# Designing Simpler Mechanisms: Transparency and As-If Dominance

Lea Nagel*
Department of Economics
Harvard University

Roberto Saitto*
Kellogg School of Management
Northwestern University

April 27, 2026

[Click here for latest version]

**Abstract**

We show agents can adopt a simplifying heuristic to avoid dominated strategies. It prescribes playing strategies optimal against all unconditional behaviors of the others. A mechanism is *as-if dominant strategy* if the adoption of this heuristic is self supporting. That is, no agent can ever profit by unilaterally strategizing further. As-if dominance is more permissive than dominance, allowing for substantially *more transparent* mechanisms. This is useful, as we show transparency makes the heuristic's prescriptions easier to follow. Our approach rationalizes the auction choice of prominent online platforms, such as eBay. Further, it provides a unified explanation for a range of experimental findings. Finally, under some conditions met in our applications, as-if dominance ensures that coarse rationality emerges from a generational learning process.

## 1 Introduction

Designing mechanisms that are simple to play is important for the practical efficacy of mechanism design. Traditionally, *dominant strategy* mechanisms are considered a

*This paper was previously circulated under the title "As-If Dominant Strategy Mechanisms." 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, Johannes Hörner, Shengwu Li, Andrew Mackenzie, Antonio Penta, Al Roth, Ilya Segal, 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, ESSET 2025, CESifo Area Conference on Behavioral Economics 2025, and at various seminars and workshops for their helpful comments.

natural way to achieve this end, as they provide each agent with a strategy which is optimal regardless of what others do. In this sense, their predictions rely on rationality only, and do not require agents to form correct beliefs. However, a growing body of evidence shows that actually recognizing a strategy is always optimal can be cognitively demanding, leading to surprisingly low rates of rational behavior.[1] Moreover, several of those papers document that increasing the *transparency* of a mechanism—by allowing agents to observe more about other agents' past moves—can significantly improve performance, even at the cost of breaking dominant strategy incentive compatibility. In fact, Kagel and Levin (2009, p. 234) observe 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."

In this paper, we propose a theoretical framework that formalizes when we should expect increases in transparency to be beneficial, and apply it to identify simpler mechanisms in settings ranging from online auctions to school choice problems.

Our starting point is the observation that an agent who wants to avoid dominated strategies has to engage in significant contingent thinking—which is known cognitively demanding (see, e.g., Esponda and Vespa, 2014; Li, 2017). First, she must deduce all possible behaviors of the others from the description of the mechanism—i.e., its extensive form—which requires reasoning about what information others can condition their behavior on. Second, she must condition on each of these behavior profiles when evaluating her own alternatives.

We argue in Section 3 that agents can sometimes adopt a heuristic, which we call *coarse rationality*, to simplify the first task. It consists in playing, if available, a strategy that best responds to all coarse behaviors of the others—behaviors that do not condition on information disclosed over the course of play. Indeed, deducing the others' coarse behaviors does not require an agent to think about what they can condition their actions on. Moreover, we show in Theorem 1 that every strategy prescribed by coarse rationality—which we call *coarsely optimal*—is undominated. This implies that in dominant strategy mechanisms, coarse rationality suffices to always best respond, making it a natural heuristic to adopt. Building on this, in Section 4 we argue that a larger class of mechanisms, which we call *as-if dominant strategy*, makes the common adoption of coarse rationality self-supporting: They are such that (i) each agent has a coarsely optimal strategy, and (ii) if all agents are coarsely rational, then coarse ra-

[1] See, e.g., Bó and Hakimov (2020); Breitmoser and Schweighofer-Kodritsch (2022); Cason et al. (2006); Chen and Sönmez (2006); Echenique et al. (2016); Hakimov and Raghavan (2025); Kagel and Levin (1993, 2009); Klijn et al. (2019); Masuda et al. (2022); Rees-Jones (2018).

tionality suffices to always best respond. That is, *as if* the mechanism was dominant strategy, no agent can gain from unilaterally being more sophisticated.

The usefulness of as-if dominance is due to the fact that it admits significantly more transparent mechanisms than dominance, which can simplify the second contingent reasoning task. Indeed, coarse rationality may still need to condition on a large number of hypothetical behavior profiles of the others—it actually coincides with standard rationality in static dominant strategy mechanisms, which are known to be challenging (see, e.g., Kagel and Levin, 2009; Li, 2017). In Proposition 3 we show that increasing transparency can address this: it shrinks the sets of hypotheticals an agent needs to condition on to test coarse optimality, and increases the set of deviations that are *obviously dominated*, in the sense of Li (2017).[2] This provides a new perspective on the simplicity of obviously dominant strategy mechanisms, as they are more transparent versions of their strategically equivalent static counterparts. At the same time, it explains experimental findings that transparency continues to improve performance even within the class of obviously dominant mechanisms (Breitmoser and Schweighofer-Kodritsch, 2022; Esponda and Vespa, 2014, 2023). However, Proposition 4 shows that preserving dominance under transparency increases requires a demanding local obvious dominance condition that rarely holds in practice. As-if dominance offers a solution to this design dilemma: it maintains coarse rationality as a self-supporting heuristic while allowing for significantly more transparency.

Our main application—asynchronous auctions—illustrates this. Consider an online platform designing a single-item auction that, to encourage and retain participation, is as simple as possible. To attract bidders from different time zones, the platform is restricted to use an asynchronous format allowing bidders to bid at their convenience within some time window. Classical results identify the sealed-bid second-price auction (Myerson, 1981; Vickrey, 1961) as a natural candidate: truthful bidding is a dominant strategy, the item is assigned efficiently, and with an optimal reserve price expected revenues are maximized. Yet, it is well established that agents find sealed-bid second-price auctions difficult to play (see, e.g., Breitmoser and Schweighofer-Kodritsch, 2022; Kagel and Levin, 1993; Li, 2017). While any increase in transparency is incompatible with preserving dominance in this setting, we can use our framework to identify a simpler alternative. In particular, Theorem 2 characterizes the *Online English Auction* as maximally transparent within the set of as-if dominant strategy asynchronous second-price auctions. This format is novel to the literature but matches the design of prominent auction platforms such as eBay:

[2]Intuitively, a deviation is obviously dominated if the agent does not need to condition on the behavior of the others to see that it is always worse.

upon accessing the platform, a bidder can observe everything that happened so far, except the value of the currently winning bid. Our result rationalizes the popularity of this format in settings where synchronous bidding is impractical, preventing the use of ascending auctions. It maximizes transparency while preserving coarse rationality as a self-supporting heuristic, thus reducing contingent reasoning demands and strengthening trust in the platform. We also apply as-if dominance to study asynchronous VCG auctions for the purpose of funding a public project.

Beyond asynchronous auctions, we show in Section 6 that the Ausubel (2004) auction with public dropouts and sequential implementations of the Deferred Acceptance (DA) and Top Trading Cycle (TTC) algorithms are all as-if dominant strategy mechanisms, each strictly more transparent than its dominant-strategy counterpart. This let us provide a unified explanation for experimental evidence for the superior performance of these mechanisms, despite breaking dominant strategy incentives (see, e.g, Bó and Hakimov, 2020; Bó and Hakimov, 2023; Hakimov and Raghavan, 2025; Kagel and Levin, 2009; Klijn et al., 2019). In particular, they make coarse rationality easier to verify, and do so by the same principles underlying the simplicity gains of obviously dominant strategy mechanisms identified in Section 3.

Finally, in Section 7, we show that, under conditions satisfied by all our applications, as-if dominance ensures coarse rationality emerges endogenously from a generational learning process in which agents discard persistently outperformed strategies. Intuitively, coarse rationality eventually becomes both positively and normatively compelling: only coarsely optimal strategies are played, and being coarsely rational suffices to always best respond. The same conditions imply (weak) dominance solvability. In particular, the sequential DA and TTC are dominance solvable.

While our approach can help designers identify effective transparent mechanisms, we abstract away from some important practical considerations. First, the form in which information is disclosed is important for agents' ability to process it.[3] Second, privacy concerns may constrain transparency—for example, in school choice settings, a designer may prefer to reveal only aggregate statistics such as school-specific cutoffs.[4] Third, time constraints may prevent frequent communication between agents and the mechanism. Finally, in settings with sophisticated agents and repeated interaction—e.g., government procurement—transparency may facilitate collusion. This makes as-if dominance best-suited for settings with sporadic interaction.

[3]E.g., eBay highlights key information and allows bidders to access additional details if they wish.

[4]The role of cutoffs in school choice is considered in, e.g., Azevedo and Leshno (2016); Hakimov and Raghavan (2025); Leshno and Lo (2020).

### 1.1 Other related literature

First, we relate to a literature studying mechanisms whose normal form allows agents to have a bounded depth of rationality (Börgers and Li, 2019; Crawford, 2021; de Clippel et al., 2019; Gorelkina, 2018; Jackson et al., 1994; 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 be chosen to reduce the set of opponent behaviors that an agent must consider, thereby making coarse optimality easier to verify.

Second, we relate to a literature on dynamic implementations of dominant strategy rules. Schummer and Velez (2021) restrict attention to sequential elicitation procedures in which each agent is called to play exactly once and observes all previously submitted reports, and show that non-truthful equilibria can arise under incomplete information; Wang and Zheng (2025) study extensive-form transformations that preserve dominant strategy incentives.

Third, we relate to the theory of ex-post implementation originating from Bergemann and Morris (2005, 2008). As-if dominant strategy mechanisms are ex-post incentive compatible, but even an ex-post perfect incentive compatible mechanism may not be an as-if dominant strategy mechanism (cf. Ausubel, 2004). As-if dominance also does not nest with obvious ex-post equilibrium, partition obvious Nash equilibrium and everywhere dominant equilibrium—proposed, respectively, by Li (2017), Zhang and Levin (2017) and Mackenzie and Zhou (2022).

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.

## 2 Preliminaries

We consider finite private-value environments.[5] Let $N$ denote a set of agents $n$, and $\Theta$ a product set of type profiles $\theta$. Let $\Theta_n$ and $\Theta_{-n}$ be the type space of agent $n$ and of all other agents, with generic elements $\theta_n$ and $\theta_{-n}$, respectively.[6] Each type $\theta_n \in \Theta_n$ induces a complete preference ordering $\leq_{\theta_n}$ over the outcome space $\Omega$. Given $\omega', \omega'' \in \Omega$, we write $\omega' =_\theta \omega''$ if $\omega' =_{\theta_n} \omega''$ for all agents $n$.

[5] In particular, in what follows, all sets are nonempty and finite. As we point out throughout the appendix, most of the analysis goes through even when our finiteness assumptions are relaxed.

[6] Throughout, we adopt similar notational conventions whenever dealing with product sets.

### 2.1 Extensive-form mechanisms

An *extensive-form mechanism* is a pair $\Gamma = (X, \mathcal{I})$ consisting of a finite tree $X$ and an information partition $\mathcal{I}$.[7] A formal definition is in Appendix A. Throughout, let $\preceq$ denote the precedence relation on the set of histories $h$. Let $A(h)$ be the non-singleton set of actions available at $h$, and let $\mathcal{I}(h)$ denote the information set containing $h$. The mechanism may include a *chance* player $c$, whose information sets are singletons.

Given a mechanism $(X, \mathcal{I})$, an agent or chance player's *behavior* is a $\mathcal{I}-$measurable map from histories at which she moves into available actions, while an agent's *strategy* is a map from types into behaviors. Let $\Sigma$ and $S$ be the product set of all behavior profiles $\sigma = (\sigma_n, \sigma_{-n}) = (\sigma_{-c}, \sigma_c)$ and strategy profiles $s = (s_n, s_{-n})$, respectively. Let $z(\sigma)$ be the unique terminal history reached under $\sigma$, and $\omega(\sigma)$ be the corresponding implemented outcome. With a slight abuse of notation, we write $\sigma_n' = \sigma_n''$ if they always reach the same terminal history, i.e. $z(\sigma_n', \sigma_{-n}) = z(\sigma_n'', \sigma_{-n})$ for all $\sigma_{-n} \in \Sigma_{-n}$.

For expositional convenience, we restrict attention to mechanisms $(X, \mathcal{I})$ in which, whenever two histories $h', h'' \in H$ at which the same agent moves can be given the same action set by adding *redundant actions*, we have $A(h') = A(h'')$ already.[8]

### 2.2 Rationality and benchmark solution concepts

We say that an agent is *rational* if she plays an *undominated strategy* (cf. Luce and Raiffa, 1957); that is, a strategy $s_n$ for which there is no type $\theta_n$ and behavior $\sigma_n$ such that $\omega(s_n(\theta_n), \sigma_{-n}) \leq_{\theta_n} \omega(\sigma_n, \sigma_{-n})$ for all $\sigma_{-n}$, with at least one inequality strict.[9]

Within this framework, an important class of mechanisms are *dominant strategy mechanisms*, in which rational agents can use a single strategy to eliminate all dominated ones (cf. Li and Dworczak, 2024). In such mechanisms, each agent $n$ has a *dominant strategy*—a strategy $s_n$ such that $\omega(\sigma_n, \sigma_{-n}) \leq_{\theta_n} \omega(s_n(\theta_n), \sigma_{-n})$ for all types $\theta_n$, behaviors $\sigma_n$, and others' behavior profiles $\sigma_{-n}$. We refer to a strategy profile composed solely of dominant strategies as a *dominant strategy equilibrium*.

Li (2017) proposes a refinement of dominance aimed at reducing the contingent reasoning required of agents. Given a strategy $s_n$, a type $\theta_n$ and a history $h$ at which agent $n$ may move under $s_n(\theta_n)$, a *deviating action* $a \in A(h) \setminus \{s_n(\theta_n)(h)\}$ is *obviously*

[7] Our definition of information partition includes perfect recall.

[8] Formally, a mechanism $(X', \mathcal{I}')$ differs from $(X, \mathcal{I})$ only for the *addition of redundant actions* if it shares the same reduced normal form and $X$ is a sub-tree of $X'$. The restriction is without loss as it could equivalently be incorporated into the definitions of transparency and of coarse behaviors in Section 3, with no impact on the analysis. That is, before applying any of our definitions, if necessary, one should first add redundant actions until the stated condition holds.

[9] This notion of rationality is stronger than only requiring 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. A strategy is *dominated* if it is not undominated (see, e.g., Bernheim, 1984; Pearce, 1984).

*dominated* if $\omega(z) \leq_{\theta_n} \omega(s_n(\theta_n), \sigma_{-n})$ for all terminal historues $z$ with $(h', a) \preceq z$ for some $h' \in \mathcal{I}(h)$, and all $\sigma_{-n}$ compatible with reaching $\mathcal{I}(h)$.[10] A strategy is *obviously dominant* if all deviating actions are obviously dominated. A profile is an *obviously dominant equilibrium* if it consists solely of obviously dominant strategies.

A strategy profile $s$ is an *ex-post equilibrium* if $\omega(\sigma_n, s_{-n}(\theta_{-n}), \sigma_{-c}) \leq_{\theta_n} \omega(s(\theta), \sigma_{-c})$ for all agents $n$, profiles $\theta$, behaviors $\sigma_n$ and chance behaviors $\sigma_c$.

## 3 A Dilemma for Mechanism Design

In this section, we introduce a heuristic—*coarse rationality*—that reduces the amount of contingent reasoning required to avoid dominated strategies. In particular, coarse rationality suffices for optimal play in dominant strategy mechanisms. We then identify a design dilemma: increasing transparency can further reduce the contingent reasoning required to behave coarsely rational, but is typically incompatible with maintaining dominant-strategy incentives. We conclude by observing that a weaker requirement may nevertheless suffice for coarse optimality to be adopted by all agents.

### 3.1 Coarse rationality

As discussed in foundational literature—see, e.g., Arrow (1951); Brandenburger et al. (2008); Kohlberg and Mertens (1986); Luce and Raiffa (1957)—avoiding dominated strategies is a reasonable guide for behavior. In practice, however, this may be cumbersome to do even in dominant strategy mechanisms, in which a single strategy can be used to eliminate all dominated ones. In fact, rational play requires agents to correctly perform the following tasks:

(i) Deduce all others' behaviors from the mechanism's extensive form.

(ii) Condition on those behaviors when comparing alternatives.

In this subsection, we show that agents can instead rely on a heuristic notion of rationality that relaxes task (i) while still guaranteeing to avoid dominated strategies. We begin with the following result.[11]

Proposition 1. *Every tree admits a coarsest information partition.*

Given a mechanism $(X, \mathcal{I})$, let $\mathscr{I}$ denote the coarsest information partition on $X$. A behavior $\sigma_n \in \Sigma_n$ is *coarse* if it is measurable with respect to $\mathscr{I}$. Let $\Sigma(\mathscr{I})$ be the

[10] That is, there are $\sigma_n \in \Sigma_n$ and $h' \in \mathcal{I}(h)$ with $h' \prec z(\sigma_n, \sigma_{-n})$.

[11] Our proof of Proposition 1 relies on Zorn's Lemma and does not need any finiteness assumption. In contemporaneous work, Cohen and Li (2025) provide a constructive proof for the finite case.

set of all coarse behavior profiles. Then, a natural way to simplify task (i) is to only deduce the set $\Sigma_{-n}(\mathscr{I})$ of *coarse behaviors* for the others. Indeed, deducing the set of coarse behaviors minimizes agent $n$'s contingent reasoning effort while being able to explain all observable play, in the following sense. First, $\Sigma_{-n}(\mathscr{I})$ is the smallest set of behaviors compatible with some information partition on $X$: deducing it requires no reasoning about what others can condition their behavior on. Second, despite this minimality, coarse behaviors are rich enough to account for all observable play.[12]

PROPOSITION 2. *Given any behavior profile, it is possible to replace the behaviors of an arbitrary set of agents with coarse behaviors without affecting which terminal history is reached.*

We now show that focusing on the coarse behaviors of the others can be a useful heuristic. 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 profiles $\sigma_{-n} \in \Sigma_{-n}(\mathscr{I})$. Let $\mathscr{C}$ denote the set of coarsely optimal strategy profiles. We say that agent $n$ is *coarsely rational* if, whenever $\mathscr{C}_n$ is nonempty, she plays a coarsely optimal strategy. Importantly, although coarsely optimal strategies are best responses to all coarse behaviors of others, they need not themselves be coarse: they may condition on the (non-hypothetical) information an agent actually observes along the path of play. We now state the main result of this subsection.

THEOREM 1. *Every coarsely optimal strategy is undominated.*

Thus, whenever a coarsely optimal strategy exists, achieving coarse rationality suffices to ensure play rationally. Further, by Proposition 2, it is follows that, as play unfolds, coarsely optimal strategies always remain potential best replies. Despite these appealing properties, if coarse rationality were not able to systematically guarantee optimal play, agents may find it worthwhile to engage in more sophisticated strategic thinking. However, Theorem 1 implies that this is not an issue for dominant strategy mechanisms, in which coarse rationality guarantees to best respond. Together with its simplicity advantages, this makes coarse rationality a compelling guide for behavior in dominant strategy mechanisms.

Observe that while coarse rationality allows agents to have a simpler understanding of what others can do, alleviating task (i), it may still require substantial contingent reasoning. Indeed, it does not simplify task (ii), as agents still have to correctly condition on each coarse behavior profile of the others to compare a coarsely optimal

[12] In fact, the coarse and full behavior space (of the others) have the same projection onto histories, with the former neglecting interactions between behaviors that remain hypothetical. This is similar in spirit to the behavioral ideas underlying (sequential) cursed equilibrium and analogy-based expectation equilibrium (cf. Cohen and Li, 2025; Eyster and Rabin, 2005; Jehiel, 2005).

strategy to any potential deviations. In fact, coarse rationality coincides with standard rationality in static dominant strategy mechanisms, in which deviations from theoretical predictions have long been documented (e.g., Esponda and Vespa, 2023; Kagel and Levin, 2001, 2009; Li, 2017). In the next subsection, we identify a design strategy that further simplifies the contingent reasoning required to behave coarsely rational, by reducing the burden of task (ii).

