



Menu mechanisms [☆]

Andrew Mackenzie ^{a,*}, Yu Zhou ^b

^a Department of Microeconomics and Public Economics, Maastricht University, Maastricht, the Netherlands

^b Graduate School of Economics, Kyoto University, Kyoto, Japan

Received 14 March 2021; final version received 20 June 2022; accepted 25 June 2022

Available online 3 July 2022

Abstract

We investigate *menu mechanisms*: dynamic mechanisms where at each history, an agent selects from a menu of his possible assignments. We consider both ex-post implementation and full implementation for a strengthening of dominance that covers off-path histories, and provide conditions under which menu mechanisms provide these implementations of rules. Our results cover a variety of environments, including matching with contracts, labor markets, auctions, school choice, marriage, object allocation, and elections. © 2022 The Author(s). Published by Elsevier Inc. This is an open access article under the CC BY license (<http://creativecommons.org/licenses/by/4.0/>).

JEL classification: D47; D82; C78

Keywords: Menu mechanism; Strategy-proofness; Robust implementation

[☆] Andrew Mackenzie gratefully acknowledges financial support from the Dutch Research Council (NWO), grant number VI.Veni.201E.004. Yu Zhou gratefully acknowledges financial support from the Joint Usage/Research Center at ISER of Osaka University and from the Grant-in-aid for Research Activity of the Japan Society for the Promotion of Science (19K13653, 20H05631, 20KK0027). We thank Masaki Aoyagi, Inácio Bó, Albin Erlanson, Makoto Hagiwara, Scott Kominers, Takashi Kunimoto, Jordi Massó, Nozomu Muto, Marek Pycia, Al Roth, Jim Schummer, Alex Teytelboym, William Thomson, and Rakesh Vohra; seminar participants at Shinshu University, the Centre for Mathematical Social Sciences, Bielefeld University, University of Rochester, Yokohama National University, Sorbonne Economics Center, Georgetown University, Northwestern University, North Carolina State University, the 2020 Lausanne Matching and Market Design Workshop, the 2020 Conference on Mechanism and Institution Design, the 2021 China Meeting of the Econometric Society, the 2021 World Congress of the Game Theory Society (GAMES), the 2021 Conference of the European Association for Research in Industrial Economics (EARIE), and the 2021 Workshop on Aggregation Across Disciplines; and the editor Marzena Rostek, the associate editor, and two anonymous referees.

* Corresponding author.

E-mail addresses: a.mackenzie@maastrichtuniversity.nl (A. Mackenzie), zhouyu_0105@hotmail.com (Y. Zhou).

<https://doi.org/10.1016/j.jet.2022.105511>

0022-0531/© 2022 The Author(s). Published by Elsevier Inc. This is an open access article under the CC BY license (<http://creativecommons.org/licenses/by/4.0/>).

1. Introduction

1.1. Overview

Among the many mechanisms that have been designed by economists, dominant strategy implementations have a striking record for real-world application—from auctions to school choice procedures to labor market clearinghouses and more. For every problem that can be solved with such a mechanism, the prototypical example is the *direct mechanism*, where the agents simultaneously report all private information to a central administrator. In practice, however, there have been some prominent concerns about miscalculations by these administrators (both deliberate and accidental),¹ and there are several advantages to using dynamic mechanisms instead: for example, they can reduce the volume of bits that must be communicated for implementation (Segal, 2010), make it easier for agents who have trouble with contingent reasoning to identify dominant strategies (Li, 2017), and improve credibility and transparency when the administrator might deliberately miscalculate outcomes such as by shill bidding in auctions (Akbarpour and Li, 2020; Hakimov and Raghavan, 2020). Though these advantages generally come at the cost of dominant strategy implementation—after all, a dynamic mechanism generally has many more strategies than its associated direct mechanism—the contribution of this paper is to identify a broad class of dynamic mechanisms across a variety of environments whose implementations are robustly incentive compatible in an even *stronger* sense.

In particular, for finite environments with private values and no consumption externalities, we investigate *menu mechanisms*: dynamic mechanisms where at each history, an agent selects from a menu of his possible assignments. Every rule has menu mechanisms that imitate the direct mechanism. Moreover, many prominent rules are effectively described with a menu mechanism—in particular, with an algorithm for calculating outcomes where agents behave desirably in a menu mechanism. Familiar menu mechanisms, and familiar rules with algorithms easily associated with menu mechanisms, include

- the cumulative offers process (Hatfield and Milgrom, 2005) in matching with contracts;
- the salary adjustment process (Kelso and Crawford, 1982) in labor markets;
- the Crawford-Knoer auction (Crawford and Knoer, 1981; Demange et al., 1986) for auctions with unit demand, including the English auction for one-object auctions;
- student-proposing deferred acceptance (Gale and Shapley, 1962) in school choice, including male- and female-proposing deferred acceptance in marriage environments;
- Gale’s top trading cycles (reported in Shapley and Scarf, 1974), serial dictatorship (see, for example, Svensson, 1999), the broader class of hierarchical exchange rules (Pápai, 2000), and the even broader class of trading cycles rules (Pycia and Ünver, 2017; Bade, 2020) in object allocation; and
- direct menu mechanisms in a variety of environments, including for voting by committees (Barberà et al., 1991) in two-candidate elections.

