, by Lemma A.8 it has an elimination guide $(\mathcal{R}_k)_{k=1}^K$. Now, define S^k as the strategy rectangle that satisfies condition (C) for k . Then, by Lemma A.9, we have that all $s'_n \in S_n^{k-1} \setminus S_n^k$ are dominated with respect to S^{k-1} . Since $S^1 = S$ and $S^K = \{s\}$, we are done. \square

PROOF OF THEOREM 7. By Lemma A.10, we know that the singleton that contains the as-if dominant strategy profile s is an admissible strategy rectangle. By Lemma A.7, we know that all admissible strategy rectangles contain s . To complete the proof, we need to show that if S^∞ is an admissible strategy rectangle, then it contains no element other than s .

For all $n \in N$, let $S_n^0 := S_n^\infty \setminus \{s_n\}$. That is, the strategies $s'_n \in S_n^0$ are such that there exists some type $\theta_n \in \Theta_n$ such that $s'_n(\theta_n)$ is not equivalent to $s_n(\theta_n)$. Suppose, towards a contradiction, that S_n^0 is nonempty for some $n \in N$. We show that then

⁶⁴Suppose $\omega(s'_n(\theta_n), s'_{-n}(\theta_n), \sigma_c) <_{\theta_n} \omega(s''_n(\theta_n), s'_{-n}(\theta_{-n}), \sigma_c)$ for some $s'_{-n} \in S'_{-n}$. Then, there is I_n in $\mathcal{R}_{k'}$ with $k \geq k'$ such that $s'_n(\theta_n)$ is not I_n -equivalent to $s_n(\theta_n)$. Since within a round play is simultaneous and starting from the next rounds all continuation strategies are equivalent to the continuation of the as-if dominant strategy profile, by Proposition 5 we can construct a coarse σ_{-n} such that $h \in I_n$ is reached under $(s'_n(\theta_n), \sigma_{-n})$, and $\omega(s_n(\theta_n), \sigma_{-n} | h) <_{\theta_n} \omega(s'_n(\theta_n), \sigma_{-n} | h)$, a contradiction of the assumption that s_n is a perfect as-if dominant strategy.

there exists at least one $n \in N$ for whom a strategy in S_n^0 (and hence S_n^∞) is dominated with respect to S^∞ , which yields a contradiction.

First, let k be the largest integer for which S^∞ satisfies condition (C). We must have $1 \leq k < K$ since $k = K$ would imply that $S_n^0 = \emptyset$ for all n , and condition (C) always holds for $k = 1$.

By the definition of k , there must be a strategy $s'_n \in S_n^0$, a type θ'_n and an information set $I_n \in \mathcal{R}_k$ such that I_n can be reached under $s'_n(\theta'_n)$ for some behavior of the others and $s'_n(\theta'_n)$ is not I_n -equivalent to $s_n(\theta'_n)$. Pick $s''_n \in S'_n$ such that s''_n agrees with s'_n at all information sets that are in some $\mathcal{R}_{k'}$ with $k' < k$, and with s_n for all information sets that are in some \mathcal{R}'_k with $k' \geq k$.

By the same arguments as in the proof of Lemma A.10, as-if dominance implies that we have $\omega(s''_n(\theta_n), s'_{-n}(\theta_{-n}), \sigma_c) \leq_{\theta_n} \omega(s'_n(\theta_n), s'_{-n}(\theta_{-n}), \sigma_c)$ for all $\theta \in \Theta$ and $s'_{-n} \in S_{-n}^\infty$, with at least one inequality holding strict.

(Case 1.) If $s''_n \in S_n^\infty$, then s'_n is dominated with respect to S^∞ , a contradiction.

(Case 2.) Suppose that $s''_n \notin S_n^\infty$. Let $S^\infty = \bigcap_\ell S^\ell$, with $S^1 = S$. If $s''_n := s_n^{\ell_0} \notin S_n^\infty$, there is a round of elimination ℓ_1 such that $s_n^{\ell_0} \in S_n^{\ell_1-1}$ but $s_n^{\ell_0}$ was dominated by some strategy $s_n^{\ell_1} \in S_n^{\ell_1-1}$ with respect to S^{ℓ_1-1} . If $s_n^{\ell_1} \notin S_n^\infty$, then there was some round of elimination $\ell_2 > \ell_1$ such that $s_n^{\ell_1} \in S_n^{\ell_2-1}$ but $s_n^{\ell_1}$ was dominated by some strategy $s_n^{\ell_2} \in S_n^{\ell_2-1}$ with respect to S^{ℓ_2-1} . By finiteness, and since $\{s\} \subset S^\infty$ by Lemma A.7 there is some integer $\ell_M < \infty$ such that $s_n^{\ell_M} \in S_n^\infty$. We now show that $s_n^{\ell_M}$ must dominate s'_n with respect to S^∞ . First, for all $m \geq 1$,

$$\omega(s_n^{\ell_{m-1}}(\theta_n), s'_{-n}(\theta_{-n}), \sigma_c) \leq_{\theta_n} \omega(s_n^{\ell_m}(\theta_n), s'_{-n}(\theta_{-n}), \sigma_c)$$

for all $\theta \in \Theta$, $s'_{-n} \in S_{-n}^{\ell_{m-1}}$ and $\sigma_c \in \Sigma_c$. Since $S^\infty \subset S^{\ell+1} \subset S^\ell$ for all $\ell \geq 1$, this implies that

$$\omega(s''_n(\theta_n), s'_{-n}(\theta_{-n}), \sigma_c) = \omega(s_n^{\ell_0}(\theta_n), s'_{-n}(\theta_{-n}), \sigma_c) \leq_{\theta_n} \omega(s_n^{\ell_M}(\theta_n), s'_{-n}(\theta_{-n}), \sigma_c)$$

for all $\theta \in \Theta$, $s'_{-n} \in S_{-n}^\infty \subset S_{-n}^{\ell_M-1}$ and $\sigma_c \in \Sigma_c$. This implies that s'_n is dominated—by $s_n^{\ell_M}$ —with respect to S^∞ , a contradiction. \square

PROOF OF COROLLARY 2. The finiteness assumptions are satisfied in both settings. As we show in the Appendix, both the Ausubel Auction with public dropouts and the sequential DA Algorithm satisfy the sufficient conditions of Theorem 7, and are hence perfect as-if dominant strategy mechanisms. \square