### 3.2 The transparency dilemma

We now show that more transparent mechanisms—i.e., those that reveal more information about others' past moves along the path of play—reduce the contingent reasoning required to recognize a coarsely optimal strategy. At the same time, we show that increasing the transparency of a mechanism is usually incompatible with maintaining dominant strategy incentives. This raises the question of whether coarse rationality can be expected to be adopted as a heuristic, leading to a design dilemma.

To ease exposition, we focus on mechanisms that share the same tree in the main text. In Appendix B, we extend the analysis to mechanisms that do not share a tree but nonetheless provide agents with an *equivalent experience*.[13]

Definition 1. A mechanism is *more transparent* than another if they share the same tree and the information partition of the former is strictly finer than that of the latter.

Additional transparency may be desirable for several reasons. For instance, using more transparent mechanism is the key to achieving design goals such as auditability, verifiability, trustworthiness, and credibility (see, e.g., Akbarpour and Li, 2020; Banchio et al., 2025; Grigoryan and Möller, 2023; Hakimov and Raghavan, 2025; Pycia and Ünver, 2020; Saitto, 2026; Woodward, 2020). In what follows, we focus on the simplicity advantages. In particular, we argue that transparency reduces the contingent reasoning required for coarse rationality.

In what follows, fix a coarsely optimal strategy $s_n$, a type $\theta_n$ and a history $h$ at which $\theta_n$ can move under $s_n(\theta_n)$. We evaluate the contingent reasoning effort to recognize $s_n$ as coarsely optimal on the basis of two assumptions that mirror assumptions commonly adopted in the simplicity literature (cf. Li, 2017). First, upon reaching $h$, agent $n$ tests the coarse optimality of $s_n(\theta_n)$ only against continuation behaviors that first deviate from $s_n(\theta_n)$ at $h$ and coarse behaviors of the others that are consistent with reaching $\mathcal{I}(h)$. Formally, a *test* for coarse rationality is a pair

[13] For example, the more general analysis allows us to compare mechanisms such as a sealed-bid second-price auction and an ascending clock auction, or a static and sequential random dictatorship. Intuitively, comparison is possible if the mechanisms are representable on the same tree.

$(\sigma_n, \sigma_{-n}) \in \Sigma_n \times \Sigma_{-n}(\mathscr{I})$ such that $\sigma_n(h) \neq s_n(\theta_n)(h)$ and $z(\sigma_n, \sigma_{-n})$ visits $\mathcal{I}(h)$. Second, while agents understand how each terminal history maps into outcomes, they find it costly to identify the terminal histories reached by two alternative continuations, for a fixed conjecture about others' behavior. Intuitively, as she considers different tests $\sigma$, she deduces the true map $\sigma \mapsto (z(\sigma), z(s_n(\theta_n), \sigma_{-n}))$.[14] Formally, let $\mathscr{Z}$ be the set of all functions $\zeta$ mapping tests into pairs $(z', z'')$ of terminal histories visiting $\mathcal{I}(h)$, with the property that $\zeta(\sigma') = \zeta(\sigma'')$ whenever $z(\sigma') = z(\sigma'')$ and $z(s_n(\theta_n), \sigma'_{-n}) = z(s_n(\theta_n), \sigma''_{-n})$. Given a set T of tests, a map $\zeta \in \mathscr{Z}$ is T$-$*consistent* if $\zeta(\tau) = (z(\tau), z(s_n(\theta_n), \tau_{-n}))$ for all $\tau \in$ T. Together, these two assumptions give rise to a natural notion of sets sufficient to verify coarse optimality. A *sufficient test set* is a set T of tests such that $\omega(\tau) \leq_{\theta_n} \omega(s_n(\theta_n), \tau_{-n})$ for all $\tau \in$ T implies $\omega(z') \leq_{\theta_n} \omega(z'')$ for all $(z', z'')$ such that there is a test $\sigma$ and a T$-$consistent map $\zeta \in \mathscr{Z}$ with $\zeta(\sigma) = (z', z'')$. Smaller sufficient test sets require lower contingent reasoning effort, as $\theta_n$ must reason through fewer hypothetical contingencies to verify coarse optimality of $s_n(\theta_n)$.

Proposition 3 (The benefits of transparency). *If a mechanism $\Gamma^A$ is more transparent than a mechanism $\Gamma^B$, the following holds true.*

*(i) If $s_n$ is coarsely optimal in $\Gamma^B$, it is also coarsely optimal in $\Gamma^A$.*

*(ii) For every coarsely optimal strategy $s_n \in \mathscr{C}_n^B$, type $\theta_n$ and history $h$ at which agent $n$ may move under $s_n(\theta_n)$:*

*(a) Every sufficient test set for $\theta_n$ in $\Gamma^B$ admits a subset that is a sufficient test set for $\theta_n$ in $\Gamma^A$, with proper inclusion whenever $\mathcal{I}^A(h) \neq \mathcal{I}^B(h)$ and $\theta_n$ is not indifferent among all outcomes that may still realize after $\mathcal{I}^B(h)$ is reached.*

*(b) The set of obviously dominated deviating actions in $\Gamma^B$ is a subset of that in $\Gamma^A$.*

Intuitively, increasing transparency makes it easier to identify and verify a coarsely optimal strategy. Part (i) of the proposition ensures all coarsely optimal strategies stay coarsely optimal. Part (ii.a) shows that verifying coarse optimality requires conditioning on fewer hypothetical contingencies at every point.Part (ii.b) shows that, additionally, transparency enlarges the set of deviations that can be recognized as dominated without conditioning on others' behavior. A practical implication of Proposition 3 is that greater transparency can make it easier for agents to follow coarse rationality. In particular, we would predict more transparent dominant-strategy mechanisms—in which coarse rationality ensures optimal play—to exhibit higher rates of rational

[14] Observe that since she tests $s_n$ against deviations only when they actually depart from it, we can analyze the verification problems at successive histories separately without loss of generality.

behavior. This is consistent with a large set of existing experimental evidence. In particular, Breitmoser and Schweighofer-Kodritsch (2022) and Esponda and Vespa (2023) show that additional transparency increases rates of rational play even within the already restrictive class of obviously dominant strategy mechanisms. To the best of our knowledge, such findings cannot be captured by any existing notion of simplicity.

Unfortunately, the next result shows that increases in transparency rarely preserve dominant-strategy incentives. Given a mechanism $(X, \mathcal{I})$, we say that a finer information $\mathcal{I}'$ partition for $X$ *separates* two histories $h'$ and $h''$ if it places them in different information sets, $\mathcal{I}'(h') \neq \mathcal{I}'(h'')$, while $\mathcal{I}(h') = \mathcal{I}(h'')$ under the original partition.

Proposition 4 (The costs of transparency). *Let $s$ be a dominant strategy equilibrium of some mechanism $(X, \mathcal{I})$. Then, the following statements hold true.*

*(i) A strategy $s_n$ remains dominant for agent $n$ under a finer information partition $\mathcal{I}'$ that separates two histories $h'$ and $h''$ at which some agent $m \neq n$ moves only if in $(X, \mathcal{I})$,*

$$\omega(\sigma_n, \sigma'_m, \sigma_{-nm}) \preceq_{\theta_n} \omega(s_n(\theta_n), \sigma''_m, \sigma_{-nm})$$

*whenever $h' \preceq z(\sigma_n, \sigma'_m, \sigma_{-nm})$ and $h'' \preceq z(s_n(\theta_n), \sigma''_m, \sigma_{-nm})$.*

*(ii) Unless $s$ is obviously dominant when all information sets are singletons, there exists a more transparent mechanism $(X, \mathcal{I}')$ for which $s$ remains an ex-post equilibrium but is no longer dominant.*

On the intensive margin, part (i) shows dominance can be preserved under transparency increases only if a restrictive local obvious dominance condition holds. For instance, in a sealed-bid second-price auction, revealing any information about a bidder's bid to another bidder breaks dominance. On the extensive margin, part (ii) shows transparency is bounded above unless an obviously dominant strategy mechanism can be obtained.[15] Such mechanisms do not exist for many applications (see, e.g., Ashlagi and Gonczarowski, 2018; Bade and Gonczarowski, 2016; Thomas, 2020; Troyan, 2019). Moreover, obviously dominant strategy mechanisms typically require substantial coordination of play, which may be impractical.[16]

Taken together, Propositions 3 and 4 generate a nontrivial design dilemma: increasing the transparency makes coarse rationality less demanding, yet is often incompatible with maintaining dominant strategies. This poses a problem, since coarse

[15] This closely relates to Theorem 6 in Mackenzie (2020), who shows that, in a mechanism of perfect information, a strategy is dominant if and only if it is obviously dominant.

[16] Pycia and Troyan (2023) show that obviously dominant strategy mechanisms are equivalent to generalized millipede games, in which agents move sequentially and potentially many times.

rationality alone can no longer guarantee optimal play, and the heuristic loses its normative force. Moreover, the natural alternative, ex-post equilibria are widely documented to be poor predictors of behavior.[17] At the same time, a growing experimental literature suggests that increasing transparency beyond the dominant-strategy can in fact improve adherence to theoretical predictions (Bó and Hakimov, 2020; Bó and Hakimov, 2023; Hakimov and Raghavan, 2025; Kagel and Levin, 2009).

How should this dilemma be resolved? In the next section, we show that there exists a broader class of mechanisms for which coarse rationality is a compelling heuristic, and which allows for substantially greater transparency.

## 4 As-If Dominance

We argued that dominant strategy mechanisms make coarse rationality a compelling heuristic, but severely restrict transparency. In this section, we identify a richer class of mechanisms that permits significantly greater transparency, while ensuring that coarse rationality is a self-supporting heuristic.

### 4.1 As-if dominance as self-supporting coarse rationality

For a given mechanism, we say that coarse rationality is a *self-supporting heuristic* if

(i) each agent has at least one coarsely optimal strategy; and

(ii) every coarsely optimal strategy profile is an ex-post equilibrium.

Observe that condition (ii) is nontrivial since coarsely optimal strategies need not themselves be coarse. A mechanism satisfying (i) and (ii) in which all agents adopt coarse rationality to guide behavior makes this heuristic as compelling as dominance: it both prevents playing dominated strategies and ensures that agents best respond to one another. We call such mechanisms as-if dominant strategy mechanisms.

Definition 2. A mechanism is an *as-if dominant strategy mechanism* if coarse rationality is a self-supporting heuristic. Its coarsely optimal strategies are called as-if dominant.

Intuitively, if coarse rationality is self-supporting and all agents adopt it as a guide for behavior, then no agent can profit from unilaterally refining her beliefs about others' behavior. In this sense, the common adoption of coarse rationality forms a "meta equilibrium:" (i) it has a non-vacuous prescription, and (ii) no agent can unilaterally

[17] See Crawford et al. (2013) for a survey of experimental evidence on non-equilibrium behavior and related theoretical work.

profit from being more sophisticated. Observe that as-if dominant strategy mechanisms can also be characterized as those in which the set of strategy profiles robust to coarse deviations by other agents forms a product set of ex-post equilibria. In particular, this implies a connection with level-$k$ thinking (Stahl and Wilson 1995; Nagel 1995; de Clippel et al. 2019). Taking level-0 behavior to be coarse, assume level-$k$ agents either (i) best-respond to a full-support belief over all behaviors that can be played by level-$(k-1)$ agents and preferences are sufficiently rich, or (ii) best-respond to a full-support belief over all behaviors that can be played by some lower level.[18] Then, as-if dominance ensures that it suffices for all agents to be of at least level-1 to guarantee that an ex-post equilibrium is played.

In Section 7, we further show that—under regularity conditions satisfied in our applications—as-if dominant strategy mechanisms are those in which coarse rationality emerges endogenously from an adaptive generational learning process.

A direct corollary of Theorem 1 is that every dominant strategy mechanism is as-if dominant. In fact, the dominant strategy outcomes coincide with the as-if dominant strategy ones. However, the class of as-if dominant strategy mechanisms is richer and allows for substantially greater transparency. Indeed, unlike dominance, coarse optimality of a strategy is preserved under transparency increases (cf. Proposition 3). Consequently, as-if dominance is preserved under transparency increases as long as the set of coarsely optimal strategy profiles continues to form a product set of ex-post equilibria.

### 4.2 Two illustrative examples

Before turning to applications, we present two examples. The first illustrates the greater flexibility of as-if dominance relative to dominance.

Example 1 (Open-outcry English auction). Several bidders are in a room bidding for an item, for which each has a private value $\theta_n \in \{0, 1, \ldots, v\}$. Bidders have quasi-linear payoffs and prefer winning the item at a price equal to their value to losing it. Starting from zero, the auctioneer asks for offers in $\{1, \ldots, v\}$.[19] At the end of each round of offers, the item is tentatively assigned to one of the highest bidders, with ties broken according to some predetermined order. Bidders can only submit offers exceeding their previous ones. If at some point no new offer is made, the item is assigned to the tentative winner at her last offer. Since bidders observe each others' actions and can engage in jump bidding, this auction *is not* a dominant strat-

[18] Richness, as formally introduced in Section 7, requires that each coarse behavior for agent $n$ is coarsely optimal for some type $\theta_n \in \Theta_n$. This condition holds in all applications we consider below.

[19] Bidders respond simultaneously.

egy mechanism.[20] However, it is an as-if dominant strategy mechanism: a coarsely rational bidder always makes the minimum offer that can make her the tentative winner, provided this offer is below her valuation. This provides a justification for the long-standing practice—under private values—of modeling the English auction as (equivalent to) a second-price sealed-bid auction, by focusing on cutoff strategies (see, e.g., Li, 2017; McAfee and McMillan, 1987; Milgrom, 1989).

As-if dominance is thus strictly more permissive than dominance. As the next example shows, at the same time, it is strictly more demanding than ex-post incentive compatibility: not every mechanism admitting a coarsely optimal ex-post equilibrium is as-if dominant.

Example 2 (Sequential second-price auction). Several bidders are in the same room bidding for an item, for which each has a private value $\theta_n \in \{0, 1, \dots, v\}$. Bidders have quasi-linear payoffs. The auctioneer approaches bidders sequentially, asking each of them to publicly announce a bid $b_n \in \{0, \dots, v\}$. After all bids are announced, the item is assigned to the highest bidder at a price equal to the second-highest bid, with ties broken according to the order of play. In this mechanism, each bidder has a coarsely optimal strategy: bid her true valuation. Moreover, truthful bidding by all agents is an ex-post equilibrium. However, coarse rationality is not self-supporting, because many other coarsely optimal strategies exist. For instance, it is coarsely optimal to bid truthfully if no previous bid exceeds one's valuation, and zero otherwise. If the first bidder recognizes that this (non-coarse) behavior is available to later bidders, she may submit a bid of $v$ to deter competition and obtain the item for free. Conversely, a bidder may fear that later bidders could retaliate by raising the payment, tempting her to shade. Thus, the mere existence of a coarsely optimal strategy profile that is an ex-post equilibrium does not ensure that coarse rationality is self-supporting.

## 5 Asynchronous Auctions

In many auction settings, it is impractical to gather all agents together to elicit their preferences. In fact, to facilitate participation, auctions increasingly allow agents to choose themselves when to do so. Prominent examples of such asynchronous formats are online auctions, which allow agents to bid at their most convenient time within a given time window. This presents a relevant constraint for practical market design. Indeed, it rules out the use of clock auctions, which have been praised for

[20] Intuitively, this is because others could condition their behavior on the size of price increments.

their strategic simplicity and leaves designers without a clear benchmark for simple asynchronous formats. In this section, we apply our framework to fill this gap. In particular, we characterize the auction format adopted by prominent online platform as maximally transparent among the set of asynchronous second-price auctions.

### 5.1 Auction environment

Let $Y$ be a finite set of feasible allocations $y$, where each allocation specifies a binary assignment $y_n \in \{0,1\}$ for every agent $n$. Implementing $y \in Y$ entails a cost $c(y) \geq 0$. The zero allocation is always feasible and costless. There is also a nonempty finite set $D$ of feasible dates at which agents can bid. An outcome in $\Omega$ specifies an allocation $y$, a vector $p$ of payments $p_n \in [0,\infty)$, and a vector $\delta$ of sets $\delta_n \subseteq D$ of bidding dates. An agent's type $\theta_n$ is identified with a pair $(v_n, W_n)$, where $v_n \in V$ represents the agent's willingness to pay to be included in the allocation, and $W_n \in \mathcal{W}$ is a preference ordering over $\mathcal{D}$, the collection of all non-empty subsets of dates at which she may interact with the auction. We assume that $V$ is a finite subset of the nonnegative reals containing zero, and that $\mathcal{W}$ contains an ordering under which all subsets of $D$ are equally ranked, and let $\Theta_n := V \times \mathcal{W}$. Each type's preferences $\preceq_{\theta_n}$ over outcomes depend only on her quasi-linear payoff $v_n y_n - p_n$ and the set of access dates $\delta_n \in \mathcal{D}$, and are separable in these two components.

The auctioneer seeks to implement a (Pivot) VCG outcome but, to ensure participation, is restricted to *asynchronous auctions*, which accommodate agents' preferences over when to bid. Intuitively, these are auctions in which bidders can choose when and how often to enter the auction, and revise their bids upward over time.

Definition 3. An *asynchronous VCG auction* is a mechanism $(X, \mathcal{I})$ satisfying the following. The tree $X$ is such that, along every terminal history $z \in Z$:

1. Chance moves and selects a reserve price $r$ from some set $R$, where $R = \{0\}$ unless the set of feasible allocations is $\{y \mid \sum_n y_n \leq 1\}$, in which case $R$ is an arbitrary nonempty subset of the real numbers. Then, in some order, each agent $n$ chooses a first *access date* (or to never access).

2. For all $d \in D$, from earliest to latest, if some agent chose $d$ as her access date, chance determines an order in which all such agents are called to play.

    a. Upon moving, an agent $n$ may submit a new *standing bid* $b_n \in V$ that strictly exceeds her previous bids. If $b_n < \max V$, she also chooses either a new access date that strictly exceeds $d$ or to never enter again. For all bidders, the initial standing bid is zero.

b. Chance assumes a bidder's value equals her standing bid if it weakly exceeds $r$ and zero otherwise. Then, it picks a *tentative allocation* among all those maximizing social surplus, unless the previous tentative allocation maximizes social surplus, in which case it is maintained. The initial tentative allocation is the zero one. A standing bid is *winning* if the agent who submitted it is included in the tentative allocation, and *losing* otherwise.

3. Once the recursive process described in 2. has terminated, the auction ends. The final tentative allocation is implemented. Each bidder included in the allocation is charged the maximum of her VCG payment—assuming her private value equals her last standing bid—and the reserve price $r$. Others are charged zero.

The information partition $\mathcal{I}$ is such that:

4. Agents select their first access dates simultaneously.

5. During the recursive process described in 2., the following holds for all $d \in D$.

   a. All agents who selected $d$ as their access date move simultaneously.

   b. If agent $n$ does not submit a new standing bid upon accessing at $d$, no other agent can ever observe she accessed the auction at $d$.

   c. If agent $n$ accesses at $d$, she can never observe whether another agent $m$ previously chose to access again at $d' \geq d$ or to never access again.

   d. Agents cannot observe actions that are equivalent to not submitting a bid. That is, if agent $n$ submits a standing bid upon accessing at $d$ that she knows cannot affect the final allocation, no other agent observes her access or bid. All other bids are said to be *admissible*.