¹ For example, a prominent lawsuit alleged that the central clearinghouse for the resident labor market in the United States had the purpose and effect of allowing hospitals to collude to suppress wages (Jung v. Association of American Medical Colleges, 2002), and the City of Boston unintentionally miscalculated school choice outcomes due to a coding error (Dur et al., 2018).

We emphasize that for many of these examples, while the algorithm is not novel, and while the direct mechanism that gathers all private information and then calculates outcomes by simulating desired behavior in the menu mechanism is not novel, the incentive compatibility of the menu mechanism itself has not been considered previously. Our main results imply that *all* of the menu mechanisms given above—or naturally derived from the algorithm of a rule given above—are robustly incentive compatible; see Section 5.

To capture robust incentives in a dynamic mechanism, we use the solution concept of *everywhere-dominance*: a strengthening of dominance that covers off-path histories. This solution concept, which to our knowledge is novel, requires that from each history—even if the history shares an information set with other histories—the continuation strategies form a dominant strategy equilibrium. This is a very demanding solution concept, requiring each agent's strategy to *always* be a best response—regardless of his beliefs about how his peers have played thus far, how his peers will play henceforth, and any other information he might lack about his current history—even if he has already deviated from this strategy (for example, due to error).

Formally, we provide sufficient conditions for a menu mechanism to provide either (i) an ex-post everywhere-dominant implementation, or (ii) both an ex-post everywhere-dominant implementation and a full everywhere-dominant implementation. Both implementation notions require that each type profile has an everywhere-dominant strategy equilibrium that achieves the desired outcome; loosely, ex-post moreover requires that each agent's strategy only depends on his own private information, while full moreover requires that all equilibria achieve the desired outcome.

In order to guarantee these robust everywhere-dominant implementations, we use the following conditions:

- For the environment, *richness* requires that each agent might have any strict ranking of his assignments (though he may also have other rankings), and *strictness* requires that agents are never indifferent.
- For the rule, *strategy-proofness* has the usual definition, and *group strategy-proofness* is the usual strong version requiring that no coalition of agents can jointly misreport to make one member better off without making another member worse off.
- For the menu mechanism, *non-repeating* requires that each agent can either never select a previous choice or never select a previous rejection, and *reaction-proofness* loosely requires that whenever one agent can deduce something about another agent's choices, the latter's assignment has already been determined.
- For the type-strategy profile—which we refer to as the *convention—preferential* requires that each agent always selects a most-preferred assignment and breaks ties consistently, and *compatibility with the rule* requires that the desired outcome is always achieved if all agents conform.

We discuss these conditions in the context of our illustrative example in Section 1.2. Our main results state that (i) if all of our conditions are satisfied except possibly *strictness* and *group strategy-proofness*, then we have an ex-post everywhere-dominant implementation, and (ii) if all of our conditions are satisfied, then we have both an ex-post and full everywhere-dominant implementation:

Theorem 1. *For each rich environment, each strategy-proof rule, each non-repeating and reaction-proof menu mechanism, and each preferential convention that is compatible with the*

rule, the menu mechanism is an ex-post everywhere-dominant implementation of the rule via the convention.

Theorem 2. *For each rich and strict environment, each group strategy-proof rule, each non-repeating and reaction-proof menu mechanism, and each preferential convention that is compatible with the rule, the menu mechanism is both an ex-post everywhere-dominant implementation of the rule via the convention and a full everywhere-dominant implementation of the rule.*

We remark that if either theorem applies to a given menu mechanism, then it also applies to any associated menu mechanism that simply conceals some information from the agents. For example, if either theorem applies to a menu mechanism that satisfies perfect recall, then it also applies to the associated menu mechanism where no actions are observable. Indeed, there is a tension between everywhere-dominance and *obvious* dominance (Li, 2017), in the sense that concealing information facilitates the former while revealing information facilitates the latter; see Section 2.5 for details.


As corollaries, our theorems provide novel results for matching with contracts, labor markets, auctions, school choice, marriage, object allocation, and elections; see Section 5.² Taken together, our results show that like direct mechanisms, menu mechanisms can systematically provide robust dominant strategy implementations, but unlike direct mechanisms, they may be able to realize the various advantages of dynamic mechanisms—such as privacy, simplicity, credibility, and transparency. Indeed, we formalize the observation that menu mechanisms improve upon the privacy of direct mechanisms (Appendix B), while recent experiments provide evidence that menu mechanisms are simpler for participants than direct mechanisms (see Section 1.3).

Before proceeding, we caution that even though switching from direct mechanisms to menu mechanisms can realize the above benefits at no cost in terms of incentives, that does not mean that there are no costs at all. Perhaps most importantly, menu mechanisms generally process messages serially, and may therefore lead to bottlenecks of the kind that were documented in the decentralized telephone market for clinical psychologists (Roth and Xing, 1997). Whether or not the advantages of dynamic mechanisms are more desirable than the speed of direct mechanisms ultimately must depend on context; our results simply show that dominant strategy implementations are widely available for both choices.


1.2. Illustrative example

In this section, we illustrate our model and main results in the context of serial dictatorship for object allocation environments. In particular, suppose that there are at least as many (indivisible) objects as agents, suppose that each agent must consume one object, and suppose that monetary transfers are not possible. Suppose moreover that the environment is both *rich* and *strict*: each agent's private information type could specify any strict preference ranking over the objects, but could not specify a ranking with indifference between multiple objects. A serial dictatorship rule is associated with a priority order over the agents, and its outcomes are calculated inductively as follows: at each type profile, each agent receives his favorite object among those that are not assigned to an agent with higher priority. It is well-known that these rules are not only *strategy-proof*, but *group strategy-proof* (Svensson, 1999; Pápai, 2000).

² Though we cannot directly apply our theorems to labor markets or auctions as they violate *richness*, we can apply them indirectly by enriching the environment and then pruning the associated menu mechanism; see Appendix G.



(a) Menu mechanism G_1 .



(b) Menu mechanism G_2 .

Fig. 1. Serial dictatorship menu mechanisms for two agents and two objects. (a) Both an ex-post perfect implementation and a full subgame perfect implementation, but not a dominant strategy implementation. (b) Both an ex-post everywhere-dominant implementation and a full everywhere-dominant implementation.

We first consider the case where the agents in $\{1, 2\}$ are to consume the objects in $\{a, b\}$ and agent 1 has top priority. In this case, the simplest menu mechanism asks 1 to select any object, then assigns this selection to 1 and assigns the other object to 2. For the purposes of illustration, however, we will consider an alternative menu mechanism, which we call G_1 : 2 publicly selects any object (and 1 can observe this choice), then 1 selects any object, and finally 1 is assigned his selection while 2 is assigned the other object. Since there is perfect information, we say that G_1 is a *public* menu mechanism, and since no agent can select the same object twice, we say that G_1 is *non-repeating*. (More generally, *non-repeating* requires that each agent can either never select a previous choice or never select a previous rejection; the final example in this section, G_4 in Fig. 2 (b), features both cases.)

In Fig. 1 (a), there are four identical copies of G_1 in the middle row. In each column, the top row specifies a type profile and the bottom row specifies the associated outcome according to the rule. A convention suggests a strategy to each agent on the basis of his own type, and the unique *preferential* convention suggests that at each history the player should select his most-preferred object. In each column, the strategy profile suggested by the *preferential* convention for the associated type profile is highlighted over that column's copy of G_1 , and it is easy to verify that the convention is *compatible with the rule*: at each type profile, if the convention's suggestion is followed, then the outcome specified by the rule is reached.

What about incentives? In each column, the type profile and the menu mechanism together form a game, and it is easy to verify that the strategy profile specified by the convention is

a subgame perfect equilibrium. Because this occurs while each agent's strategy depends only on his own type, we say that G_1 is an *ex-post perfect implementation* of the rule. Moreover, though in each column there is a second subgame perfect equilibrium where 2 deviates from the convention, nevertheless at each type profile every subgame perfect equilibrium leads to the rule's outcome; thus we say that G_1 is a *full subgame perfect implementation* of the rule. In fact, we establish that this double implementation is not a peculiar feature of serial dictatorship, but rather holds far more generally as a logical consequence of the properties that we have highlighted with italics previously in this discussion (Proposition 2). This strong conclusion requires *group strategy-proofness*, but *ex-post perfect implementation* can be guaranteed for any *strategy-proof* rule using the other hypotheses even if we do not require all preference rankings to be strict (Proposition 1).

To conclude our analysis of G_1 , notice that at each type profile, the conventional strategy profile is *not* a dominant strategy equilibrium: if 1 selects whatever 2 selects, then 2 can profitably deviate from the convention by selecting his less-preferred object. It is therefore certainly not what we call an *everywhere-dominant strategy equilibrium*, which requires that from each history the continuation strategies form a dominant strategy equilibrium. That said, notice that 1 is only able to select whatever 2 selects because he observes whatever 2 selects.

Let G_2 denote the menu mechanism constructed from G_1 by placing both histories of agent 1 into the same information set (Fig. 1 (b)). Observe that for each type profile, the strategy profile specified by the convention is an everywhere-dominant strategy equilibrium, and though there is another such equilibrium where 2 deviates from the convention, nevertheless each everywhere-dominant strategy equilibrium leads to the rule's outcome. Thus in the same way that G_1 provides a double subgame perfect implementation, G_2 provides a double everywhere-dominant implementation.

To what extent does this double everywhere-dominant implementation generalize? A pessimist might observe that G_2 is effectively just a direct mechanism, and hypothesize that in general the only menu mechanisms that achieve such a double implementation are effectively direct mechanisms—in the sense that along each play, each agent dynamically reveals a complete strict ranking—with unobservable actions. By contrast, an optimist might hypothesize that we can always achieve double everywhere-dominant implementation whenever we start from a double subgame perfect implementation and then make all actions unobservable.

In fact, both the pessimistic hypothesis and the optimistic hypothesis are incorrect. To see this, consider the case where the agents in $\{1, 2\}$ are to consume the objects in $\{a, b, c\}$ and agent 1 has top priority. The pessimist is incorrect because menu mechanism G_3 (Fig. 2 (a)) achieves double everywhere-dominant implementation, but is not effectively a direct mechanism. The optimist is incorrect because G_4 (Fig. 2 (b)) does not achieve double everywhere-dominant implementation, even though all actions are unobservable and the associated public menu mechanism achieves double subgame perfect implementation. (Indeed, for G_4 agent 1 is able to *infer* the selection of 2 at the initial history even though he cannot observe it, creating the same issue discussed for G_1 .)