An asynchronous VCG auction in which *no* information about previous bids is disclosed to bidders corresponds to a standard VCG auction and therefore admits a dominant strategy equilibrium. Notably, the reserve price is the only information that can be revealed while preserving dominant-strategy incentive compatibility—any further information breaks the dominant-strategy property. At the other extreme, in an asynchronous second-price auction in which *all* information about previous bids is disclosed, truthful bidding is only an ex-post equilibrium. Between these extremes there are many other disclosure policies an online platform may adopt. In what follows, we explore the extent to which as-if dominance can relax the transparency constraints dominance imposes in two different settings.

### 5.2 Selling an item online

We first consider the design problem of an online platform selling a single indivisible item to multiple potential bidders. In this context, an allocation is feasible if and only if $\sum y_n \leq 1$, and implementation is costless, i.e., $c(y) \equiv 0$. The asynchronous VCG auction therefore reduces to a *asynchronous second-price auctions* : at any point in time, there is at most one winning bid, and corresponding winning bidder. We begin by characterizing the set of as-if dominant strategy formats available to the platform.

Proposition 5. *An asynchronous second-price auction is an as-if dominant strategy mechanism if and only if the following holds for all $\beta$: Until it is certain the final selling price is at least $\beta$, no agent observes whether someone else previously submitted a bid of at least $\beta$.*

In Appendix G, we show that all as-if dominant strategy asynchronous second-price auctions implement the same outcome as a standard sealed-bid auction—possibly with a reserve. Thus, they achieve both efficiency and—in a regular and symmetric independent-private-value environment (see, e.g., Elkind, 2007)—profit maximization. From Proposition 5, we conclude that the restrictions imposed by as-if dominance on the information partition of the asynchronous second-price auction concern only what is disclosed about the current winning bid. This is in sharp contrast to the restriction imposed by traditional dominance, under which only the reserve price can be revealed. In practice, many disclosure policies reveal strictly more than the reserve price while respecting the bounds implied by Proposition 5. One example is what we call the *Online English Auction*. This format, while novel to the theoretical literature, closely resembles those widely used by prominent online platforms such as eBay.[21]

Definition 4. An asynchronous second-price auction is an *Online English Auction* if its information partition is such that, whenever an agent moves, she observes only the following (in addition to her own previous moves).

(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 values of all previously submitted admissible losing bids.

[21]Other platforms using similar formats 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.gsaauctions.gov/), all accessed on February 10, 2024. For exposition, we abstract away from the possibility that bidders may revise their maximum bid downwards as long as it remains above the currently displayed price. This variant is as-if dominant, but, by design, prevents reliably revealing information regarding the highest bid, which only simplifies our characterization.

THEOREM 2. *The Online English Auction achieves maximal transparency within the set of as-if dominant strategy asynchronous second-price auction .*

Theorem 2 provides a theoretical rationale for the popularity of this format. The disclosure policy preserves strong incentives for truthful bidding while maximizing transparency subject to the as-if dominant strategy constraint. As discussed in Section 3, this additional transparency can be advantageous both in terms of simplicity and trust in the platform. Together with the above-mentioned fact that as-if dominance ensures the Online English Auction can achieve both efficiency and profit maximization, in Online Appendix G we show that it admits equilibria in which private sellers have no incentive for shill bidding. Taken together, these results show that the Online English Auction offers a natural alternative to clock auctions when requiring synchronous participation is impractical.

### 5.3 Funding a public project

We now turn to the problem of designing an asynchronous VCG auction to fund a non-excludable public project. Examples are the construction of a bridge or the renewal of a public park. In this environment, $Y$ contains only two allocations: The status quo, with $y_n \equiv 0$, and the public project being undertaken, with $y_n \equiv 1$ for all $n$. For simplicity, we assume $V := \{0, 1, \dots, \overline{v}\}$ for some positive integer $\overline{v}$, and that the project entails a fixed integer cost $c > 0$. Moreover, we assume agents are *relatively small*, $\overline{v} + |N| - 2 \leq c < (|N| - 1)\overline{v}$. In words, if all but at most two agents value the project at most one, the project is strictly efficient only if both remaining agents have positive value for it. However, the project is always strictly efficient if all but at most one agent have maximum value for it. The designer's objective is to implement the socially efficient allocation while minimizing expenditures. Accordingly, we restrict attention to mechanisms without a reserve.

PROPOSITION 6. *If agents are relatively small, an asynchronous VCG auction for a public project is as-if dominant strategy if and only if, whenever agent $n$ moves, she cannot exclude that the sum of all the standing bids of the others is at most $c - \overline{v} + 1$, unless it is already certain that the project is undertaken and all agents $i \neq n$ pay nothing.*

As for the previous subsection, all as-if dominant strategy asynchronous VCG auctions implement the same outcome as a standard VCG auction. However, the above theorem illustrates that the transparency constraints imposed by as-if dominance depend sharply on the setting. When funding a public project with relatively small agents, the gap between as-if dominance and traditional dominance is considerably smaller than when selling an item online. In particular, in the former, no agent can

ever exclude any pair $(y, p_n)$ initially associated with a given coarse strategy, limiting the benefits of additional transparency beyond the dominant strategy benchmark.

## 6 Other Applications

In this section, we consider three additional settings to illustrate the potential usefulness of as-if dominant strategy mechanisms that are strictly more transparent than their dominant strategy counterparts: multi-unit auctions, as well as one-sided and two-sided matching. Our results are consistent with a range of experimental evidence, to which we thus provide a unified explanation.

### 6.1 Efficient multiunit auctions

Suppose there are $L \geq 2$ units of a homogeneous good to be auctioned among at least two agents (e.g., pollution permits or water rights). An outcome in $\Omega$ is a pair $(y, p)$ specifying the units $y_n \in \{0, 1, \ldots, L\}$ assigned to each agent and the payment $p_n \in [0, \infty)$ she makes, with $\sum_n y_n \leq L$. An agent's type space $\Theta_n$ consists of all vectors $\theta_n$ of $L$ non-increasing marginal values $\theta_n^\ell \in \{0, 1, \ldots, u\}$, where $u$ is a positive integer. Agent $n$'s preferences $\leq_{\theta_n}$ over outcomes are represented by the standard quasi-linear payoff function $(y, p) \mapsto \sum_{\ell \leq y_n} \theta_n^\ell - p_n$.

As shown by Gkatzelis et al. (2017), in this environment no obviously dominant strategy mechanism can implement the efficient allocation.[22] Together with Proposition 4, this implies that a designer who seeks to allocate the units efficiently can only offer only limited transparency if she insists on using dominant strategy mechanisms. We thus apply our framework to derive practical design guidance for this setting.

Ausubel (2004) proposes a class of ascending multiunit auctions implementing the same outcome as the classic sealed-bid VCG auction (see Ausubel, 2004; Nagel and Saitto, 2025 for formal definitions). After chance selects a priority order among all permutations of agents, the auction proceeds in rounds. Agents are initially endowed with a demand of $L$ units. In each round $r \in \{0, \ldots, u-1\}$, they report how many units they remain (weakly) willing to buy for a unit price of $r+1$. At the end of round $r$, an agent with positive demand wins a number of units equal to the increase in non-negative net-of-her-demand excess supply, at a price of $r$. The auction terminates at the end of round $r < u-1$ if there is excess supply. If excess supply is strict, units are rationed at price $r$ according to the priority order, subject to the constraint that agent $n$ cannot be assigned more units than the ones she was demanding at the start

[22] In Online Appendix F, we provide a direct proof focusing on the VCG outcome. The impossibility also follows from Proposition 7.3 in Gkatzelis et al. (2017).

of round $r$. If there is still excess demand at the end of round $u-1$, all remaining units are rationed at price $u$, once again following the priority order and subject to the constraint that no bidder can be assigned more units than she last demanded.

As discussed in Ausubel (2004), such an auction can be implemented with different information disclosure policies. In the *Ausubel auction with private dropouts*, agents never observe any information apart from the round they are currently in. This is the only information partition which makes the mechanism dominant strategy. By contrast, in the *Ausubel auction with public dropouts* reports are made simultaneously within rounds, but agents can observe their competitor's past demand reports. In both formats, agents do not observe the priority order chosen by chance. Our framework predicts that the simplicity benefits of this additional transparency need not be outweighed by the loss of dominant-strategy implementability.

Proposition 7. *The Ausubel auction with public dropouts is an as-if dominant strategy mechanism, and truthful reporting in all rounds is the unique as-if dominant strategy.*

Kagel and Levin (2009) compare the standard sealed-bid VCG auction, the Ausubel auction with private dropouts, and the version with public dropouts. Consistent with our prediction, they find significantly higher rates of truthful bidding in the latter. Concretely, in the first 12 auctions with $L = 2$ units, truthful bidding on the higher-valued unit among non-winning bidders is 18% in the sealed-bid Vickrey auction, 43% in the Ausubel auction with private dropouts, and 83% when dropouts are public. Results are similar for the lower-valued unit, with the gap narrowing somewhat as bidders gain experience. As Kagel and Levin (2009) suggest, the additional transparency provided by the dropout information appears to simplify the bidding task, even at the cost of dominant-strategy implementability. Proposition 7 provides a theoretical justification: increased transparency preserves coarse rationality as a self-supporting heuristic, while making coarsely optimal strategies easier to identify.

### 6.2 School choice and housing allocation

We now consider two classical matching environments: school choice and housing allocation (cf. Gale and Shapley, 1962; Shapley and Scarf, 1974). In both, agents' type spaces consist of all strict preference orderings over a set of outcomes.

In *school choice problems*, outcomes are schools—each endowed with a capacity and a priority ranking over agents—and an outside option.[23] We focus on *student-optimal*

[23]We model schools as nonstrategic. Their strict rankings over agents are selected by chance over all possible ones (e.g., as a function of zip codes, sibling priority, and a random tie breaker). We assume a school's ranking can deem an agent as unacceptable.

*stable matchings*: The agents' most preferred matching among those in which no agent prefers another school that also prefers her to at least one currently assigned student. In *housing problems*, outcomes are houses, each initially owned by a distinct agent. We focus on *core allocations*: The unique allocation such that no coalition of agents can can profitably deviate by reallocating their houses among themselves instead.

In both settings, there are direct mechanisms in which truthful reporting is the unique dominant strategy equilibrium and that implement the outcomes we are interested in. In particular, once agents have simultaneously reported their preference orderings, a centralized clearinghouse can use the *Deferred Acceptance* (DA) and *Top Trading Cycle* (TTC) algorithms, to find the student-optimal stable matching and the core allocation, respectively (cf. Dubins and Freedman, 1981; Gale and Shapley, 1962; Roth, 1982; Shapley and Scarf, 1974). In contrast, Ashlagi and Gonczarowski (2018) and Troyan (2019) show that neither outcome can be implemented by an obviously dominant strategy mechanism, implying that there is once again a binding limit on the transparency attainable in dominant strategy mechanisms. This is practically relevant: experimental evidence shows that even when the direct mechanisms are explained in advance, participants often fail to recognize truth-telling as dominant (see, e.g., Chen and Sönmez, 2006; Klijn et al., 2019; Rees-Jones, 2018). What if we instead use more transparent mechanisms, asking agents to take active part in the algorithm?

The *sequential DA* and *sequential TTC* algorithms work as follows. All agents begin unmatched. In each round, unmatched agents simultaneously report a most preferred outcome to apply to among those they have not yet applied to. Within a round of sequential TTC, an agent is called to play repeatedly until she applies to a house that is not yet permanently assigned.

(i) In sequential DA, each school tentatively accepts applicants up to capacity, rejecting its least favorite applicants (old or new) when necessary. If an agent applies to her outside option, she is permanently matched to it.

(ii) In sequential TTC, any agent belonging to a cycle of applications is permanently assigned the house she applied to. The applications of all other agents are kept on hold unless that house is assigned, in which case they are rejected.

In both mechanism, at end of each round, rejected agents become (or stay) unmatched. If no agent is rejected, the mechanism ends and the current matching is implemented. The information partitions are as follows. When called to play in sequential DA, an agent can observe everything that occurred in previous rounds.[24] In

[24] Agents can only infer schools' preferences from previous rounds. The mechanism can be adapted to reveal from the start if an agent is unacceptable to a school (e.g., via public cutoffs).

sequential TTC, an agent can observe all and only the applications the others submitted to houses that were not yet permanently assigned.[25] While neither mechanism has a dominant strategy equilibrium (cf. Klijn et al., 2019), we have the following result.

Proposition 8. *In their respective environments, the sequential DA and TTC algorithms are as-if dominant strategy mechanisms. Moreover, in both mechanisms, truthfully reporting in all rounds is the unique as-if dominant strategy.*

The sequential mechanisms we proposed are implement the same outcomes as their dominant strategy counterparts in as-if dominant strategies, but are significantly more transparent.[26] Thus, our framework predicts that increasing transparency up to the level provided by these mechanisms should improve adherence to the theoretical predictions. This is consistent with, and provides a unified theoretical justification for, a growing body of experimental evidence (see, e.g., Bó and Hakimov, 2020, 2022; Klijn et al., 2019) documenting superior performance under sequential mechanisms relative to their dominant strategy alternatives.

## 7 Learning Coarse Rationality

We defined as-if dominant strategy mechanisms as those in which coarse rationality is self-supporting. In this section, we show that, under some additional regularity conditions satisfied in all our applications, as-if dominance guarantees coarse rationality emerges from a particular adaptive inter-generational learning process. We then prove a relationship between as-if dominance and dominance solvability.

### 7.1 Preliminary definitions

The main results of this section restrict attention to as-if dominant strategy mechanisms which satisfy a few regularity conditions, always satisfied in our applications. To state them, it is convenient to introduce some preliminary concepts.

#### *7.1.1 Time linearity*

Intuitively, time-linearity means the information partition is such that agents cannot infer additional information about the status of play by looking at their watch. That

[25] Observe that sequential TTC would have the same reduced normal form if agents were restricted to apply to houses that have not yet been assigned. In fact, the former is obtained from the latter by adding redundant actions.

[26] As we discuss in Appendix B, both the static and sequential mechanisms can be represented on the same tree, and thus can be formally compared using our more general transparency order.

is, we do not implicitly assume agents cannot condition their behavior on the passing of time. Formally, the information partition of a mechanism $(X, \mathcal{I})$ is *linear in time* if there is a partition $\mathcal{Q}$ of the set of non-terminal histories satisfying the following.

(i) If a time cell $Q \in \mathcal{Q}$ and some $I \in \mathcal{I}$ are not disjoint, then $I \subseteq Q$.

(ii) Within a time cell, play is simultaneous: If different players move at $h \in Q$ and $I \subseteq Q$, then $h$ precedes $I$ if and only if $(h, a)$ precedes $I$ for all $a \in A(h)$.

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

A sufficient condition for time linearity is that information sets "do not cross:" If $h' \in I'$ precedes $h'' \in I''$, then no $h''' \in I''$ precedes some $h'''' \in I'$.

### *7.1.2 Conditional coarse optimality and strategic equivalence*

We now define coarse optimality conditional on reaching a certain branch of a mechanism $(X, \mathcal{I})$, and two notions of strategic equivalence. Formally, two information sets at which agent $n$ moves belong to the same *branch* $\mathbb{I}_n$ if agent $n$'s experience prior to reaching them is the same.[27] Let $\Sigma(\mathbb{I}_n)$ be the set of behaviors $\sigma$ that visit some member of $\mathbb{I}_n$, and observe that we have $\Sigma(\mathbb{I}_n) = \Sigma_n(\mathbb{I}_n) \times \Sigma_{-n}(\mathbb{I}_n)$ by perfect recall.

Then, we say that $\sigma'_n \in \Sigma_n(\mathbb{I}_n)$ is $\mathbb{I}_n$*−coarsely optimal* for $\theta_n \in \Theta_n$ if

$$\omega(\sigma_n, \sigma_{-n}) \preceq_{\theta_n} \omega(\sigma'_n, \sigma_{-n})$$

for all coarse $\sigma_{-n} \in \Sigma_{-n}(\mathbb{I}_n) \cap \Sigma_{-n}(\mathscr{S})$, and all $\sigma_n \in \Sigma_n(\mathbb{I}_n)$. Observe that this definition does allow agent $n$ to condition on the different information sets in $\mathbb{I}_n$.

For $\sigma'_n, \sigma''_n \in \Sigma_n(\mathbb{I}_n)$, we write $\sigma'_n \sim_{\mathbb{I}_n} \sigma''_n$ if they are *strategically equivalent conditional on reaching* $\mathbb{I}_n$: that is, $\omega(\sigma'_n, \sigma_{-n}) =_\theta \omega(\sigma''_n, \sigma_{-n})$ for all $\sigma_{-n} \in \Sigma_{-n}(\mathbb{I}_n)$ and $\theta \in \Theta$.

## 7.2 As-if dominance and coarse rationality as a learning outcome

Our approach is related to the idea that agents can progressively learn to avoid iteratively dominated strategies (see, e.g., Marx, 1999). Consider a sequence of generations, each composed of infinitely many replicas of the incomplete information game defined by $\Gamma$. Let $S(0) := S$. The play of generation $\gamma \in \mathbb{N} := \{1, 2, \ldots\}$ is a *restriction* of the strategy space, i.e. a nonempty rectangular subset $S(\gamma)$ of $S$. Let $\Sigma(\gamma)$ be the set of behavior profiles $\sigma \in \Sigma$ such that $\sigma = (s(\theta), \sigma_c)$ for some strategy profile $s \in S(\gamma)$, type profile $\theta \in \Theta$ and chance behavior $\sigma_c \in \Sigma_c$.

[27] That is, they are preceded by the same sequence of information sets and corresponding played actions. Observe that the set of all branches partitions the information sets.

Definition 5. Given a mechanism $(X, \mathcal{I})$, a sequence of play $(S(\gamma))_{\gamma \in \mathbb{N}}$ is *consistent with generational learning* if it satisfies the following conditions.

Full Sampling of Recent Savvy Best Replies. For all $\gamma \in \mathbb{N}$, the set $S_n(\gamma)$ contains all strategies $s_n \in S_n(\gamma - 1)$ such that:

(i) For each $\theta_n \in \Theta_n$, there is $\sigma_{-n}^{\theta_n} \in \Sigma_{-n}(\gamma - 1)$ such that $\omega(\sigma_n, \sigma_{-n}^{\theta_n}) \leq_{\theta_n} \omega(s_n(\theta_n), \sigma_{-n}^{\theta_n})$ for all $\sigma_n \in \Sigma_n$.

(ii) There are no $\theta_n \in \Theta_n$ and $\sigma_n \in \Sigma_n(\gamma - 1)$ with $\omega(s_n(\theta_n), \sigma_{-n}) \leq_{\theta_n} \omega(\sigma_n, \sigma_{-n})$ for all $\sigma_{-n} \in \Sigma_{-n}(\gamma - 1)$, with the inequality sometimes strict.

Eventual Savviness. For all $\underline{\gamma} \in \mathbb{N}$, there is $\overline{\gamma} > \underline{\gamma}$ such that, for all $\gamma' > \overline{\gamma}$, if $s_n \in S_n(\gamma')$, then there are no $\sigma_n \in \Sigma_n$ and $\theta_n \in \Theta_n$ with

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

for all $\sigma_{-n}$ for which there is some $\gamma$ with $\underline{\gamma} \leq \gamma < \gamma'$ such that $\sigma_{-n} \in \Sigma_{-n}(\gamma)$, with the inequality sometimes strict.

A path of play consistent with generational learning is *regular* if $S_n(\gamma)$ contains no dominated strategy, for all agents $n$ and generations $\gamma$.

The full sampling condition should be interpreted as an implication of generations being large: If a particular strategy is optimal for some belief consistent with the previous generation's play, and no alternative strategy has proven consistently better, there will be agents in the current generation whose beliefs make it optimal to play that strategy. Eventual savviness closely resembles the adaptive learning ideas from Milgrom and Roberts (1990) and Marx (1999). Intuitively, it requires that strategies which are consistently outperformed by some alternative are eventually discarded. Regularity assumes agents start from a rational base, and learn a simpler heuristic as generations pass. Observe that (regular) generational learning does not rely on the assumption that agents know others' payoffs, but only that they can observe past play. For instance, they may have access to discussion forums (online auctions), or generational knowledge (school applications). We next show that under some regularity conditions, as-if dominance ensures coarse rationality emerges from generational learning, in the sense that it eventually becomes both positively and normatively compelling.

Definition 6. Given a mechanism, coarse rationality *emerges from (regular) generational learning* if there is at least one (regular) sequence of play $(S(\gamma))_\gamma$ consistent with generational learning and, for all such sequences, there is $\gamma'$ such that for all $\gamma > \gamma'$:

(i) All and only coarsely optimal strategies are played, i.e. $S(\gamma) = \mathscr{C}$.

(ii) Coarse rationality is necessary and sufficient for optimal play. That is, $s_n \in \mathscr{C}_n$,

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

for all $\theta \in \Theta$, profiles $s_{-n} \in S_{-n}(\gamma)$ and behaviors $\sigma_n \in \Sigma_n$ and $\sigma_c \in \Sigma_c$, with at least one inequality holding strict whenever $\sigma_n$ is not coarsely rational for $\theta_n$.

Observe that as-if dominance is a necessary condition for coarse rationality to emerge. The main result of this section is a partial converse to this statement. We call an as-if dominant strategy mechanism *regular* if agents are indifferent among all their as-if dominant strategies whenever the others play undominated strategies.

THEOREM 3. *Consider an as-if dominant strategy mechanism satisfying the following.*

TIME-LINEARITY. *The information partition is linear in time.*

RICHNESS. *For all $\sigma_n \in \Sigma_n(\mathscr{I})$, there is $\theta_n \in \Theta_n$ with $s_n(\theta_n) = \sigma_n$ for some $s_n \in \mathscr{C}_n$.*

PERFECTION. *For all agents $n$, types $\theta_n \in \Theta_n$ and branches $\mathbb{I}_n$, there is at least one $\sigma_n \in \Sigma_n(\mathbb{I}_n)$ that is $\mathbb{I}_n-$coarsely optimal for $\theta_n$. Moreover, for all such behaviors $\sigma_n$,*

*(i) there is $\theta'_n \in \Theta_n$ and $s_n \in \mathscr{C}_n$ such that $s_n(\theta'_n) \in \Sigma_n(\mathbb{I}_n)$ and $s_n(\theta'_n) \sim_{\mathbb{I}_n} \sigma_n$, and*

*(ii) for all $\sigma'_n \in \Sigma_n(\mathbb{I}_n)$,*

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

*for all $\theta_{-n} \in \Theta_{-n}$, $s_{-n} \in \mathscr{C}_{-n}$ and $\sigma_c \in \Sigma_c$ such that $(s_{-n}(\theta_{-n}), \sigma_c) \in \Sigma_{-n}(\mathbb{I}_n)$.*

*If there is a unique as-if dominant strategy profile, then coarse rationality emerges from generational learning. Moreover, as long as the mechanism is regular, coarse rationality emerges from regular generational learning.*

The richness condition ensures that the type space is rich enough for every coarse behavior to be coarsely optimal for some type. Perfection requires the mechanism to be everywhere as-if dominant, even off-path. In particular, if an agent deviates from as-if dominance, sub-condition (i) ensures that no new as-if dominant behavior arises, while (ii) ensures that as-if dominance is preserved. In all dominant strategy mechanisms satisfying richness, coarse rationality emerges from regular generational learning. Moreover, the three regularity conditions in Theorem 3 hold in all of our applications. In the mechanisms from Section 6, there is a unique as-if dominant strategy profile, while those from Section 5 are regular.

Corollary 1. *In all the as-if dominant strategy mechanisms from Section 6, coarse rationality emerges from generational learning. In the Online English Auction, coarse rationality emerges from regular generational learning.*

### 7.3 As-if dominance and (weak) dominance solvability

In this subsection, we explore the relationship between as-if dominance and (weak) dominance solvability. It turns out that the conditions under which coarse rationality emerges from generational learning are also (almost) sufficient for dominance solvability. Before stating the result, we formally define dominance solvability.

Given a finite mechanism, an *order of elimination of dominated strategies* is a decreasing sequence $(S^k)_k$ of nonempty product sets $S^k \subset S$ with $S^0 := S$ and such that (i) for all $k$ and $n$, all $s_n \in S_n^{k-1} \setminus S_n^k$ are weakly dominated in the game defined by $S^{k-1}$, and (ii) there are no weakly dominated strategies in the game defined by $\bigcap_k S^k$. An order of elimination is *regular* if $S^1$ contains no strategy that is dominated in $S$.

A finite mechanism $(X, \mathcal{I})$ is *(regular) dominance solvable* if all (regular) orders of eliminations of dominated strategies terminate in the same product set $S^\infty \subseteq S$ such that any pair of strategy profiles $s', s'' \in S^\infty$ is *outcome-equivalent*, i.e.

$$\omega(s'(\theta), \sigma_c) =_\theta \omega(s''(\theta), \sigma_c)$$

for all chance behaviors $\sigma_c \in \Sigma_c$ and type profiles $\theta \in \Theta$. In this case, the set $S^\infty$ is called the *(regular) solution* of the mechanism.

Theorem 4. *Consider an as-if dominant strategy mechanism satisfying time-linearity, richness and perfection. Then, the following two claims hold true.*

*(i) If the as-if dominant strategy profile is unique, then the mechanism is weak-dominance solvable, and its solution is the as-if dominant strategy profile.*

*(ii) If the mechanism is regular and all as-if dominant strategy profiles are outcome-equivalent, then it is regular weak-dominance solvable, and its regular solution is the set of as-if dominant strategy profiles.*

The conditions of case (i) are satisfied by all application settings in Section 6, while the conditions of case (ii) are satisfied by the ones in Section 5. Hence, similarly to the previous subsection, we obtain the following corollary.

Corollary 2. *The as-if dominant strategy mechanisms from Section 6 are dominance solvable, and their solution is their as-if dominant strategy outcome. The Online English Auction is regular dominance solvable, and its regular solution is its as-if dominant strategy outcome.*

In particular, the above result implies Theorem 2 in Ausubel (2004), which states that the Ausubel auction with public dropouts is weak dominance solvable. Additionally, our result also implies that both the sequential TTC and DA algorithm are weak dominance solvable. To the best of our knowledge, these are novel findings.

## 8 Conclusion

We show agents can resort to a simplifying heuristic, coarse rationality, to avoid dominated strategies. The common adoption of coarse rationality is self-supporting—that is, no one can ever gain from unilaterally being more sophisticated—in the class of as-if dominant strategy mechanisms. These include dominant strategy mechanisms, but allow for mechanisms with significantly more transparency, which we show makes coarse rationality simpler to follow. Moreover, we show that, under some conditions met in our applications, as-if dominance implies that coarse rationality emerges from a generational learning process, and implies weak dominance solvability.

We characterize the Online English Auction as maximally transparent within the class of as-if dominant strategy asynchronous second-price auctions. This format, while novel to the literature, matches the design choices of prominent online platforms such as eBay. Our results thus provide a theoretical rationale for the popularity of this auction format in settings where gathering agents together to collect their bids is impractical. We also apply our framework to several other canonical settings, from multiunit auctions to matching problems. In particular, we show that the Ausubel auction with public dropouts and the sequential implementations of the DA and TTC algorithms are as-if dominant strategy mechanisms. This lets us provide a unified explanation for a wide range of experimental evidence. First, our results are consistent with the evidence that more transparent dominant strategy mechanisms are simpler to play, even when starting with an obviously dominant strategy one (Breitmoser and Schweighofer-Kodritsch, 2022; Esponda and Vespa, 2014; Li, 2017). Second, they are consistent with a growing experimental literature reporting how additional transparency can improve predictive reliability even when it breaks dominant strategy incentives (Bó and Hakimov, 2020; Bó and Hakimov, 2023; Hakimov and Raghavan, 2025; Kagel and Levin, 2009; Klijn et al., 2019).

## References

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

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.

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*.

Banchio, M., Yang, F., and Skrzypacz, A. (2025). Dynamic threats to credible auctions. Working paper.

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.

Cohen, S. and Li, S. (2025). Sequential cursed equilibrium.

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 Iriberri, 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.

Esponda, I. and Vespa, E. (2023). Contingent thinking and the sure-thing principle: Revisiting classic anomalies in the laboratory. *Review of Economic Studies, forthcoming*.

Eyster, E. and Rabin, M. (2005). Cursed equilibrium. *Econometrica*, 73(5):1623–1672.

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.

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.

Hakimov, R. and Raghavan, M. (2025). Improving transparency and verifiability in school admissions: Theory and experiment. *Management Science*, 0(0):null.

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. (2001). Behavior in multi-unit demand auctions: Experiments with uniform price and dynamic vickrey auctions. *Econometrica*, 69(2):413–454.

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. (2024). Are simple mechanisms optimal when agents are unsophisticated? Working paper.

Li, S. (2017). 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.

Marx, L. M. (1999). Adaptive learning and iterated weak dominance. *Games and Economic Behavior*, 26(2):253–278.

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.

Milgrom, P. (1989). Auctions and bidding: A primer. *Journal of Economic Perspectives*, 3(3):3–22.

Milgrom, P. and Roberts, J. (1990). Rationalizability, learning and equilibria in games with strategic complementarities. *Econometrica*, 58(6):1255–1277.

Myerson, R. B. (1981). Optimal auction design. *Mathematics of Operations Research*, 6(1):58–73.

Nagel, L. and Saitto, R. (2025). A measure of complexity for strategy-proof mechanisms.

Nagel, R. (1995). Unraveling in guessing games: An experimental study. *The American Economic Review*, 85(5):1313–1326.

Okamoto, N. (2018). An efficient ascending-bid auction for multiple objects: Comment. *The American Economic Review*, 108(2):555–560.

Pearce, D. G. (1984). Rationalizable strategic behavior and the problem of perfection. *Econometrica*, 52(4):1029–1050.

Pycia, M. and Troyan, P. (2023). Obviously strategyproof mechanisms in general environments. Working paper.

Pycia, M. and Ünver, M. U. (2020). Arrovian efficiency and auditability in the allocation of discrete resources.

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.

Saitto, R. (2026). Credible multiunit auction design. Working paper.

Schummer, J. and Velez, R. A. (2021). Sequential preference revelation in incomplete information settings. *American Economic Journal: Microeconomics*, 13(1):116–47.

Shapley, L. and Scarf, H. (1974). On cores and indivisibility. *Journal of Mathematical Economics*, 1(1):23–37.

Stahl, D. O. and Wilson, P. W. (1995). On players' models of other players: Theory and experimental evidence. *Games and Economic Behavior*, 10(1):218–254.

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.

Wang, W. and Zheng, Z. (2025). On the dominance of truth-telling in gradual mechanisms.

Woodward, K. (2020). Self-auditable auctions.

Zhang, L. and Levin, D. (2017). Bounded rationality and robust mechanism design: An axiomatic approach. *American Economic Review*, 107(5):235–39.

## A Extensive-Form Mechanisms

A *mechanism* is a pair $\Gamma = (X, \mathcal{I})$, consisting of a tree $X$ and an information partition $\mathcal{I}$ for $X$, which we now define formally. A *tree* is a triple $X = (H, P, \omega)$ composed of:

(i) A set $H$ of sequences $h$ from some set $A$ of *actions*, called *histories*. It is downward closed with respect to the *precedence relation* defined as $h' \preceq h''$ if and only if $h'' = (h', h)$. If $h'$ is an infinite sequence and $H$ contains all $h \prec h'$, then $h' \in H$. It contains a unique *root*, denoted $\varnothing$, with no predecessor.

  (a) The set $Z \subseteq H$ of *terminal histories* is the set of histories $z$ with no successor.

  (b) For all $h \in H \setminus Z$, the set $A(h)$ of *available actions* contains all $a \in A$ with $(h, a) \in H$. It always contains at least two elements.

(ii) A map $P\colon H \setminus Z \to N \cup \{c\}$. Agent $n$ is *called to play*, or *moves*, at $h$ if $P(h) = n$.

(iii) An *outcome function* $\omega\colon Z \to \Omega$.

An *information partition* for $X$ is a partition $\mathcal{I}$ of $H \setminus Z$, whose elements are called *information sets* and satisfy the following. Let $\mathcal{I}(h)$ denote the information set containing $h$. First, if $\mathcal{I}(h') = \mathcal{I}(h'')$, then $P(h') = P(h'')$ and $A(h') = A(h'')$. Second, if $P(h) = c$, then $\mathcal{I}(h)$ is a singleton. Third, agents have *perfect recall*. Formally, let agent $n$'s *experience* at $h$ be the sequence of all information sets and actions that she respectively moves at and takes along $h$, in the order they occur. Then, agent $n$'s experience at any two given histories at which she is called to play is the same whenever they

lie in the same information set. In the text and whenever not specified otherwise, we restrict attention to *finite* mechanisms, i.e. such that $H$ is finite.

Given a mechanism, a *behavior* for agent $n$ is a rule $\sigma_n$ mapping each history at which agent $n$ moves into an available action that is $\mathcal{I}-$*measurable*: If $\mathcal{I}(h') = \mathcal{I}(h'')$, then $\sigma_n(h') = \sigma_n(h'')$. Behaviors for the chance player are defined analogously. We extend behaviors from histories to information sets in the obvious way. A behavior profile $\sigma = (\sigma_n, \sigma_{-n})$ specifies a behavior for each agent and a behavior for the chance player. Let $z(\sigma)$ be the terminal history reached by $\sigma$, and define $\omega(\sigma) := \omega(z(\sigma))$. A *strategy* for agent $n$ is a map $s_n$ from $\Theta_n$ into her behavior space $\Sigma_n$.

## B Extending the Transparency Ranking

We extend the transparency partial order to mechanisms that, while not sharing the same tree, provide an *equivalent experience* to agents. This allows to compare, for instance, clock auctions with sealed-bid second-price auctions, and the static and sequential matching algorithms from Section 6. It is also indifferent to how simultaneous play is represented on the tree. All results from the main text extend to this generalized ranking. No finiteness assumption is needed for this section.

### B.1 Equivalent experiences

To formally define experiential equivalence, we introduce some additional pieces of notation. Fix a mechanism $\Gamma$. A history is *virtually terminal* if all its weak non-terminal successors lie in singleton information sets, and all terminal ones map into the same outcome. Let $\tilde{z}\colon \Sigma \to H$ map each behavior profile $\sigma \in \Sigma$ into the unique virtually terminal history preceding $z(\sigma)$ that does not itself have a virtually terminal predecessor. Two behaviors $\sigma_n', \sigma_n'' \in \Sigma_n$ are *virtually identical* if $\tilde{z}(\sigma_n', \sigma_{-n}) = \tilde{z}(\sigma_n'', \sigma_{-n})$ for all $\sigma_{-n} \in \Sigma_{-n}$. Let $\varsigma_n(\sigma_n)$ denote the set of behaviors that are virtually identical to $\sigma_n$. We call such an equivalence class a *virtual behavior* and denote a generic one by $\varsigma_n$. The same construction is applied to the chance player. We use $\varsigma$ to denote virtual behavior profiles, and $\varsigma_{-n}$ to denote virtual behavior profiles for the others. As we show next, the maps $\tilde{z}$ and $\omega$ can naturally be extended to virtual behavior profiles.

Proposition B.1. *If $h$ is not preceded by a virtually terminal history, then a behavior profile $\sigma$ satisfies $h \preceq z(\sigma)$ if and only if $h \preceq z(\sigma')$ for all virtually identical behavior profiles $\sigma'$.*

A virtual profile $\varsigma_{-n}$ *separates* virtual behaviors $\varsigma_n'$ and $\varsigma_n''$ if $\tilde{z}(\varsigma_n', \varsigma_{-n}) \neq \tilde{z}(\varsigma_n'', \varsigma_{-n})$. If $\varsigma_{-n}$ separates $\varsigma_n'$ and $\varsigma_n''$, let $\tilde{\Sigma}_{-n}(\varsigma_n', \varsigma_n'', \varsigma_{-n})$ be defined as the set of all virtual profiles of the others compatible with reaching the information set containing the last history weakly preceding both $\tilde{z}(\varsigma_n', \varsigma_{-n})$ and $\tilde{z}(\varsigma_n'', \varsigma_{-n})$.

Definition B.1. Two mechanisms $\Gamma^1$ and $\Gamma^2$ are *experientially equivalent* if there is a pair $(\phi, \gamma)$ of product maps $\phi = (\phi_n)_n$ and $\gamma = (\gamma_n)_n$ which *transforms* the former into the latter. That is, the following conditions hold true.

(i) Each $\phi_n$ bijects $n$'s two virtual behavior spaces, with $\omega^2 \circ \phi \equiv \omega^1$.

(ii) Each $\gamma_n$ bijects $n$'s two coarse behavior spaces, with $\varsigma_n^2 \circ \gamma_n \equiv \phi_n \circ \varsigma_n^1$.

(iii) For all agents, $\varsigma_{-n}$ separates $\varsigma_n'$ and $\varsigma_n''$ if and only if $\phi_{-n}(\varsigma_{-n})$ separates $\phi_n(\varsigma_n')$ and $\phi_n(\varsigma_n'')$. Moreover, whenever separation occurs,

$$\phi_{-n}(\tilde{\Sigma}_{-n}^1(\varsigma_n', \varsigma_n'', \varsigma_{-n})) \equiv \tilde{\Sigma}_{-n}^2(\phi_n(\varsigma_n'), \phi_n(\varsigma_n''), \phi_{-n}(\varsigma_{-n})).$$

Intuitively, condition (i) requires the two mechanisms to have virtually identical behavior spaces. In particular, whenever an agent has to separate two behaviors, she has the same information about what others can do. Condition (ii) additionally requires they have the same unconditional structure, and in particular that two actions correspond to the same *intention* in one mechanism if and only if this is the case in the other. This captures, for instance, the idea that transforming standard sequential matching pennies into a game in which the second mover wants to match what the first one does if and only if the latter plays heads changes the rules of the game, as well as agents' experience. Finally, condition (iii) requires agents have the same information about what others can do whenever separating two of their behaviors. A sufficient condition for two mechanisms to be experientially equivalent is to differ in how they order simultaneous play only.

Lemma B.1. *Experiential equivalence is reflexive, symmetric and transitive. Moreover, it preserves as-if dominance.*

### B.2 The extended transparency ranking

We now extend our transparency partial order to an *extended transparency ranking*.

Definition B.2. A mechanism $\Gamma^1$ is *(weakly) more transparent* than a mechanism $\Gamma^2$ if there exists a tree $X$ and information partitions $\mathcal{I}^1$ and $\mathcal{I}^2$ satisfying the following.