To distinguish between G_3 and G_4 , let us say that a mechanism is *reaction-proof* if and only if whenever one agent (the “observer”) can infer that another agent (the “deviator”) has deviated from an agreed upon strategy, at this point the deviator is safe from retaliation in the sense that his assignment has already been determined. More precisely, the observer's inference involves a pair of histories for the observer that do not share an information set, and we require only that the deviator's assignment has already been determined at one of these histories; see Section 2.3 for the formal definition. Observe that G_3 is *reaction-proof* while G_4 is not. In fact, we establish that whenever we satisfy the hypotheses for double subgame perfect implementation, and whenever

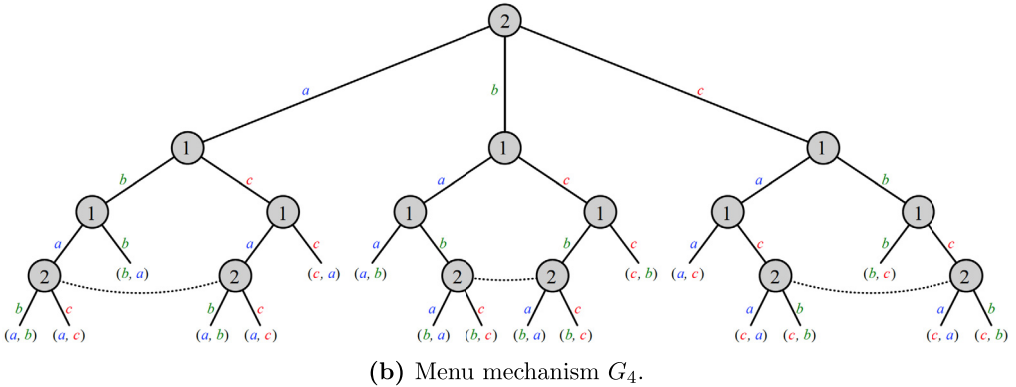



Fig. 2. Serial dictatorship menu mechanisms for two agents and three objects. (a) Double everywhere-dominant implementation, but not effectively a direct mechanism. (b) Not double everywhere-dominant implementation, but the associated public menu mechanism provides double subgame perfect implementation.

we are able to then thicken information sets to achieve *reaction-proofness*, the resulting menu mechanism is a double everywhere-dominant implementation (Theorem 2). As with our result for public menu mechanisms, this strong conclusion requires *group strategy-proofness*, but *ex-post everywhere-dominant implementation* can be guaranteed for any *strategy-proof* rule using the other hypotheses even if we do not require all preference rankings to be strict (Theorem 1).

1.3. Literature

Our paper is closely related to the literature on our leading examples; see Section 5. Moreover, our paper is related to recent results on special classes of menu mechanisms, implementation theory, and recent experiments.

First, our results complement recent results in the literature, which we describe using our language. First, Kawase and Bando (2021) prove that for each deferred acceptance proposal game associated with a public menu mechanism, honesty is a subgame perfect equilibrium; this is an implication of our Proposition 1. Interestingly, Kawase and Bando (2021) also consider the games where (i) only the side of the market that processes proposals is strategic, and (ii) both sides are strategic; though these games can also be described using menu mechanisms, the associated rules are not *strategy-proof*, and therefore our results do not apply. Second, Bó and Hakimov (2019) prove that for menu mechanisms derived from the deferred acceptance algorithm, honesty is a robust ordinal perfect Bayesian equilibrium, and Bó and Hakimov (2020b) extend this result to *pick-an-object mechanisms* for one-sided matching markets; these are menu mechanisms with some additional structure where each agent necessarily consumes the last ob-

ject he selected. These results are similar to our Proposition 1, but involve a natural subclass of our menu mechanisms and a different implementation notion.

Second, our paper is related to three topics in the broader literature on implementation theory: ex-post perfect implementation, double implementation, and obviously strategy-proof implementation. First, ex-post perfect implementation is a focal notion of robust implementation for dynamic mechanisms which has been used to analyze auctions (Ausubel, 2004; Ausubel, 2006; Sun and Yang, 2014; Drexl and Kleiner, 2015) and voting (Kleiner and Moldovanu, 2017; Gershkov et al., 2017; Kleiner and Moldovanu, 2019). Because we require *strictness* and *strategy-proofness*, our results only apply to restricted versions of these settings: (i) auctions where agents have unit demand, under the restriction that preferences are strict; and (ii) elections with two candidates. Indeed, our paper primarily complements these previous contributions by applying to matching environments. Second, double implementation refers to two kinds of implementation simultaneously (Maskin, 1979), and to our knowledge we are the first to consider full subgame perfect implementation (Moore and Repullo, 1988) in this context.³ Third, everywhere-dominance and *obvious* dominance (Li, 2017) both strengthen dominance in the context of dynamic games; we compare the two in Section 2.5.

Finally, recent experiments suggest that menu mechanisms may provide simplicity benefits, as measured by the likelihood of subjects conforming to the convention. In particular, the evidence suggests that while obviously strategy-proof mechanisms generally outperform both menu mechanisms and direct mechanisms (Bó and Hakimov, 2020b), menu mechanisms outperform direct mechanisms for both deferred acceptance (Klijn et al., 2019; Bó and Hakimov, 2020a) and top trading cycles (Bó and Hakimov, 2020b); thus menu mechanisms can provide simplicity benefits even when obviously strategy-proof implementations are not available.

2. Model

2.1. Environments

We begin by introducing a generic environment in our model: a finite setting with incomplete information, private values, and no consumption externalities.