(i) The mechanism $\Gamma^\ell$ is experientially equivalent to $(X, \mathcal{I}^\ell)$, for all $\ell \in \{1, 2\}$.

(ii) The information partition $\mathcal{I}^1$ is (weakly) finer than $\mathcal{I}^2$.

We start proving that the ranking is well-defined, in the following formal sense.

Lemma B.2. *Fix two mechanisms $\Gamma^1$ and $\Gamma^2$. Suppose that $(X^a, \mathcal{I}^\ell)$ and $(X^b, \mathcal{J}^\ell)$ are mechanisms experientially equivalent to $\Gamma^\ell$, for all $\ell \in \{1,2\}$ . Then, $\mathcal{I}^1$ is (weakly) finer than $\mathcal{I}^2$ if and only if $\mathcal{J}^1$ is (weakly) finer than $\mathcal{J}^2$.*

This proves that the extended transparency ranking reduces to the partial order considered in the main text when we fix the tree. More generally, the following holds.

Theorem B.1. *The extended transparency ranking is a partial preorder defined on the proper class of all extensive-form mechanisms.*

Using the extended transparency ranking, we can compare all mechanisms considered in the main text. A standard clock auction is more transparent than the corresponding strategically equivalent second-price sealed bid auction. They can be represent on the same tree, where first chance picks a priority order and then bids are elicited in rounds of binary questions. Regardless of what others do, an agent always continues to be called to play until they refuse to further raise their bid. Rounds are simultaneous, and no information about the priority order or previous rounds is provided, with the exception that, in the clock auction, when at the end of some round at most one bidder is left standing this is publicly revealed. Similarly, the multiunit VCG auction is less transparent than the Ausubel auction with private dropouts, which is less transparent than its public dropouts version. Moreover, the sequential matching mechanisms from Section 6 are more transparent than their static counterpart.

## C Proofs

In this section, we provide proofs for the main results of the paper. All other omitted proofs, including those of several lemmas we state here, are in the Online Appendix.

### C.1 Proof of Proposition 1

We provide a proof using Zorn's lemma, which does not rely on any finiteness assumptions. In contemporaneous work, Cohen and Li (2025) provide a constructive proof for the finite-length case. The proof relies on the following lemmas.

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

Lemma C.2. *The finest common coarsening of two information partitions is an information partition itself.*

We now prove 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. By Lemma C.1, the information partition $\mathcal{I}^C$ that satisfies $\mathcal{I}^C(h) := \{h' \in H \mid \text{there is } \mathcal{I}' \in C \text{ such that } h' \in \mathcal{I}'(h)\}$ for all non-terminal $h$ is an upper bound for $C$. Hence, by Zorn's Lemma, there is a maximal information partition $\mathcal{I}$. Next, we show $\mathcal{I}$ is the (unique) coarsest information partition for $X$. Suppose not: Then, there is an information partition $\mathcal{I}'$ and a 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 $\mathcal{I}$. By Lemma C.2, $\mathcal{I}' \vee \mathcal{I}$ is an information partition for $X$ that is strictly coarser than $\mathcal{I}$, contradicting its maximality. *Q.E.D.*

### C.2 Proof of Proposition 2

The proof does not require any finiteness assumption. Fix a mechanism $(X, \mathcal{I})$, a behavior profile $\sigma \in \Sigma$, and a subset $M \subseteq N$ of agents. We construct a profile $\sigma' \in \Sigma$ that agrees with $\sigma$ for all agents $n \notin M$ and for chance, prescribes a coarse behavior for all agents $n \in M$, and satisfies $z(\sigma') = z(\sigma)$. For each $n \in M$, define $\sigma'_n$ as follows: at every information set $\mathscr{I}(h)$ at which $n$ moves along $z(\sigma)$, set $\sigma'_n(\mathscr{I}(h)) = \sigma_n(h)$, and prescribe an arbitrary action at all other information sets of $\mathscr{I}$ at which $n$ moves. By perfect recall, agent $n$ visits each information set of $\mathscr{I}$ at most once along $z(\sigma)$, and all histories in $\mathscr{I}(h)$ share the same action set, so this is well-defined. By construction, $\sigma'_n$ is $\mathscr{I}-$measurable and hence coarse, and agrees with $\sigma_n$ at every history preceding $z(\sigma)$ at which $n$ is called to play. Since the terminal history is fully determined by the actions taken along the path of play, we have $z(\sigma') = z(\sigma)$. *Q.E.D.*

### C.3 Proof of Theorem 1

The theorem is a consequence of the following lemma. It holds for arbitrary type spaces and all mechanisms of *conditionally finite length*, i.e. such that, for all $\sigma_c \in \Sigma_c$, there is a finite integer $\ell(\sigma_c)$ such that, for all $\sigma_{-c} \in \Sigma_{-c}$, history $z(\sigma_{-c}, \sigma_c)$ has length at most $\ell(\sigma_c)$. In particular, no cardinality assumption on the mechanism's width is needed. Given a mechanism $\Gamma$, we say $\sigma_n$ is *undominated* for $\theta_n$ if there is no $\sigma'_n$ such that $\omega(\sigma_n, \sigma_{-n}) \leq_{\theta_n} \omega(\sigma'_n, \sigma_{-n})$ for all $\sigma_{-n} \in \Sigma_{-n}$, with at least one strict inequality.

Lemma C.3. *Consider a mechanism of conditionally finite length. If agent $n$ has coarsely optimal strategy $s_n$ with $s_n(\theta_n) = \sigma_n$, then $\sigma_n$ is an undominated behavior for $\theta_n$.*

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

Let $B_{-n}$ be the set of behaviors for the others, $\sigma_{-n}$, such that $\omega(\sigma_n'', \sigma_{-n}) <_{\theta_n} \omega(\sigma_n', \sigma_{-n})$, i.e., such that $\theta_n$ strictly prefers to follow $\sigma_n'$ over $\sigma_n''$. Since $\sigma_n'$ dominates $\sigma_n''$, the set $B_{-n}$ is not empty. Moreover, there is some $\sigma_c$ for which $B_{-n}$ is not empty even when restricting to behavior profiles in which chance's behavior is $\sigma_c$. Slightly abusing notation, let $B_{-n}$ be as defined above but fixing chance's behavior to $\sigma_c$.

Next, for all $\sigma_{-n}$, let $\eta(\sigma_{-n})$ 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 $h' \in \mathscr{I}(h'')$, then $\sigma_{-n}(h') = \sigma_{-n}(h'')$. Thus, $\eta(\sigma_{-n})$ is the length of the longest history $h \preceq z(\sigma_n', \sigma_{-n})$ for which $\sigma_{-n}$ is consistent with reaching both $(\sigma_n', \sigma_{-n})$ and $(\sigma_n'', \sigma_{-n})$ while behaving coarsely at all histories preceding $h$. Observe that, for all $\sigma_{-n} \in B_{-n}$, it must be that $\eta(\sigma_{-n})$ is strictly less than the length of $z(\sigma_n', \sigma_{-n})$. Otherwise, we can construct a coarse $\sigma_{-n}^c \in \Sigma_{-n}(\mathscr{I})$ with $\sigma_{-n}^c(h) = \sigma_{-n}(h)$ 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 $z(\sigma_n, \sigma_{-n}^c) = z(\sigma_n, \sigma_{-n})$ for all $\sigma_n \in \{\sigma_n', \sigma_n''\}$, contradicting coarse optimality.

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

Next, let $h'$ be the unique strict predecessor of $z(\sigma_n', \sigma_{-n})$ of length $\ell$. By construction, $h'$ is well defined, neither chance nor agent $n$ is called to play at $h'$, and there is $h'' \prec z(\sigma_n'', \sigma_{-n})$ such that $\mathscr{I}(h') = \mathscr{I}(h'')$ and $\sigma_{-n}(h') \neq \sigma_{-n}(h'')$. Let $\sigma_{-n}'$ and $\sigma_{-n}''$ be defined such that $\sigma_{-n}'(h) = \sigma_{-n}''(h) = \sigma_{-n}(h)$ at all $h \notin \mathscr{I}(h')$ at which $n$ is not called to play, but $\sigma_{-n}'(\mathscr{I}(h')) = \sigma_{-n}(h')$ while $\sigma_{-n}''(\mathscr{I}(h')) = \sigma_{-n}(h'')$. Since $\mathscr{I}$ satisfies perfect recall, we have $z(\sigma_n', \sigma_{-n}') = z(\sigma_n', \sigma_{-n})$ and $z(\sigma_n'', \sigma_{-n}'') = z(\sigma_n'', \sigma_{-n})$. Further, by construction, we have $\ell < \eta(\sigma_{-n}')$ and $\ell < \eta(\sigma_{-n}'')$, since these behaviors behave coarsely at least one more step than $\sigma_{-n}$. Then, by definition of $\eta$, it must be that $\{\sigma_{-n}', \sigma_{-n}''\}$ and $B_{-n}$ are disjoint sets: Since $\sigma_n'$ dominates $\sigma_n''$ for type $\theta_n$, it follows

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

Finally, observe that $\sigma_{-n}(h') \neq \sigma_{-n}(h'')$ implies that $\mathcal{I}(h') \neq \mathcal{I}(h'')$. Thus, we can define $\sigma_{-n}'''$ 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 $\sigma_{-n}'''(h'') = \sigma_{-n}(h')$. 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 $\theta_n$, behavior $\sigma_n''$ is dominated by $\sigma_n'$. *Q.E.D.*

The proof then follows from observing that we consider finite mechanisms, which are a special case fo conditionally finite-length mechanisms, and that a strategy is undominated if and only if it always prescribes an undominated behavior. *Q.E.D.*

### C.4 Proof of Proposition 3

Given a history $h$ at which agent $n$ is called to play, a behavior $\sigma_n \in \Sigma_n$ is compatible with reaching $h$ if $(h', \sigma_n(h')) \preceq h$ for all $h' \prec h$ at which $n$ moves. A profile $\sigma_{-n} \in \Sigma_{-n}$ is compatible with reaching $h$ if there exists $\sigma_n$ such that $h \prec z(\sigma_n, \sigma_{-n})$. Moreover, $\sigma_{-n} \in \Sigma_{-n}$ is compatible with reaching $\mathcal{I}(h)$ if there is $h' \in \mathcal{I}(h)$ such that $\sigma_{-n}$ is compatible with reaching $h'$. We use the following lemma.

Lemma C.4. *Every mechanism sharing the same tree has the same set of coarse behaviors. Moreover, the set of coarse behavior profiles of the others compatible with reaching an information set $\mathcal{I}(h)$ at which $n$ moves is strictly increasing in $\mathcal{I}(h)$ with respect to set inclusion.*

Given the first part of Lemma C.4, statement (i) follows since behaviors in $\Gamma^B$ are also measurable in $\Gamma^A$. From now on, suppose $\theta_n$ can move at $h$ under $s_n$, which is coarsely optimal in $\Gamma^B$. Moreover, to simplify notation, whenever conditioning on a history $h$ at which $n$ is called to play, we treat two behaviors $\sigma_n', \sigma_n''$ which only differ on histories that are not weak successors of $h$ as the same behavior. Similarly, we treat profiles $\sigma_{-n}$ which only differ on histories that do not visit $\mathcal{I}(h)$ as the same.

#### *C.4.1 Proof of statement (ii)(a)*

If $\theta_n$ is indifferent among all outcomes that can still realize after $\mathcal{I}^B(h)$ is reached, the claim holds trivially as no test is needed to verify coarse optimality. Thus, suppose this is not the case, and let $\mathrm{T}^B$ be a sufficient test set in $\Gamma^B$.

Claim C.1. *For all $h' \in \mathcal{I}^B(h)$, there is some $\tau \in \mathrm{T}^B$ such that $h' \prec z(\tau)$.*

Proof. Suppose the statement is false. Then, there is $h' \in \mathcal{I}(h)$ such that no $\tau \in \mathrm{T}^B$ satisfies $h' \prec z(\tau)$. Let $\bar{z}$ and $\underline{z}$ be two terminal histories visiting $\mathcal{I}^B(h)$ with $\omega(\underline{z}) <_{\theta_n} \omega(\bar{z})$. Let $\sigma' = (\sigma_n', \sigma_{-n}') \in \Sigma_n \times \Sigma_{-n}(\mathscr{I})$ be such that $\sigma_n'(h') \neq s_n(\theta_n)(h')$ and $h' \prec z(\sigma')$. Next, let $\zeta^B$ lead to the correct terminal histories, except for tests that generate the same pair as $\sigma'$. In that case, set $\zeta^B(\tau') = (\bar{z}, \underline{z})$. Observe that $\zeta$ is $\mathrm{T}^B-$consistent, as there is no $\tau \in \mathrm{T}^B$ such that $h' \prec z(\tau)$, and instead $h' \prec z(\sigma')$. *Q.E.D.*

Now, let $\mathrm{T}^{B\setminus A}$ be the set of all $\tau \in \mathrm{T}^B$ such that $\tau_{-n}$ is compatible with reaching $h' \in \mathcal{I}^B(h) \setminus \mathcal{I}^A(h)$, and define $\mathrm{T}^A := \mathrm{T}^B \setminus \mathrm{T}^{B\setminus A}$.

Claim C.2. *The set $\mathrm{T}^A$ is a sufficient test set in $\Gamma^A$.*

Proof. Since $\Gamma^A$ is more transparent than $\Gamma^B$, all $\tau \in \mathrm{T}^A$ are tests in $\Gamma^A$. Now suppose $\omega(\tau) \leq \omega(s_n(\theta_n), \tau_{-n})$ for all $\tau \in \mathrm{T}^A$ and let $\sigma'$ be a test in $\Gamma^A$. Then, $\sigma'_{-n} \in \Sigma_{-n}(\mathscr{S})$ is compatible with reaching some history $h' \in \mathcal{I}^A(h)$. By Proposition 2, we can choose a coarse $\sigma''_n$ such that $z(\sigma''_n, \sigma'_{-n}) = z(\sigma')$. Then, $\sigma'' = (\sigma''_n, \sigma'_{-n})$ is a test in both $\Gamma^A$ and $\Gamma^B$. Next, let $\zeta^A$ be $\mathrm{T}^A-$consistent. Then, we can define $\zeta^B$ such that, for all tests $\sigma$ in $\Gamma^B$, we have $\zeta^B(\sigma) = \zeta^A(\sigma)$ if $\sigma_{-n}$ is consistent with reaching $\mathcal{I}^A(h)$, and $\zeta^B(\sigma) = (z(\sigma), z(s_n(\theta_n), \sigma_{-n}))$ otherwise. By construction, $\zeta^B$ is $\mathrm{T}^B-$consistent. In fact, no test $\tau \in \mathrm{T}^B$ such that $\tau_{-n}$ is not compatible with reaching $\mathcal{I}^A(h)$ constrains $\zeta^B$, and the latter coincides with the true map on tests that pass though $\mathcal{I}^B(h) \setminus \mathcal{I}^A(h)$. Then, we can write $\zeta^A(\sigma') = \zeta^A(\sigma'') = \zeta^B(\sigma'') = (z', z'')$. Since $\mathrm{T}^B$ is a sufficient test set in $\Gamma^B$, it follows that $\omega(z') \leq_{\theta_n} \omega(z'')$. *Q.E.D.*

By Claim C.1, $\mathrm{T}^A$ is strictly smaller than $\mathrm{T}^B$ whenever $\mathcal{I}^B(h) \neq \mathcal{I}^A(h)$. By Claim C.2, the statement follows. *Q.E.D.*

### *C.4.2 Proof of statement (ii)(b)*

Suppose that in $\Gamma^A$, the deviation $a \in A(h) \setminus \{s_n(\theta_n)(h)\}$ is not obviously dominated by $s_n(\theta_n)$ for $\theta_n$. Then, there is some $\sigma_{-n} \in \Sigma^A_{-n}$ and terminal history $z$ with $h' \prec z$ for some $h' \in \mathcal{I}(h)$ such that $\omega(s_n(\theta_n), \sigma_{-n}) <_{\theta_n} \omega(z)$. By Proposition 2, we can assume without loss that $\sigma_{-n}$ is coarse. Moreover, since $\mathcal{I}^B$ is coarser than $\mathcal{I}^A$, we have $h' \in \mathcal{I}^B(h)$. Then, $a$ is not obviously dominated in $\Gamma^B$. *Q.E.D.*

## C.5 Proof of Proposition 4

### *C.5.1 Proof of statement (i)*

By perfect recall and since $h'$ and $h''$ are initially in the same information set, for all deviations $\sigma_n$, profiles $\sigma_{-nm}$ for agents other than $n$ and $m$, and behaviors $\sigma'_m$ and $\sigma''_m$ satisfying the stated conditions, there is a single behavior $\sigma'''_m$ with $\omega(\sigma_n, \sigma'''_m, \sigma_{-nm}) = \omega(\sigma_n, \sigma'_m, \sigma_{-nm})$ and $\omega(s_n(\theta_n), \sigma'''_m, \sigma_{-nm}) = \omega(s_n(\theta_n), \sigma''_m, \sigma_{-nm})$. *Q.E.D.*

### *C.5.2 Proof of statement (ii)*

Lemma C.5. *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 $(X, \mathcal{I}')$.*

We now prove the statement using an argument that mirrors the proof of Theorem 6 Mackenzie (2020). Let $s$ be a dominant strategy equilibrium for $(X, \mathcal{I})$. By

Lemma C.5, $s$ is an ex-post equilibrium when $X$ is equipped with its finest information partition. It thus suffices show that if $\mathcal{I}$ has all singleton information sets, then $s$ is obviously dominant. Indeed, 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 $\sigma_{-n}$, we have that $h$ is reached under $(s_n(\theta_n), \sigma'_{-n})$ and $(s_n(\theta_n), \sigma''_{-n})$ if and only if $\sigma'_{-n}$ and $\sigma''_{-n}$ 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 $\sigma_n$ 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 $\sigma'''_{-n}$ such that $\omega(\sigma_n, \sigma'''_{-n}) = \omega(\sigma_n, \sigma'_{-n})$ and $\omega(s_n(\theta_n), \sigma'''_{-n}) = \omega(s_n(\theta_n), \sigma''_{-n})$. *Q.E.D.*

### C.6 Proof of Proposition 5

We start with a preliminary lemma, which is also used in the proof of Theorem 2.

Lemma C.6. *In an asynchronous second-price auction, a strategy $s_n$ is coarsely optimal if and only if, for all $\theta_n = (v_n, W_n)$ and $\sigma_{-n}$, the following hold true along $z(s_n(\theta_n), \sigma_{-n})$.*

*(i) The dates at which agent $n$ accesses the platform constitute a set which is maximal with respect to her preference ordering $W_n$ over when to submit bids.*

*(ii) The last bid submitted by agent $n$ equals her valuation $v_n$, unless she is called to play at $h \prec z(s_n(\theta_n), \sigma_{-n})$ with $\mathcal{I}(h)$ such that agent $n$ can infer that the item cannot be sold for a price below $\beta \geq v_n$. In this case, her last submitted bid is either weakly smaller than $v_n$ or strictly smaller than $\beta$.*

#### *C.6.1 Proof of the if direction*

Suppose $(X, \mathcal{I})$ is an asynchronous second-price auction satisfying the condition of the statement. By Lemma C.6, each agent has at least one coarsely optimal strategy. In particular, it is coarsely optimal to enter in a way that maximizes her preferences over when to bid and to eventually bid her valuation. Thus, the mechanism satisfies condition (i) of coarse rationality being self-supporting. It remains to show condition (ii) is satisfied. Suppose not. Then, there is some agent $n$ and $s_{-n} \in \mathscr{C}_{-n}$, such that for some $\theta \in \Theta$ and $\sigma_c \in \Sigma_c$ there exists $\sigma_n \in \Sigma_n$ 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 $s_n$ of agent $n$. Recall that, for any agent $i$ and type $\theta_i \in \Theta_i$, any coarsely optimal strategy profile $s$ must specify a behavior $s_i(\theta_i)$ that satisfies the conditions in Lemma C.6. Let $r$ be the reserve price chosen by $\sigma_c$ and let $\theta_i = (v_i, W_i)$, for all agents $i$. We consider two cases.

(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 $v_n$. By Lemma C.6, if an agent $m$ is allocated the object under $\omega(s_n(\theta_n), s_{-n}(\theta_{-n}), \sigma_c)$, agent $m$ must have bid at least $v_n$ at some point. Further, we must have $v_m \geq v_n \geq r$. But then, again by Lemma C.6, agent $n$ cannot get allocated the object at a price strictly below $v_n$ under $\omega(\sigma_n, s_{-n}(\theta_{-n}), \sigma_c)$, a contradiction. Using Lemma C.6 once again, if no one is allocated the object under $\omega(s_n(\theta_n), s_{-n}(\theta_{-n}), \sigma_c)$, then $r \geq v_n$, which once again contradicts that agent $n$ is allocated the object at a price strictly below $v_n$ 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 $p_n$, whereas $\omega(\sigma_n, s_{-n}(\theta_{-n}), \sigma_c)$ allocates agent $n$ the object at some payment $p'_n < p_n$. Observe that, by Lemma C.6, we must have $p_n \leq v_n$. If $r = p_n$ or $r > p'_n$, we have an immediate contradiction. Hence, assume $p_n > p'_n \geq r$. By the rules of the auction, no agent other than $n$ can have submitted a bid strictly above $p_n$ along $z(s(\theta), \sigma_c)$. Then, the condition satisfied by the information partition implies that, whenever an agent $i \neq n$ submitted a bid, she did not know whether a bid strictly above $p_n$ had been submitted (even if possibly agent $n$ did). Moreover, since some agent $i \neq n$ must have submitted $p_n > r$, by Lemma C.6, we must have $\max_{i \neq n} v_i = p_n$. But, again by the rules of the auction, no agent other than $n$ can have submitted a bid strictly above $p'_n$ along $z(\sigma_n, s_{-n}(\theta_{-n}), \sigma_c)$. Given the condition on the information partition, this contradicts Lemma C.6, since we know $\max_{i \neq n} v_i = p_n > p'_n$.

As desired, every coarsely optimal strategy profile is an ex-post equilibrium. *Q.E.D.*

### *C.6.2 Proof of the only if direction*

Suppose there is some $h$ at which $n$ moves with $\mathcal{I}(h)$ such that $n$ observes that a bid of at least $\beta$ has been made by some other agent, but she does not know for certain that the final selling price will be at least $\beta$. Observe that this implies agent $n$ cannot have bid weakly above $\beta$ up to $h$, and that there exists some history $h' \in \mathcal{I}(h)$ such that $r < \beta$ and $\beta$ is strictly higher than any other submitted bid. By perfect recall, it is without loss to assume $h' = h$. For all $i$, let $\beta_i$ be the last bid submitted by agent $i$ along $h$, with the convention that $\beta_i = 0$ if agent $i$ submitted no bid. Observe that, by the above, there is some agent $m \neq n$ with $\beta_i < \beta \leq \beta_m$ for all $i \neq m$.

Next, let $\theta \in \Theta$ satisfy the following. First, all agents are indifferent over all subsets of bidding dates. Second, all agents $i \neq n$ have values $v_i = \beta_i$. Third, $v_n = \beta$. The following are implications of Lemma C.6. First, there is a is coarsely optimal strategy $s_n$ that prescribes the following behavior to $\theta_n$: Play in a way that makes it possible to reach $\mathcal{I}(h)$. If $\mathcal{I}(h)$ is reached, do not raise the bid $\beta_n < \beta$ and never enter

again. If upon entering $\mathcal{I}(h)$ is revealed to be unreachable, bid $\beta = v_n$ and never enter again. Second, for all $i \notin \{n, m\}$, there is a coarsely optimal strategy such that $s_i(\theta_i)$ is coarse and consistent with reaching $h$, and such that bidder $i$ never raises her bid again after bidding $\beta_i = v_i$. Third, there is a coarsely optimal strategy $s_m$ such that $s_m(\theta_m)$ prescribes to enter only at $\max T$ and bid $\beta_m = v_m$ there.

To complete the proof, let $\sigma_c$ be consistent with reaching $h$. Since under $s_m(\theta_m)$, agent $m$ makes her first access only at $\max D$ and all bidders $i \notin \{n, m\}$ only submit bids strictly below $\beta$, there is no $h' \in \mathcal{I}(h)$ such that $h' \prec z(s(\theta), \sigma_c)$. Then, the construction of $s_n(\theta_n)$ implies that along $z(s(\theta), \sigma_c)$ agent $n$ bids $\beta$, and bidder $m$ cannot win the item at a profit at $\omega(s(\theta), \sigma_c)$. Now let $\sigma_m$ be a behavior in which agent $m$ behaves consistently with reaching $h$, and plans to never enter again. Then, $h$ is reached under $(s_{-m}(\theta_{-m}), \sigma_m, \sigma_c)$, and the construction of $s_n(\theta_n)$ is such that agent $n$ does not raise her bid $\beta_n < \beta$ and never enters again. Hence, bidder $m$ wins the item for a profit at $\omega(s_{-m}(\theta_{-m}), \sigma_m, \sigma_c)$. But this implies that part (ii) of condition for coarse rationality to be self-supporting is violated. *Q.E.D.*

### C.7 Proof of Theorem 2

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 is some agent $n$ and a history $h$ at which she is called to play such that $\mathcal{I}(h)$ is a proper subset of $\mathcal{I}^E(h)$. Moreover, at all histories, they have the same set of admissible bids. 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 $\mathcal{I}^E(h)$ is uniquely identified with a possible value $\beta$ for the currently winning bid. This implies 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, $\beta'$, is the value of the maximum between the reserve price and the highest admissible losing bid. Then, slightly abusing notation, if $\beta' \notin \mathcal{I}(h)$, we have a contradiction of Proposition 5. If $\beta' \in \mathcal{I}(h)$, then there is $\beta \in \mathcal{I}^E(h) \setminus \mathcal{I}(h)$ such that the lowest value for the currently winning bid in $\mathcal{I}(\beta)$ is strictly above $\beta'$. Once again, this contradicts Proposition 5. *Q.E.D.*

### C.8 Preliminary lemmas for Section 7

Given a mechanism that is time linear, the *depth* of a time cell $Q$ is the largest integer $\delta$ such that there are $\delta$ distinct time cells $Q_1, \ldots, Q_\delta = Q$ with $Q_\ell \preceq Q_{\ell+1}$ for all $\ell < \delta$.

Lemma C.7. *Consider a mechanism whose information partition is closed with respect to time observable. If $Q' \preceq Q''$ and $Q' \neq Q''$, then $Q'$ has strictly smaller depth than $Q''$.*

LEMMA C.8. *If $(X, \mathcal{I})$ has an information partition that is closed with respect to time observables, then it admits an elimination guide $(\mathcal{E}_k)_{k=1}^K$; that is, a partition of $\mathcal{Q}$ such that, for all $k \in \{1, \ldots, K\}$, $Q' \in \mathcal{E}_k$, $h \in Q'$, and $a \in A(h)$,*

*(i) either $(h, a) \in Q'$;*

*(ii) or $(h, a) \in \left(\bigcup_{0 \leq \ell < k} \mathcal{E}_\ell\right) \cup Z$, where $\mathcal{E}_0 := \varnothing$.*

*Moreover, all time cells $Q \in \mathcal{E}_k$ have depth $K - k + 1$.*

LEMMA C.9. *Consider an as-if dominant strategy mechanism satisfying richness. Then, only coarsely optimal strategies ex-post best reply to all others' coarsely optimal strategy profiles.*

### C.9 Proof of Theorem 3

Consider a mechanism $(X, \mathcal{I})$ satisfying the conditions of the theorem. By as-if dominance and Lemma C.9, it suffices to show that there is a sequence of play $(S(\gamma))_{\gamma \in \mathbb{N}}$ consistent with generational learning, and that for all such sequences, there is a $\gamma'$ such that for all $\gamma > \gamma'$, we have $S(\gamma) = \mathscr{C}$.

CLAIM C.3. *If the as-if dominant strategy profile is unique and a sequence of play $(S(\gamma))_{\gamma \in \mathbb{N}}$ is consistent with generational learning, then $\mathscr{C} \subseteq S(\gamma)$ for all generations $\gamma$. Moreover, the same is true if both the mechanism and the sequence of play are regular.*

PROOF. For every $n$, Theorem 1 implies all strategies in $\mathscr{C}_n$ are undominated. Further, by condition (ii) of as-if dominance and full-sampling of savvy best replies, we have $\mathscr{C} \subseteq S(1)$. We proceed by induction. Suppose $\mathscr{C} \subseteq S(\gamma)$. Condition (ii) of as-if dominance implies every $s_n \in \mathscr{C}_n$ is a best response to some $(s_{-n}(\theta_{-n}), \sigma_c) \in \Sigma_{-n}(\gamma)$ with $s_{-n} \in \mathscr{C}_{-n}$ and $\theta_{-n} \in \Theta_{-n}$. Moreover, richness implies $\Sigma(\mathscr{I}) \subseteq \Sigma(\gamma)$. This implies any $s_n'$ that is weakly better than $s_n \in \mathscr{C}_n$ for all $\theta_n \in \Theta_n$ and $\sigma_{-n} \in \Sigma_{-n}(\gamma)$ must be coarsely optimal. If the as-if dominant strategy profile is unique, we are done. Otherwise, by regularity of the mechanism and the sequence of play, we have $\omega(s_n(\theta_n), \sigma_{-n}) =_{\theta_n} \omega(s_n'(\theta_n), \sigma_{-n})$ for all $\sigma_{-n} \in \Sigma_{-n}(\gamma)$. By full sampling of recent savvy best replies, it follows that we also have $\mathscr{C} \subseteq S(\gamma + 1)$. *Q.E.D.*

Let $\underline{S}$ and $\overline{S}$ denote the limit inferior and superior of the sequence of play $(S(\gamma))_{\gamma \in \mathbb{N}}$ as $\gamma$ goes to infinity, respectively. Observe that Claim C.3 implies $\mathscr{C} \subseteq \underline{S}$. To complete the proof it remains to show that $\overline{S} \subset \mathscr{C}$. To this end, observe that, by Lemma C.8, the mechanism admits an elimination guide $\mathcal{E} = (\mathcal{E}_k)_{k=1}^K$. Let $\overline{\Sigma}$ be the set of behavior profiles $\sigma \in \Sigma$ such that there is $\theta \in \Theta$ and $s \in \overline{S}$ with $s(\theta) = \sigma$.

Slightly abusing notation, in what follows, we sometimes regard the elements of an elimination guide $\mathcal{E}$ as sets of information sets. Given a mechanism satisfying the

assumptions of Theorem 3 and an elimination guide $\mathcal{E}$, for all $k \in \{0, \ldots, K\}$, let $\mathcal{E}^k := \bigcup_{\ell \leq k} \mathcal{E}_\ell$, and for all branches $\mathbb{I}_n$, let $\Sigma(\mathbb{I}_n \cap \mathcal{E}^k)$ denote the set of behaviors $\sigma$ such that $z(\sigma)$ visits an information set in $\mathbb{I}_n \cap \mathcal{E}^k$. Note that $\Sigma_n(\mathbb{I}_n \cap \mathcal{E}^k) = \Sigma_n(\mathbb{I}_n)$ whenever $\mathbb{I}_n \cap \mathcal{E}^k$ is nonempty. Given two behaviors $\sigma'_n, \sigma''_n \in \Sigma_n(\mathbb{I}_n)$, we write $\sigma'_n \sim_{\mathbb{I}_n \cap \mathcal{E}^k} \sigma''_n$ if they are strategically equivalent conditional on reaching $\mathbb{I}_n \cap \mathcal{E}^k$, i.e., $\omega(\sigma'_n, \sigma_{-n}) =_\theta \omega(\sigma''_n, \sigma_{-n})$ for all $\sigma_{-n} \in \Sigma_{-n}(\mathbb{I}_n \cap \mathcal{E}^k)$ and $\theta$.

Definition C.1. Consider a mechanism satisfying the assumptions of Theorem 3, with elimination guide $\mathcal{E}$. Then, a product set $\tilde{\Sigma} \subseteq \Sigma$ *satisfies* $k-$*reduction* if, for all agents $n$, branches $\mathbb{I}_n$ and $\sigma_n \in \tilde{\Sigma}_n \cap \Sigma_n(\mathbb{I}_n)$, there are a type $\theta_n \in \Theta_n$ and a strategy $s_n \in \mathscr{C}_n$ such that $s_n(\theta_n) \in \Sigma_n(\mathbb{I}_n)$ and $s_n(\theta_n) \sim_{\mathbb{I}_n \cap \mathcal{E}^k} \sigma_n$.

We now use a backward-induction argument to prove that $\overline{\Sigma}$ must satisfy the $k-$reduction condition for all $k \in \{1, \ldots, K\}$. Since play within time cells is simultaneous, each agent is called to play at most once within each $Q$. Moreover, fixing $k$ and $I_n \subseteq Q \in \mathcal{E}_k$, and some $\sigma_{-n}$ compatible with reaching $I_n$, the following holds. Suppose agent $i \neq n$, after $I_n$ has been reached, for some continuation behavior of agent $n$, can be called to play in some elimination round $\mathcal{E}_\ell$ with $\ell < k$. Then, there is a unique branch $\mathbb{I}_i$ that contains all information sets in $\bigcup_{\ell<k} \mathcal{E}_\ell$ at which $i$ is first called to move after $Q \supseteq I_n$ is reached, regardless of how $n$ behaves from $I_n$ onward.

Claim C.4. *Suppose $\overline{\Sigma}$ satisfies $k-$reduction. Then, it also satisfies $(k+1)-$reduction.*

Proof. Take any $\mathbb{I}_n$ with nonempty intersection with $\mathcal{E}^{k+1}$, and fix some $\theta_n \in \Theta_n$. By sub-condition (i) of perfection, there is some $\theta'_n$ and $s_n \in \mathscr{C}_n$ such that $s_n(\theta'_n)$ is $\mathbb{I}_n-$coarsely optimal for $\theta_n$, and, by sub-condition (ii) among behaviors compatible with reaching $\mathbb{I}_n$, a best response to all $(s_{-n}(\theta_{-n}), \sigma_c)$ with $s_{-n} \in \mathscr{C}_{-n}$ and $\theta_{-n} \in \Theta_{-n}$ compatible with reaching $\mathbb{I}_n$. Next, let $I_n \in \mathbb{I}_n \cap \mathcal{E}_{k+1}$ and fix some $\sigma_{-n} \in \overline{\Sigma}_{-n}$ compatible with reaching $I_n$. Observe that there is always at least one such $\sigma_{-n}$ since all coarse behaviors are in $\overline{\Sigma}_{-n}$ by Claim C.3 and richness. Suppose that, for some continuation behavior of $n$, agent $i \neq n$ can be called to play at some information set in $\mathcal{E}^k$. As we observed before stating the claim, there exists a unique $\mathbb{I}_i$ such that, after $Q \supseteq I_n$ is reached, for any continuation behavior of $n$, the first information set lying in $\mathcal{E}^k$ at which agent $i$ is called to play is contained in $\mathbb{I}_i$. Moreover, by $k-$reduction, there are $\theta_i \in \Theta_i$ and $s_i \in \mathscr{C}_i$ such that $s_i(\theta_i) \in \Sigma_i(\mathbb{I}_i)$ and $s_i(\theta_i) \sim_{\mathbb{I}_i \cap \mathcal{E}^k} \sigma_i$. Since play within time cells is simultaneous, it follows that there are $s_{-n} \in \mathscr{C}_{-n}$ and $\theta_{-n} \in \Theta_{-n}$ such that $\omega(\sigma_n, \sigma_{-n}) =_{\theta_n} \omega(\sigma_n, s_{-n}(\theta_{-n}), \sigma_c)$ for all $\sigma_n$ compatible with reaching $I_n$ and all $\theta_n \in \Theta_n$.[28] Therefore, the above implies we can conclude that, conditional on

[28] If $i$ is never called to play again after $I_n$ is reached, by richness and the fact that play within time cells is simultaneous, we can choose $\theta_i$ and $s_i \in \mathscr{C}_i$ that generate a suitable coarse behavior.

$I_n$ being reached, behavior $s_n(\theta_n')$ is optimal for $\theta_n$ against all $\sigma_{-n} \in \overline{\Sigma}_{-n}$ compatible with reaching $I_n$. Since $I_n$ is an arbitrary element of $\mathbb{I}_n \cap \mathcal{E}^{k+1}$, this implies that $s_n(\theta_n')$ is optimal for $\theta_n$ against all $\sigma_{-n} \in \overline{\Sigma}_{-n}(\mathbb{I}_n \cap \mathcal{E}^{k+1})$.