Definition. An *environment* is a tuple $(N, (X_i)_{i \in N}, X, (\Theta_i)_{i \in N})$, where

- N is a nonempty and finite set of *agents*;
- for each $i \in N$, X_i is a nonempty and finite set of *assignments*, for which we let
 - (i) \mathcal{R}_i denote the set of (complete and transitive) *preference relations* on X_i , and
 - (ii) $\mathcal{P}_i \subseteq \mathcal{R}_i$ denote the set of those that are *strict* (that is, antisymmetric);
- $X \subseteq \times X_i$ is a nonempty set of *outcomes*, where each outcome consists of an assignment for each agent; and

³ Double implementation has previously been investigated for full Nash implementation with full undominated Nash implementation (Yamato, 1993), full Nash implementation with full dominant strategy implementation (Saijo et al., 2007), and full dominant strategy implementation with full ex-post Nash implementation (Hagiwara, 2020). We remark that a sufficient condition for double dominant implementation, *strict strategy-proofness*, has recently been considered in the context of auctions (Escudé and Sinander, 2020).

- for each $i \in N$, Θ_i is a nonempty set of types, where each type $\theta_i \in \Theta_i$ determines a preference relation $R_i(\theta_i) \in \mathcal{R}_i$; we let $P_i(\theta_i)$ denote the asymmetric part of $R_i(\theta_i)$ and let $I_i(\theta_i)$ denote the symmetric part of $R_i(\theta_i)$.

We let Θ denote $\times \Theta_i$, and refer to each $\theta \in \Theta$ as a *type profile*; we let $R(\theta) \equiv (R_i(\theta_i))_{i \in N}$ denote the associated *preference profile*. We assume there are no consumption externalities and thus sometimes abuse notation, letting $R_i(\theta_i)$ denote not only a binary relation on X_i but also the associated binary relation on X . For convenience, whenever we refer to a generic environment we implicitly assume all of this notation.

Many familiar economic settings—such as matching with contracts, labor markets, auctions, school choice, object allocation, and two-candidate elections—can be articulated as environments with additional structure. For our illustrative example (Section 1.2), we define an object allocation environment as follows: (i) N is the set of agents, (ii) there are at least as many objects as agents, and for each $i \in N$, X_i is the set of objects, (iii) X is the set assignment profiles at which no two agents consume the same object, and (iv) for each $i \in N$, Θ_i is the set of strict rankings of objects. For details about all of our applications—including how they fit the other primitives and hypotheses that we introduce in this section—see Appendix G.

Our results involve two hypotheses about environments, both of which regard the preferences an agent might have:

Definition. *Hypotheses for environments.* We say that an environment is

- *rich* if and only if for each $i \in N$ and each $P_i \in \mathcal{P}_i$, there is $\theta_i \in \Theta_i$ such that $R_i(\theta_i) = P_i$; and
- *strict* if and only if for each $i \in N$ and each $\theta_i \in \Theta_i$, $R_i(\theta_i) \in \mathcal{P}_i$.

Richness rules out many preference restrictions by requiring that each agent might have any strict ranking of his assignments; this is generally satisfied in matching environments, but violated in auction environments (where lower payments must be preferred) and labor market environments (where higher salaries must be preferred). *Strictness* rules out indifference, which is common in matching environments but uncommon in auction environments. That said, we remark that *strictness* holds for auctions where each admissible valuation is generic, in the sense that it belongs to an open ball whose members all induce the same ranking of assignments—for example, a single-object auction with an integer bid increment and non-integer valuations.

2.2. Rules

In a given environment, the agents wish to condition the outcome on their collective private information according to a (social choice) rule:

Definition. Fix an environment. A *rule* is a function $f : \Theta \rightarrow X$.

For our illustrative example, each priority order over the agents specifies a *serial dictatorship rule*, which assigns to each type profile the outcome determined by the following *serial dictatorship algorithm*: assign the agent with highest priority his most-preferred object, then assign the agent with second-highest priority his most-preferred object of those that have not yet been

assigned, and continue in this fashion until all agents have received objects. We emphasize that while the term *serial dictatorship* is often used interchangeably for the rule and the associated algorithm, for our purposes it is important to make a distinction because the algorithm is even more closely associated with the mechanism than the rule; the same remark applies to the cumulative offers process (matching with contracts), deferred acceptance (school choice), and top trading cycles (object allocation).

Our results involve two hypotheses about rules, both of which regard incentive compatibility:

Definition. *Hypotheses for rules.* Fix an environment and a rule f . We say that f is

- *strategy-proof* if and only if for each $\theta \in \Theta$, each $i \in N$, and each $\theta'_i \in \Theta_i$, we have $f(\theta_i, \theta_{-i}) R_i(\theta_i) f(\theta'_i, \theta_{-i})$; and
- *group strategy-proof* if and only if there is no $\theta \in \Theta$, $N' \subseteq N$, and $\theta'_{N'} \in \times_{N'} \Theta_i$ such that
 - (i) for each $i \in N'$, $f(\theta'_{N'}, \theta_{N \setminus N'}) R_i(\theta_i) f(\theta)$; and
 - (ii) there is $i \in N'$ such that $f(\theta'_{N'}, \theta_{N \setminus N'}) P_i(\theta_i) f(\theta)$.