Along with Claim C.3 and richness, the above implies that any strategy that prescribes $\theta_n$ a behavior $\sigma_n \in \Sigma_n(\mathbb{I}_n \cap \mathcal{E}^{k+1})$ that, conditional on reaching $\mathbb{I}_n \cap \mathcal{E}^{k+1}$, is not $\mathbb{I}_n-$coarsely optimal for $\theta_n$ cannot be in $\overline{S}_n$, for otherwise eventual savviness would be violated. Indeed, one may construct a dominating $\sigma_n'$ that agrees with $\sigma_n$ everywhere except if $\mathbb{I}_n \cap \mathcal{E}^{k+1}$ is reached, in which case it agrees with $s_n(\theta_n')$. By sub-condition (i) of perfection, if $\sigma_n \in \overline{\Sigma}_n$, then it must be that there is $\overline{\theta}_n \in \Theta_n$ and $\overline{s}_n \in \mathscr{C}_n$ such that $\sigma_n \sim_{\mathbb{I}_n \cap \mathcal{E}^{k+1}} \overline{s}_n(\overline{\theta}_n)$. This proves that $\overline{\Sigma}_n$ satisfies $(k+1)-$reduction. *Q.E.D.*

Observe that, since $\mathcal{E}^0 = \varnothing$, we have that $k-$reduction trivially holds for $k = 0$. Then, Claim C.4 implies that $\overline{\Sigma}$ satisfies $k-$reduction for all elimination rounds $k \in \{1, \ldots, K\}$. Next, for all agents $n$, letting $\mathbb{I}_n$ be the first branch at which $n$ is ever called to play, it follows that $\sigma_n \in \overline{\Sigma}_n$ only if $\sigma_n \sim_{\mathbb{I}_n} s_n(\theta_n)$ for some $s_n \in \mathscr{C}_n$ and $\theta_n \in \Theta_n$. Therefore, by condition (ii) of as-if dominance, eventual savviness, richness and Claim C.3, it must be that $\overline{S} \subseteq \mathscr{C}$, as desired. The above arguments along with finiteness imply that, for all sequences of play consistent with generational learning, $S(\gamma) = \mathscr{C}$. To complete the proof, we show that the set of sequences of play consistent with generational learning is non-empty. To this end, note that the elimination guide defines one particular sequence of play consistent with generational learning: Let $S(0) := S$, and $S(1)$ be the set of undominated strategy profiles. In particular, $S_n(1)$ excludes all strategies which, conditional on reaching $\mathcal{E}_1$, prescribe continuation behaviors that are not strategically equivalent to some coarsely optimal behavior. Recursively, for each $\gamma > 1$, let $S(\gamma)$ be the set of strategies surviving $\gamma$ rounds of iterated elimination of weakly dominated strategies if, in each round $\gamma \leq K$, we eliminate only strategies dominated conditional on reaching a time cell in $\mathcal{E}_\gamma$, and we do not eliminate any further strategies for $\gamma > K$. By construction, first-order rationality is satisfied. Moreover, no strategy that is undominated in the game defined by $S(\gamma - 1)$ is ever eliminated at round $\gamma$, which implies full sampling of recent savvy best replies also holds. Finally, after finitely many rounds $\gamma'$, we have $S(\gamma) = \mathscr{C}$ for all $\gamma > \gamma'$ so eventual savviness also holds. *Q.E.D.*

### C.10 Proof of Theorem 4

Let $(X, \mathcal{I})$ be a mechanism satisfying the assumptions of the theorem, with elimination guide $(\mathcal{E}_k)_{k=1}^K$. Let $(S^\ell)_\ell$ be an order of elimination of dominated strategies. Assume that either (i) the as-if dominant strategy profile is unique or (ii) the order of

elimination is regular. Arguing similarly as for Claim C.3 in the proof of Theorem 3, it must be that $\mathscr{C}$ is contained in all $S^\ell$. Indeed, by richness, all coarse behaviors survive elimination as long as $\mathscr{C}$ does: So, coarsely optimal strategies can only be eliminated by coarsely optimal strategies, which is impossible by either (i) or (ii) together with regularity of the mechanism. Arguing as for Claim C.4, $\mathscr{C}$ contains the intersection of all $S^\ell$. Indeed, let $k$ be the largest integer $k'$ such that the image $\tilde{\Sigma}$ of $\tilde{S} := \bigcap_\ell S^\ell$ satisfies $k'-$reduction, observing that this is trivially the case for $k = 0$. We will show if $k < K$, some strategy must be dominated, reaching a contradiction. Specifically, we can take a branch $\mathbb{I}_n$ such that $\mathbb{I}_n \cap \mathcal{E}^{k+1}$ is non-empty and some behavior $\sigma_n \in \tilde{\Sigma}_n \cap \Sigma_n(\mathbb{I}_n)$ for which there is no $s_n \in \mathscr{C}_n$ and $\theta_n \in \Theta_n$ such that $s_n(\theta_n) \in \Sigma_n(\mathbb{I}_n)$ and $\sigma_n \sim_{\mathbb{I}_n \cap \mathcal{E}^{k+1}} s_n(\theta_n)$. Then, using richness, perfection and $k-$reduction, this implies that a strategy prescribing $\sigma_n$ to some type is dominated by one that agrees with it everywhere except if $\mathbb{I}_n \cap \mathcal{E}^{k+1}$ is reached, in which case it is substituted with some coarsely optimal continuation. If such strategy is not in $\tilde{S}$ because it has already been eliminated, then we can construct a chain of eliminations which, by finiteness, eventually gives us a strategy in $\tilde{S}$ that does dominate $\sigma_n$. *Q.E.D.*

# Online Appendix

In this supplemental appendix, we report all proofs omitted from the main text. Then, in Online Appendix G, we study additional properties of asynchronous auctions. The proofs omitted from the main text are organized as follows. First, Online Appendix D contains the proofs for Appendix B. Second, Online Appendix E contains the proofs for the lemmas we omitted from Appendix C. Finally, Online Appendix F contains the proofs of Propositions 6, 7, and 8, as well as those of Corollaries 1 and 2.

## D Proofs for Appendix B

### D.1 Proof of Proposition B.1

Since $h$ is not preceded by a virtually terminal history, we have $h \preceq z(\sigma)$ if and only if $h \preceq \tilde{z}(\sigma)$. Then, the statement follows by definition of virtual identity. *Q.E.D.*

### D.2 Proof of Lemma B.1

As-if dominance is preserved because coarse behaviors are isomorphic and coarsely optimal strategies prescribe virtually isomorphic behaviors. Thus, coarsely optimal strategies remain coarsely optimal, and continue to ex-post best respond to each other. For reflexivity, observe that identity maps always transform a mechanism into itself. For symmetry, if a pair $(\phi, \gamma)$ transforms $\Gamma^1$ into $\Gamma^2$, then the pair $(\phi^{-1}, \gamma^{-1})$ transforms $\Gamma^2$ into $\Gamma^1$. For transitivity, if $(\phi^1, \gamma^1)$ transforms $\Gamma^1$ in $\Gamma^2$ and $(\phi^2, \gamma^2)$ transforms $\Gamma^2$ into $\Gamma^3$, then $(\phi^2 \circ \phi^1, \gamma^2 \circ \gamma^1)$ transforms $\Gamma^1$ into $\Gamma^3$. *Q.E.D.*

### D.3 Proof of Lemma B.2

By Lemma B.1, for each $\ell \in \{1, 2\}$ we have that $(X^a, \mathcal{I}^\ell)$ and $(X^b, \mathcal{J}^\ell)$ are experientially equivalent, and there is a pair $(\phi^\ell, \gamma^\ell)$ that transforms the former into the latter. To simplify notation, we avoid using $\gamma^\ell$ and simply identify the two mechanisms' coarse behavior spaces. We use $\varsigma^{(a,\ell)}$ and $\varsigma^{(b,\ell)}$, with the appropriate subscripts, to denote the virtual identity operators in $(X^a, \mathcal{I}^\ell)$ and in $(X^b, \mathcal{J}^\ell)$, respectively. Observe that experiential equivalence implies $\varsigma^{(b,\ell)}(\sigma) = \phi(\varsigma^{(a,\ell)}(\sigma))$ for all coarse behavior profiles $\sigma$. Without loss, suppose $\mathcal{I}^1$ is weakly finer than $\mathcal{I}^2$.

Let $h^b \in H^b$ be such that $\mathcal{J}^1(h^b)$ is not a singleton. This implies that some agent $n$ moves at $h^b$, and that $h^b$ is not a virtually terminal history in $(X^b, \mathcal{J}^1)$. Moreover, since there is more than one available action, there are coarse behaviors $\sigma'_n$ and $\sigma''_n$

compatible with reaching $\mathcal{J}^1(h^b)$ and there departing. Given a history $h_b \in \mathcal{J}^1(h^b)$, let $\sigma_{-n}(h^b)$ be a coarse profile for the others compatible with reaching $h_b$.

By construction, $\varsigma_{-n}^{(b,1)}(\sigma_{-n}(h^b))$ separates $\varsigma_n^{(b,1)}(\sigma_n')$ and $\varsigma_n^{(b,1)}(\sigma_n'')$. Then, by experiential equivalence, we have that $\varsigma_{-n}^{(a,1)}(\sigma_{-n}(h^b))$ separates $\varsigma_n^{(a,1)}(\sigma_n')$ and $\varsigma_n^{(a,1)}(\sigma_n'')$. Letting $h^a$ being the corresponding departure history, it follows that $h^a$ is not virtually terminal in $(X^a, \mathcal{I}^1)$. Since $\mathcal{I}^2$ is coarser, we also have that $h^a$ is not virtually terminal in $(X^a, \mathcal{I}^2)$. Furthermore, once again by experiential equivalence, for all $h_b \in \mathcal{J}^1(h^b)$ we have that $\varsigma_{-n}^{(a,1)}(\sigma_{-n}(h_b))$ is compatible with reaching $\mathcal{I}^1(h^a)$. By Proposition B.1, it follows that $\sigma_{-n}(h_b)$ is as well. Moreover, since $\mathcal{I}^2$ is coarser than $\mathcal{I}^1$, it is also compatible with reaching $\mathcal{I}^2(h^a)$, and the same is true for $\varsigma_{-n}^{(a,2)}(\sigma_{-n}(h_b))$.

Given what we have shown above, experiential equivalence implies that $h^b$ is not virtually terminal in $(X^b, \mathcal{J}^2)$. Moreover, we have that $\varsigma_{-n}^{(b,2)}(\sigma_{-n}(h_b))$ is compatible with reaching $\mathcal{J}^2(h^b)$. By Proposition B.1, the same holds for $\sigma_{-n}(h_b)$. Hence, we have that $\mathcal{J}^1(h^b) \subseteq \mathcal{J}^2(h^b)$ and we can conclude that $\mathcal{J}^1$ is weakly finer than $\mathcal{J}^2$.

To complete the proof, suppose $h^a \in H^a$ is such that $\mathcal{I}^2(h^a)$ contains some history $h_a \notin \mathcal{I}^1(h^a)$. This implies that $h^a$ is not virtually terminal in $(X^a, \mathcal{I}^2)$. Letting $\sigma_{-n}(h_a)$ be a coarse behavior compatible with reaching $h_a$, it follows that $\sigma_{-n}(h_a)$ is not compatible with reaching $\mathcal{I}^1(h^a)$. Similarly to before, we can choose coarse $\sigma_n'$ and $\sigma_n''$ that can reach $\mathcal{I}^2(h^a)$ and depart from each other there. Let $h^b$ be the departure history in $X^b$ of $\sigma_n'$ and $\sigma_n''$ when the others follow $\sigma_{-n}(h^a)$. To satisfy experiential equivalence, if $h^a$ is virtually terminal in $(X^a, \mathcal{I}^1)$, then $h^b$ is virtually terminal in $(X^b, \mathcal{J}^1)$ and we are done, for $h^b$ cannot be virtually terminal in $(X^b, \mathcal{J}^2)$. So, suppose this is not the case. Then, experiential equivalence implies that $\varsigma_{-n}^{(b,1)}(\sigma_{-n}(h_a))$ is compatible with reaching $\mathcal{J}^2(h^b)$ but not $\mathcal{J}^1(h^b)$. Using Proposition B.1 one last time, it follows that the former is a proper superset of the latter. *Q.E.D.*

### D.4 Proof of Theorem B.1

Given Lemma B.2, we only have to prove transitivity of the extended transparency ranking. Now, suppose $\Gamma^1$ is weakly less transparent than $\Gamma^2$, which in turn is weakly less transparent than $\Gamma^3$. Let $X^{12}$ be the tree used to compare the first pair, and $X^{23}$ the one used to compare the second one. Since $\Gamma^\ell$ is less transparent than $\Gamma^{\ell+1}$, it follows that its virtually terminal histories on the tree over which we represent them are a subset of those of $\Gamma^{\ell+1}$. Moreover, since $\Gamma^2$ is experientially equivalent to mechanisms defined on both trees, we can treat the coarse behavior spaces of the two mechanisms as the same. Then, we can represent $\Gamma^3$ on $X^{12}$ using the following information partition $\mathcal{I}$: Two histories $h'$ and $h''$ have $\mathcal{I}(h') = \mathcal{I}(h'')$ if and only if (i) neither is virtually terminal for the mechanism experientially equivalent to $\Gamma^2$, (ii)

they lie in an agent's coarse information set and (iii) in the mechanism defined on $X^{23}$ to which $\Gamma^3$ is experientially equivalent, the agent cannot tell apart the coarse profiles of the others compatible with reaching them when her coarse behaviors compatible with reaching them first depart from each other. In particular, from what we argued above, operating this way we obtain virtually terminal histories for $\Gamma^3$ weakly preceding those of $\Gamma^2$. Then, all mechanisms are represented on $X^{12}$, and, by Lemma B.2, the information partition corresponding to $\Gamma^3$ is finer than that corresponding to $\Gamma^2$. Hence, it is also finer than that corresponding to $\Gamma^1$, as desired. *Q.E.D.*

# E Omitted Proofs of Lemmas from Appendix C

## E.1 Proof of Lemma C.1

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 $\mathcal{I}^C$ is indeed a partition. Suppose $h \in \mathcal{I}^C(h') \cap \mathcal{I}^C(h'')$: Then, there are $\mathcal{I}' \in \mathcal{C}$ and $\mathcal{I}'' \in \mathcal{C}$ such that $\mathcal{I}'(h') = \mathcal{I}'(h)$ and $\mathcal{I}''(h'') = \mathcal{I}''(h)$. Assume without loss $\mathcal{I}''$ is coarser than $\mathcal{I}'$. This implies that $\mathcal{I}''(h') = \mathcal{I}''(h'') = \mathcal{I}''(h)$. Now let $h''' \in \mathcal{I}^C(h')$. Then, there is $\mathcal{I}''' \in \mathcal{C}$ such that $\mathcal{I}'''(h''') = \mathcal{I}'''(h')$. Next, assume that $\mathcal{I}'''$ is coarser than $\mathcal{I}''$, the other case being symmetric. Then, $\mathcal{I}'''(h''') = \mathcal{I}'''(h'')$ and thus $h''' \in \mathcal{I}^C(h'')$: This implies $\mathcal{I}^C(h') \subseteq \mathcal{I}^C(h'')$. A symmetric argument proves the other inclusion, yielding $\mathcal{I}^C(h') = \mathcal{I}^C(h'')$. Finally, we show that $\mathcal{I}^C$ satisfies perfect recall. Fix an agent $n$, and suppose she moves at two histories $h'$ and $h''$ with $\mathcal{I}^C(h') = \mathcal{I}^C(h'')$. 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 $n$ is 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 $\mathcal{I}^C(h'_k) = \mathcal{I}^C(h''_k)$, and thus $h'$ and $h''$ have the same experience under $\mathcal{I}^C$. *Q.E.D.*

## E.2 Proof of Lemma C.2

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. The construction of $\mathcal{I}$ immediately implies $P(h') = P(h'')$ and $A(h') = A(h'')$ whenever $\mathcal{I}(h') = \mathcal{I}(h'')$.

Now suppose by contradiction $\mathcal{I}$ does not have perfect recall. Then, there exists a terminal history $z'$ such that, for some $h' \prec z'$ at which some agent $n$ moves,

there is $h'' \in \mathcal{I}(h')$ such that her 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 $\{I^k\}_{k=1}^K$ at which $n$ moves such that $I^k \in \mathcal{I}^1$ for all odd $k$ and $I^k \in \mathcal{I}^2$ 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 $\xi(h^1)$ be the last history agent $n$ is called to play before $h^1$. (If she was never called to play before, let $\xi(h^1)$ 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—by perfect recall, we must have $\mathcal{I}^1(\xi(h^2)) = \mathcal{I}^1(\xi(h^1))$. This implies $\mathcal{I}(\xi(h^2)) = \mathcal{I}(\xi(h^1))$. Moreover, due to our selection of $h'$, agent $n$'s experience under $\mathcal{I}$ at $\xi(h^2)$ and $\xi(h^1)$ must thus be the same. But then, since $\xi(h^1)$ and $\xi(h^2)$ lie in the same information set and we argued that agent $n$ takes the same action at both prior reaching $I^1$, it follows that her experience at $h^1$ and $h^2$ is the same as well. Proceeding this way using an induction argument, we conclude agent $n$'s experience at $h'$ is the same as at $h''$, reaching a contradiction. *Q.E.D.*

### E.3 Proof of Lemma C.4

The first claim is a consequence of the fact that $\mathscr{I}$ only depends on the mechanism's tree $X$. To prove the second claim, observe the following. By Proposition 2, for all $h' \in \mathcal{I}(h)$, the set of coarse behavior profiles $\sigma_{-n}$ compatible with reaching $h'$ is nonempty. Now, consider two distinct $h', h'' \in \mathcal{I}(h)$, and let $h'''$ the last history preceding both. By perfect recall, $n$ is not called to play at $h'''$. Letting $a'$ and $a''$ be such that $(h''', a') \preceq h'$ and $(h''', a'') \preceq h''$, all coarse behavior profiles compatible with reaching $h'$ play $a'$ at $h'''$, while those compatible with reaching $h''$ play $a''$. Hence, distinct histories in $\mathcal{I}(h)$ identify disjoint sets of coarse behavior profiles for the others, and the claim follows. *Q.E.D.*

### E.4 Proof of Lemma C.5

Observe 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 $n$, there is $\theta \in \Theta$ and an $\mathcal{I}'$–measurable $\sigma_n'$ 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 $\sigma''_n$ such that $\sigma''_n(h) = \sigma'_n(h)$ at all $h$ at which $n$ is called to play under $(\sigma'_n, s_{-n}(\theta_{-n}))$. Then, $\omega(\sigma''_n, s_{-n}(\theta_{-n})) = \omega(\sigma'_n, s_{-n}(\theta_{-n}))$, contradicting that $s$ is an equilibrium of $(X, \mathcal{I}'')$. *Q.E.D.*

### E.5 Proof of Lemma C.6

From the perspective of agent $n$, each of the other agents' coarse behaviors is equivalent to choosing a set of bidding dates and 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 choose a maximally preferred set of bidding dates and eventually bid truthfully, unless at some point it is revealed to her that she cannot make a strictly positive payoff, in which case she can bid anything that ensures she does not win at a loss. The two conditions ensure agent $n$ behaves consistently with these incentives. Observe that, since agents who access on the same date move simultaneously, an agent can never exclude that the others will raise their standing bids. *Q.E.D.*

### E.6 Proof of Lemma C.7

Suppose $Q' \preceq Q''$ and $Q' \neq Q''$. Let $d$ be the depth of $Q'$, so there exist distinct time cells $Q_1, \ldots, Q_d = Q'$ with $Q_\ell \preceq Q_{\ell+1}$ for all $\ell < d$. Appending $Q''$ yields a strictly longer chain $Q_1, \ldots, Q_d = Q' \preceq Q''$. By acyclicity of $\preceq$, all time cells in this chain are distinct, so $Q''$ has depth at least $d + 1$. *Q.E.D.*

### E.7 Proof of Lemma C.8

Recall that if $Q'$ precedes $Q''$ and $Q' \neq Q''$, then, by acyclicity, $Q'$ has strictly smaller depth than $Q''$. Now, let $K$ be the maximum depth reached by some time cell. Let $\mathcal{E}_K \subseteq \mathcal{Q}$ be the set of all $Q \in \mathcal{Q}$ such that $Q' \preceq Q$ implies $Q' = Q$. Observe that $\mathcal{E}_K$ contains only the time cell containing the root of the game. Recursively, for all $k < K$ such that $\mathcal{Q} \neq \bigcup_{k<\ell} \mathcal{E}_\ell$, let $\mathcal{E}_k$ be the set of all $Q \in \mathcal{Q} \setminus \bigcup_{k<\ell} \mathcal{E}_\ell$ such that $Q' \preceq Q$ implies either $Q' = Q$ or $Q' \in \mathcal{E}_\ell$ for some $\ell > k$. This construction implies that $Q \in \mathcal{E}_k$, if and only if its depth is $K - k + 1$. Indeed, using the fact that play within a time cell is simultaneous, it is immediate to see that this is the case if $k = K$. Thus, suppose the claim holds for all $\ell$ with $k < \ell \leq K$. Using this inductive hypothesis, we show the claim then also holds for $k$. To prove the if direction, observe that, if $Q \in \mathcal{E}_k$, it must be that all time cells preceding $Q$ are in $\bigcup_{k<\ell} \mathcal{E}_\ell$. This implies the depth of

$Q$ is no larger than $K-k+1$. But since $Q \in \mathcal{E}_k$ requires $Q \notin \bigcup_{k<\ell} \mathcal{E}_\ell$, the inductive hypothesis implies that its depth is exactly $K-k+1$. To prove the only if direction, suppose for the sake of contradiction that $Q$ has depth $K-k+1$ but $Q \notin \mathcal{E}_k$. By the inductive hypothesis, $Q \notin \bigcup_{k<\ell} \mathcal{E}_\ell$. Thus, there must be some $Q' \neq Q$ such that $Q' \preceq Q$ but $Q' \notin \bigcup_{k<\ell} \mathcal{E}_\ell$. Moreover, once more by the inductive hypothesis $Q'$ has depth at least $K-k+1$. But, using Lemma C.7, this contradicts the fact that $Q$ is of depth $K-k+1$. To complete the proof, observe that the construction of $(\mathcal{E}_k)_k$ together with the fact that play within time cells is simultaneous imply that requirements (i) and (ii) of the statement hold for all $k$, $Q' \in \mathcal{E}_k$, $h \in Q'$ and $a \in A(h)$. *Q.E.D.*

### E.8 Proof of Lemma C.9

By richness, every coarse profile $\sigma_{-n}$ can be played under coarse rationality for some profile $\theta_{-n}$. Hence, any strategy that is ex-post optimal against all coarsely optimal strategy profiles of the others must also best respond to all coarse behavior profiles of the others. *Q.E.D.*

## F Other Proofs Omitted from the Main Text

### F.1 Proof of Proposition 6

We prove the two directions separately. Recall that the auction are relatively small, and thus we have $\overline{v} + |N| - 2 \leq c < (|N|-1)\overline{v}$.