Strategy-proofness is the usual requirement that for each type profile, the strategy profile where all agents honestly report their types is a dominant strategy equilibrium in the game given by the direct mechanism and the type profile. *Group strategy-proofness* requires that moreover, no coalition of agents can ever benefit by deviating from this equilibrium—in the sense that at least one member is better off while no member is worse off—through a coordinated misrepresentation of their preferences. All of our results require *strategy-proofness*, and our strongest conclusions require *group strategy-proofness*.

2.3. Menu mechanisms

We are interested in implementing a given rule with a mechanism, or an extensive game form with players in N and outcomes in X ; see Appendix A for the formal definition, which is familiar to most readers. The notation we use throughout the paper is gathered in Table 1. For convenience, whenever we refer to a generic mechanism G , we implicitly assume all of this notation. Note that a mechanism G and a type profile θ together determine a game $(G, R(\theta))$.

We are now ready to introduce menu mechanisms, where agents iteratively select from menus of their assignments. Menu mechanisms are related in spirit to *natural implementation*, where an agent’s *strategy* can be interpreted as a consumption bundle (Saijo et al., 1996); the difference is that with menu mechanisms, an agent’s *action* can be interpreted as a consumption bundle:

Definition. Fix an environment. A mechanism is moreover a *menu mechanism* if and only if

- $\mathcal{A} = \cup X_i$; and
- for each $i \in N$, each $h \in H_i$, and each $h' \in \sigma(h)$, $\alpha(h') \in X_i$. We define the *menu at h* , $X_i(h) \subseteq X_i$, by $X_i(h) \equiv \{\alpha(h') | h' \in \sigma(h)\}$.

As discussed earlier, many menu mechanisms can be derived naturally from algorithms associated with familiar rules. For our illustrative example, menu mechanism G_3 (Fig. 2 (a)) is derived naturally from the serial dictatorship algorithm described earlier. Similarly, the deferred acceptance algorithm for school choice yields menu mechanisms where at each history, the player is a student who is not tentatively accepted by any school, and the player selects a school that

Table 1
Notation for a generic mechanism.

Name	Notation	Representative element
Set of histories	H	h
Precedence relation over histories	\succsim	
Set of immediate successors of h	$\sigma(h)$	
Set of plays	Π	π
Set of terminal histories	Z	z
Player function	\mathbb{P}	
Set of histories that belong to i	H_i	
Set of actions	\mathcal{A}	
Action function	α	
Set of actions available at h	$\mathcal{A}(h)$	
Action taken at h to remain on π	$\alpha^h(\pi)$	
Action taken at h to continue toward h'	$\alpha^h(h')$	
Information partition for i	\mathbb{I}_i	\mathcal{I}_i
Set of actions available at \mathcal{I}_i	$\mathcal{A}(\mathcal{I}_i)$	
Outcome function	\mathcal{X}	

has not yet rejected him. We remark that for the Vickrey rule for one-object auctions, we are less interested in menu mechanisms derived from the synonymous algorithm (the simple “second-price” calculation) and more interested in those derived from an alternative algorithm (given by the English auction).

In general, we are able to draw stronger conclusions when less information is available to the agents. We highlight two extreme cases: all actions are observable and no actions are observable. The former plays a technical role in our proofs as we begin by analyzing mechanisms with perfect information, while the latter is simply convenient for providing examples that meet the requirements of our main results:

Definition. *Information assumptions for menu mechanisms.* Fix an environment. We say that a menu mechanism is

- *public* if and only if for each $i \in N$, $\mathbb{I}_i = \{\{h\} | h \in H_i\}$; and
- *private* if and only if for each $i \in N$ and each pair $h, h' \in H_i$, h and h' share an information set if and only if at these histories, i has encountered the same menus and taken the same actions in the same order.⁴

We remark that public and private versions of a given menu mechanism describe dramatically different institutions. For example, a public deferred acceptance menu mechanism might describe courtship in a ballroom, while a private one might describe courtship through a dating app. As another example, a public English auction menu mechanism might describe a sale at an auction house, while a private one might describe a sale online (say, on e-Bay). In general, we find it

⁴ Formally, for each $i \in N$ and each $h \in H_i$, let $\{h_1, h_2, \dots, h_t\}$ denote $\{h' \in H_i | h' < h\}$ such that $h_1 < h_2 < \dots < h_t$, and define the *experience of i at h* to be the list of menus and selections $\mathcal{E}_i(h) \equiv ((X_i(h_1), \alpha^{h_1}(h)), (X_i(h_2), \alpha^{h_2}(h)), \dots, (X_i(h_t), \alpha^{h_t}(h)), X_i(h))$. We require that for each $i \in N$ and each pair $h, h' \in H_i$, h and h' share an information set if and only if $\mathcal{E}_i(h) = \mathcal{E}_i(h')$.

convenient to think of private menu mechanisms as apps for electronic devices that occasionally notify users that they must select from a given menu.

In order to introduce our hypotheses about menu mechanisms, we first recall the familiar concepts of pure strategy profiles and plays. Moreover, we introduce notation for the play that first proceeds from the initial history to a given history, then proceeds according to a given strategy profile.

Definition. Strategy profiles, plays, and related notation. Fix an environment and a mechanism.

- For each $i \in N$, a (pure) strategy for i is a mapping $s_i : H_i \rightarrow \mathcal{A}$ such that
 - (i) for each $h \in H_i$, $s_i(h) \in \mathcal{A}(h)$; and
 - (ii) for each $\mathcal{I}_i \in \mathbb{I}_i$ and each pair $h, h' \in \mathcal{I}_i$, $s_i(h) = s_i(h')$.
 We let S_i denote the set of strategies for i . A strategy profile is a profile of strategies $s = (s_i)_{i \in N}$, and we let $S \equiv \times S_i$ denote the set of strategy profiles.
- A play is a maximal chain of histories, which gives a complete description of a sequence of choices; we write π for a play and Π for the set of plays.
- For each $h \in H$ and each $s \in S$, we define $\pi^h(s)$ to be the play that first proceeds from the initial history to h and then proceeds according to s . Moreover, we define $\mathcal{X}^h(s) \equiv \mathcal{X}(\pi^h(s))$, and for each $i \in N$ let $\mathcal{X}_i^h(s)$ denote the associated assignment. When h is the initial history, we simply write $\pi(s)$ and $\mathcal{X}(s)$.

Our results involve two hypotheses about menu mechanisms, the first regarding individual agents and the second regarding pairs of agents:

Definition. Hypotheses for menu mechanisms. Fix an environment and a menu mechanism G . We say that G is

- *non-repeating* if and only if for each $i \in N$, either
 - (i) *i has non-repeating choices*: for each pair $h, h' \in H_i$ such that $h < h'$, we have that $\{\alpha^h(h')\} \cap X_i(h') = \emptyset$, or
 - (ii) *i has non-repeating rejections*: for each pair $h, h' \in H_i$ such that $h < h'$, we have that $(X_i(h) \setminus \{\alpha^h(h')\}) \cap X_i(h') = \emptyset$; and
- *reaction-proof* if and only if for each distinct pair $i, j \in N$, each $h \in H_i$, each $s_{-i} \in S_{-i}$, each pair $s'_i, s''_i \in S_i$, and each pair $h', h'' \in H_j$ such that
 - $h' \in \pi^h(s'_i, s_{-i})$ and $h'' \in \pi^h(s''_i, s_{-i})$,
 - j plays the same number of times before h' and h'' , and
 - h' and h'' are in different information sets,
 there is $h^\# \in \{h', h''\}$ at which the assignment for i has already been determined: for each pair $\pi_1, \pi_2 \in \Pi$ such that $h^\# \in \pi_1 \cap \pi_2$, we have $\mathcal{X}_i(\pi_1) = \mathcal{X}_i(\pi_2)$.

Non-repeating requires that each agent can either never select a previous choice or never select a previous rejection. All of the menu mechanisms we have discussed thus far in Section 2 satisfy *non-repeating choices* for all agents, while *non-repeating rejections* holds for (i) agent 1 in mechanism G_4 (Fig. 2 (b)), and (ii) all agents in Vickrey’s variant of the Dutch auction, where prices descend until two agents have entered, at which point the first entrant wins (Vickrey, 1961).

Reaction-proofness requires that whenever an agent i (the “deviator”) signals something to another agent j (the “observer”)—in the sense that if j knows the peers of i have played s_{-i} since h , then he can distinguish between whether i has played s'_i or s''_i since h —then by the time j acquires this information, i is already safe from any reaction—in the sense that for each pair of histories where j acquires this information, in at least one of these histories the assignment i receives has already been determined. As discussed in Section 1.2, mechanism G_3 is *reaction-proof* while mechanism G_4 is not (Fig. 2).

Before proceeding, we illustrate that *reaction-proofness* holds for three additional examples. For simplicity we focus on private menu mechanisms, where if there is a pair of histories $h', h'' \in H_j$ that satisfy the hypotheses for *reaction-proofness*, then there is an earliest such pair that j is able to distinguish between only because he receives two distinct menus at these histories.

Example 1. Reaction-proofness and direct menu mechanisms. Consider any private menu mechanism derived from a direct mechanism (where the first agent fully reveals his ranking by selecting his favorite assignment, then his favorite assignment that he has not already selected, and so on; then the second agent does so; and so on). The result is *reaction-proof* because at each history, the player’s menu is always the set of assignments that the player has not already selected, and is thus determined by the player’s strategy alone.

Example 2. Reaction-proofness and deferred acceptance. Consider any private menu mechanism derived from the deferred acceptance algorithm (where at each history the player selects an assignment that has not yet rejected him). Even though the order in which the agents act is not trivial, the same argument used for Example 1 establishes that the menu mechanism is *reaction-proof*.

Example 3. Reaction-proofness and top trading cycles. This example involves some nuance. For each private menu mechanism derived from the top trading cycles algorithm (where at each history the player selects an unassigned object to point toward), let us say that an agent is *active* if (i) he has not yet played, or (ii) the object he previously selected has been removed. Let us say that there is a *simple order* if the player is always the active agent with lowest index. We claim that any such menu mechanism is *reaction-proof*.