#### *F.1.1 Proof of the only if direction*

Suppose the only if direction is false. Then, there is an as-if dominant strategy asynchronous VCG auction in which we can take an history $h'$ satisfying the following.

(i) Some agent $n$ moves at $h'$, and her standing bid is $b_n < \overline{v}$.

(ii) Letting $b_i$ denote the standing bid of agent $i$ along $h'$, where $b_i = 0$ if $i$ submitted no bid, we have that $\beta := \sum_{i \neq n} b_i > c - \overline{v} + 1$.

(iii) There is $j \neq n$ with $1 < b_j = \max_{i \neq n} b_i$ and $\sum_{i \neq j} b_i < c$.

(iv) For all histories $h \in \mathcal{I}(h')$, the sum of the bids of the others is at least $\beta$.

Suppose first $\beta + b_n > c$. Then, it is coarsely optimal for agent $n$ to never raise her bid again once $h'$ is reached. In fact, by (iv) she knows that $b_n$ suffices for the project to be undertaken with certainty. In particular, consider $\theta \in \Theta$ such that all agents are

indifferent between all sets of access dates, $v_i$ is equal to each agent's standing bid at $h'$ for each $i \neq n$ (or zero, if she submitted none), and $v_n = \overline{v}$ which by assumption strictly exceeds $b_n$. Then, we can construct the following coarsely optimal profile. All agents other than $n$ behave coarsely, and such that $h'$ can be reached. Agent $n$ behaves so that $h'$ can be reached, and such that she bids $\overline{v}$ if $\mathcal{I}(h')$ is not reached, and leaves the auction bidding $b_n$ otherwise. Then, agent $j$ would find it strictly profitable to lower her last standing bid by one.

Second, suppose $\beta + b_n \leq c$. Then, (ii) implies $b_n \leq c - \beta < \overline{v} - 1$. Hence, defining $b'_n = c - \beta + 1$, we have $b'_n \in V$ and $b'_n < \overline{v}$. Next, consider $\theta \in \Theta$ such that all agents are indifferent between all sets of access dates, $v_i$ is equal to each agent's standing bid at $h'$ for each $i \neq n$ (or zero, if she submitted none), and $v_n = \overline{v}$ which as we argued strictly exceeds $b'_n$. Then, we can construct the following coarsely optimal profile. All agents other than $n$ behave coarsely, and such that $h'$ can be reached. Agent $n$ behaves so that $h'$ can be reached, and such that she bids $\overline{v}$ if $\mathcal{I}(h')$ is not reached, and leaves the auction bidding $b'_n$ otherwise. But this implies $j$ has a profitable deviation to bid one unit less. In fact, $\beta + b'_n = c + 1$ and $b_j \geq 2$ by (iii) and agents being relatively small. This implies bidder $j$ pays a positive amount on the path we constructed, which can be reduced by reducing her bid and have bidder $n$ bid $\overline{v}$ instead. *Q.E.D.*

#### *F.1.2 Proof of the if direction*

If the others behave coarsely, a strategy is always optimal if and only if the agent always bids on a set of date that maximizes her preferences over when to bid, and eventually bid her valuation, unless she is revealed that a lower bid guarantees the project to be undertaken with probability one. However, the information partition is such that this only happens when the project is undertaken for sure, and the standing bid of agent $n$ is already sufficient for no other bidder to pay any price. In particular, the assumption that $c < (|N| - 1)\overline{v}$ implies that an agent can never exclude that the final bids of the others will end up to be exactly $c - v_n + 1$. In that case, the price that agent $n$ has to pay is never less than $v_n - 1$, and $v_n$ is the lowest bid guaranteeing the project is undertaken. In particular, if she instead reports $v_n - 1$, chance randomizes among implementing and not implementing the project. *Q.E.D.*

### F.2 Impossibility of obviously dominant multiunit VCG auctions

In this subsection, we provide a direct proof of the impossibility of implementing VCG outcome in obviously dominant strategies when there are multiple units. As mentioned in the main text, the impossibility also follows from Proposition 7.3 in Gkatzelis et al. (2017), which shows no *weakly group-strategyproof* mechanism can

guarantee more than a $\sqrt{2}-$approximation to optimal social welfare when there are at least two agents with non-increasing marginal values over at least two units, since obviously dominant-strategy mechanisms are weakly group-strategyproof (cf. Li, 2017).

We say that a multiunit allocation problem is *proper* if there are at least two agents, two units, and two marginal values per unit.

PROPOSITION F.1. *For proper multi-unit allocation problems with transfers, there is no mechanism that implements the Vickrey outcome in obviously dominant strategies.*

PROOF. Without loss, assume the marginal values lie in some set $\{0, 1, ..., u\}$ with $u \geq 1$. Suppose $(X, \mathcal{I})$ implements the VCG outcome in obviously dominant strategies. Let $s$ denote the obviously dominant strategy profile. Again without loss, we can assume chance plays all its behaviors with positive probability and that $\mathcal{I}$ is a perfect information partition. Take a profile $\theta' \in \Theta$ such that agents 1 and 2 have a marginal value of one for the first unit and zero for all others; all other agents have zero marginal value for all units, Let $z'$ be the terminal history 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 VCG 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 $\theta''$ 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. *Q.E.D.*

### F.3 Proof of Proposition 7

For each agent $n$, let $s_n$ denote the strategy corresponding to sincere bidding. That is, for each type $\theta_n$, let $s_n(\theta_n)$ denote the behavior that, whenever called to play, prescribes to drop the $\ell$-th unit in the $r$-th round if and only if $\theta_n^\ell \leq p^r$. First, observe that $s_n$ is the unique coarsely optimal for each agent $n$. This follows from the fact that sincere bidding is the unique dominant strategy in the Ausubel auction without dropout information (see Theorem 1 in Ausubel, 2004), which corresponds to the coarsest partition that can be defined on the auction's tree. In fact, as long as the others behave coarsely, over-reporting only exposes to a loss, which is sometimes strict (units may

be rationed in the next round and the agent may have highest priority). Similarly, under-reporting can only leave profit on the table, sometimes strictly (as the others may drop their units, and the agent may not win some units at a profitable price due to having the lowest priority). This is true even after deviating from $s_n$, as shown by Okamoto (2018). Second, the fact that sincere bidding is an ex-post equilibrium of the Ausubel auction with public dropouts then immediately implies that coarse rationality is self-supporting in that auction, completing the proof. *Q.E.D.*

### F.4 Proof of Proposition 8

For each agent $n$, let $s_n$ denote the strategy corresponding to truthful reporting. That is, for each type $\theta_n$, let $s_n(\theta_n)$ denote the behavior that, whenever called to play, applies to the option among those not yet applied to.

First, note $s_n$ is coarsely optimal for each agent $n$. If the others behave coarsely, their behavior is equivalent to submitting a rank order list over alternatives. Since truthful reporting is a dominant strategy in the static DA (Dubins and Freedman, 1981; Roth, 1982) and TTC (Shapley and Scarf, 1974), it follows that $s_n$ is coarsely optimal in the sequential DA and TTC. Moreover, $s$ only prescribes coarse behaviors, so the self-supporting condition follows, and the mechanism is as-if dominant strategy.

For uniqueness, suppose at some information set, $\theta_n$ applies to some $y'$ which she ranks below another available option $y''$. In sequential DA, it is possible $n$ has first priority at both schools (or one of them is the safe outside option), and hence $n$ can be strictly worse off by applying to $y'$ over $y''$.

For sequential TTC, since agent $n$'s applications to already-assigned houses can never be observed by the other agents, $s_n(\theta_n)$ is strategically equivalent to any behavior $\sigma_n$ follows $\theta_n$'s ranking, up to dropping or inserting any house which is revealed to have been already permanently assigned. Hence, (conditional) coarse optimality of any candidate strategy at $\mathbb{I}_n$ depends only on the order in which it applies to non-permanently-assigned houses. Hence, suppose $y'$ and $y''$ have not yet been permanently assigned in a previous round. Let $y_n$ denote agent $n$'s house. If $y' = y_n$, then $n$ will get matched to $y_n$ for certain if she applies to it. Since $y''$ was not removed as part of a cycle in the previous round, the chain of applications emanating from $y''$ must end with some $m$ who is called to play again simultaneously with $n$. Thus, $n$ can get matched to $y''$ if she applies to it—either because $m = n$, or because $m$, who is asked to make a new application, applies to $y_n$ next. This is possible since $y_n$ has not yet been removed, so $m$ cannot have applied to $y_n$ yet, and it may be her next most preferred house. In either case, $n$ is strictly worse off from applying to $y_n$ over $y''$. A symmetric argument applies if $y_n = y''$. Thus, suppose $y_n \notin \{y', y''\}$. Let $m'$

and $m''$ denote the last agents in the chain of applications emanating from $y'$ and $y''$, respectively. If either $m' \neq n$ or $m'' \neq n$, by the fact $y_n$ has not yet been removed in a cycle, $m'$ or $m''$, respectively, cannot have applied to $y_n$ yet, they both are called to play with no held application, and thus can both point to $y_n$ next. Hence, regardless of the identity of $m'$ and $m''$, agent $n$ can be assigned whichever of $y'$ and $y''$ she applies to, and hence is strictly worse off by applying to $y'$ over $y''$. *Q.E.D.*

### F.5 Proofs of Corollary 1 and Corollary 2

It is easy to see that all static dominant strategy mechanisms considered in the paper satisfy linearity, richness and perfection, and they also have a unique (as-if) dominant strategy profile. To complete the proof, we verify the same is true for the dynamic mechanisms considered.

Lemma F.1. *The Ausubel Auction, with either private or public dropouts, satisfies time-linearity, richness and perfection.*

Proof. Time linearity follows by the round structure of the auction: Two histories belong to the same time cell if and only if they correspond to the same unit price. Richness follows because the unique as-if dominant strategy $s_n$ prescribes to drop a unit if and only if the unit price is strictly above an agent's marginal value for it. Hence, each coarse behavior—which are lists specifying how many units an agent should drop in each round—corresponds to some type in $\theta_n \in \Theta_n$. In particular, there is no difference between assuming that others behave coarsely or follow $s_{-n}$. This implies the second requirement of perfection. Then, perfection follows because, by Okamoto (2018) and arguing similarly as we did when proving uniqueness in Proposition 7, the tie breaking rule makes it such that, even upon reaching a branch $\mathbb{I}_n$ after deviating from $s_n(\theta_n)$, it is uniquely optimal for $\theta_n$ to behave according to $s_n(\theta_n')$ the type $\theta_n'$ that should have behaved as $\theta_n$ did in the previous rounds, and shares the same valuations as $\theta_n$ for the units she is still demanding. *Q.E.D.*

Lemma F.2. *Both the sequential DA and TTC algorithms satisfy time-linearity, richness and perfection.*

Proof. Time-linearity follows by the round structure of the mechanisms. Two histories are part of the same time cell if and only if they are part of the same round of applications. Observe that this would not be true if agents could not observe which round they are in when called to play. Note that in the sequential TTC, an agent $n$ may make multiple applications within a single round if she first applies to already-assigned houses. However, this is not observed by opponents, and hence her

opponents' information sets will be determined solely by the applications she makes to a not-yet-assigned house. To prove richness, observe that coarse behaviors are equivalent to submitting applications following some given list. Hence, richness follows because each list is the preference order of some type, and the coarsely optimal behavior specified by the as-if dominant strategy $s_n$ is always equivalent to reporting truthfully. Finally, we prove perfection. Observe that the second requirement holds as $s$ always specifies coarse behavior profiles. To prove the first requirement, suppose $\theta_n$ has deviated from $s_n(\theta_n)$ and is now reached a branch $\mathbb{I}_n$. There is a type $\theta'_n$ for which what $\theta_n$ did so far is consistent with the as-if dominant strategy, and agrees with $\theta_n$ on all options $\theta_n$ has not applied yet. Since, by Proposition 8, $s_n(\theta'_n)$ is uniquely coarsely optimal for $\theta'_n$—up to strategic equivalence—the same is true for $\theta_n$ conditional on reaching $\mathbb{I}_n$. *Q.E.D.*

Lemma F.3. *The sequential Online English Auction satisfies time-linearity, richness and perfection. Moreover, it is regular.*

Proof. Time linearity is obtained by grouping together in the same time cell all histories at which agents move that correspond to the same access date. For richness, we can use the fact that the set $\mathcal{W}$ of preferences contains full indifference, and the fact that $\Theta_n$ is the cartesian product of the valuation set $V$ and $\mathcal{W}$. By Lemma C.6 and the fact that bids lie in $V$, this implies that we can use coarsely optimal strategies to generate all coarse behaviors. For perfection, suppose that $\theta_n = (v_n, W_n)$ reaches an branch after behaving in a way that is not consistent with coarse rationality. Arguing similarly as in Lemma C.6, continuation is $\mathbb{I}_n-$coarsely optimal if and only if it satisfies the following conditions: Successive access dates to maximize $W_n$ conditional on the entries already realized before reaching $\mathbb{I}_n$. Moreover, if agent $n$'s standing bid upon reaching $\mathbb{I}_n$ is weakly below $v_n$, she eventually bids $v_n$, unless she is revealed that the final price will be above weakly above $v_n$ , in which case she can have any last bid weakly below $v_n$. Otherwise, she never raises her bid again. As long as the others follow coarsely optimal strategies (or coarse behaviors), this is an optimal continuation, so the second requirement of perfection holds. To prove the first requirement holds, we show that there is some $\theta'_n$ for which these are continuations of its as-if dominant strategy. In particular, they are all as-if dominant strategies of the type $\theta'_n$, who has no preferences over time and whose valuation $v'_n$ is the maximum between $v_n$ and the standing bid of agent $n$ upon reaching $\mathbb{I}_n$.

Finally, for regularity, observe that it is a dominated strategy to leave the auction with a standing bid lower than an agent's true valuation—unless she is revealed she cannot win at a profit, in which she can bid anything below her value or the price threshold she observed. (However, it is not dominated to bid on a sub-optimal set of

access dates.) Since agents care only about others' final standing bids, and not about when others enter, once dominated strategies are eliminated, all coarsely optimal strategies perform the same. *Q.E.D.*

Given the above lemmas, the corollaries follow from Theorem 3 and Theorem 4, respectively. *Q.E.D.*

## G Additional Results for Section 5

### G.1 As-if dominant strategy outcomes of asynchronous second-price auctions

All as-if dominant strategy asynchronous second-price auctions have the same outcomes as the dominant strategy second price auction in which nothing but possibly the reserve price is disclosed when an agent accesses the platform. This means that not only they can achieve efficiency, but also—in regular and symmetric independent-private-value environments (see, e.g. Elkind, 2007)—weakly higher expected revenues than any other Bayesian incentive compatible mechanism.

Proposition G.1. *All as-if dominant strategy profiles of an asynchronous second-price auction 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.*

Proof. By Lemma C.6, 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 is revealed that either the reserve value or a losing bid of another agent was weakly above her value. All claims of the statement follow. *Q.E.D.*

### G.2 Shill-proofness of the Online English Auction

We say a truthful equilibrium is *fully asynchronous* if each bidder bids her true value on the first date compatible with maximizing her preferences over when to bid, and not all bidders enter always at the same time.[29] To ensure such equilibria exist, for

[29] An equilibrium is *truthful* if each bidder bids her true value upon entering, unless she knows another bidder already bid weakly more, in which case she may bid in any way that guarantees a loss.

this subsection only, we assume that all $W_n \in \mathcal{W}$ are *monotone*: If $\delta', \delta'' \in \mathcal{D}$ are such that $\delta'' \subseteq \delta'$, then $W_n$ ranks $\delta'$ weakly below $\delta''$.

A fully asynchronous truthful equilibrium is *shill-proof* if a shill bidder—chosen at random among all agents and whose objective is to maximize auction revenues, who is the only one aware of this objective—can never raise expected revenues by entering and submitting a positive bid. We assume valuations and preferences over time are chosen independently from each other. In particular, an agent valuation is drawn from a probability distribution $\pi \in \Delta(V)$. The distribution is *regular* if it has full support and, letting $V$ be ordered such that $v_k \leq v_\ell$ if and only if $k \leq \ell$, the map

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

is increasing. A reserve price is *optimal* if it maximizes expected revenue in a sealed-bid second-price auction with types drawn independently from $\pi$.[30]

Proposition G.2. *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.*

Proof. For all $v_k$, define $f(v_k) := \pi(\{v_k\})$ and $F(v_k) := \pi(\{v_\ell \mid k \leq \ell\})$ and recall that all submitted bids lie in $V$. The assumption on the prior and full asynchronicity of the equilibrium implies the date at which an agent submit her bid reveals no information about her valuation.

First, we prove statement (i). Consider the Online English Auction with optimal reserve price $r^*$. Since truthful bidding is an as-if dominant strategy, it is an ex-post equilibrium for all bidders to bid truthfully. It remains to show that a shill bidder has no incentive to submit any bid $b > 0$. Given the information partition and Proposition 5, submitting a bid at any date at which the highest bid is known to be at least $\beta \geq r^*$ is equivalent to setting a new reserve price among all bidders with values above $\beta$. Since bidding times are uninformative of types, the optimal bid solves

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

[30]Specifically, the optimal reserve is is the lowest $r \in V$ such that $v_k - (v_{k+1} - v_k)\frac{1-\pi(\{v_l|\ell \leq k\})}{\pi(\{v_k\})} > 0$ for all $v_k \in V$ with $v_k \geq r$ (see, e.g., Akbarpour and Li, 2020; Elkind, 2007).

By regularity, for all $b_k \geq \beta$,

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

whenever $\beta \geq r^*$. Hence submitting $b > \beta$ cannot be profitable in expectation, and submitting $b = \beta$ does not affect the outcome since at least one other bidder must have already submitted a bid of $\beta$. Thus the equilibrium is shill-proof.

Second, we prove statement (ii). Take a history of play along which a bidder receives more information than she would in an Online English auction. In particular, suppose the shill bidder observes that the highest bid is above $\beta$ but cannot rule out that the final price will be below $\beta$, then submitting a bid of $\beta$ strictly increases expected revenues, so the equilibrium is not shill-proof. *Q.E.D.*