Indeed, suppose that (i) i is at history h while the other agents play s_{-i} , (ii) if i plays s'_i from h onward, then j faces histories $h'_1 < h'_2 < \dots < h'_t$, (iii) if i plays s''_i from h onward, then j faces histories $h''_1 < h''_2 < \dots < h''_t$, and (iv) j faces the same sequence of menus in both cases until the end, where the two menus might differ: for each $t^* \in \{1, 2, \dots, t - 1\}$ we have $X_j(h'_{t^*}) = X_j(h''_{t^*})$. If the object owned by i is missing from $X_j(h'_t)$ or $X_j(h''_t)$ (or both), then the assignment of i has already been determined at the associated history and we are done. If instead the object owned by i belongs to both $X_j(h'_t)$ and $X_j(h''_t)$, then i has not completed a cycle before either h'_t or h''_t , and we claim that both menus must be the same. Indeed, both (i) between h and h'_t , and (ii) between h and h''_t , each peer of i faces the same sequence of menus in the same sequence of information sets after h , and thus necessarily makes the same choices in the same order. It is possible that i plays at a different number of histories (i) between h and h'_t , and (ii) between h and h''_t , because i selects objects that are removed in cycles that do not include him more often along one of the two plays, but as i never completes a cycle before either h'_t or h''_t , thus he does not impact the order in which his peers play or the sequence of menus faced by any peer, so $X_j(h'_t) = X_j(h''_t)$ as claimed. Altogether, then, the menu mechanism is *reaction-proof*.

Notice that without a simple order, the menu mechanism may not be *reaction-proof*: s'_i could cause j to play immediately, while s''_i could cause j to play in a few turns after another cycle has been completed, with the assignment of i still undetermined in both cases.

2.4. Conventions

We focus on implementation that is robust, in that it does not rely on any assumptions about the beliefs agents have about the private information of their peers. Informally, we say that a mechanism implements the rule if and only if regardless of the type profile, the desired outcome is a plausible consequence of the strategic choices of the agents. We formalize this in several ways, using both type-strategy profiles in a mechanism and strategy profiles in the associated games. To ease the discussion, we introduce the term *convention* as a suggestive shorthand for a type-strategy profile—for example, in a direct mechanism, honesty is a convention.

Definition. Conventions. Fix an environment and a mechanism. For each $i \in N$, a *type-strategy* for i is a mapping $S_i : \Theta_i \rightarrow S_i$. A *convention* is a profile of type-strategies $S = (S_i)_{i \in N}$.

For our illustrative example, the convention specifies that at each history, the player should select his most-preferred object from the menu (according to his type). More generally, we are interested in situations where such a convention yields a calculation algorithm for its rule:

Definition. Hypotheses for conventions. Fix an environment, a rule f , a menu mechanism G , and a convention S . We say that S is

- *f-compatible* (or *compatible with the rule*) if and only if for each $\theta \in \Theta$, we have $\mathcal{X}(S(\theta)) = f(\theta)$; and
- *preferential* if and only if for each $i \in N$ and each $\theta_i \in \Theta_i$, there is a *tie-breaker* $\tau_i(\theta_i) \in \mathcal{P}_i$ such that for each $h \in H_i$, $[S_i(\theta_i)](h) = \operatorname{argmax}_{\tau_i(\theta_i)}[\operatorname{argmax}_{R_i(\theta_i)} X_i(h)]$.⁵

Compatibility with the rule requires that the rule’s outcome is always reached if agents follow the convention, while the *preferential* requirement imposes that the convention simply specifies that each agent should always pick a most-preferred assignment, breaking ties consistently. It is easy to verify that both conditions are satisfied by the conventions for all of the examples we have discussed thus far.

2.5. Implementation

In order to consider incentives in a mechanism G , we first consider incentives in each game (G, R) . In particular, a *solution concept* associates each of these games with a collection of plausible strategy profiles, and we are primarily interested in the novel (to our knowledge) solution concept of *everywhere-dominance*. In order to facilitate our analysis and discussion, we also recall three standard solution concepts:

Definition. Solution concepts. Fix an environment, a mechanism, and a preference profile. Each solution concept **SC** gives a collection of strategy profiles $\mathbf{SC}(G, R) \subseteq S$. We consider:

⁵ Abusing notation, if $\operatorname{argmax}_{R_i}(X'_i)$ is a singleton $\{x\}$, we sometimes let $\operatorname{argmax}_{R_i}(X'_i)$ denote x .

- *Nash equilibrium*: $s \in \mathbf{NE}(G, R)$ if and only if for each $i \in N$ and each $s'_i \in S_i$, $\mathcal{X}(s) R_i \mathcal{X}(s'_i, s_{-i})$;
- *dominant strategy equilibrium*: $s \in \mathbf{DE}(G, R)$ if and only if for each $i \in N$, each $s'_{-i} \in S_{-i}$, and each $s'_i \in S_i$, $\mathcal{X}(s_i, s'_{-i}) R_i \mathcal{X}(s'_i, s'_{-i})$;
- *subgame perfect equilibrium*: if G has perfect information, then $s \in \mathbf{SPE}(G, R)$ if and only if for each $i \in N$, each $h \in H_i$, and each $s'_i \in S_i$, $\mathcal{X}^h(s) R_i \mathcal{X}^h(s'_i, s_{-i})$; and
- *everywhere-dominant strategy equilibrium*: $s \in \mathbf{EDE}(G, R)$ if and only if for each $i \in N$, each $h \in H_i$, each $s'_{-i} \in S_{-i}$, and each $s'_i \in S_i$, $\mathcal{X}^h(s_i, s'_{-i}) R_i \mathcal{X}^h(s'_i, s'_{-i})$.

It is easy to verify that everywhere-dominance is stronger than both dominance and subgame perfection, which in turn are both stronger than Nash. Notice that unlike subgame perfection, everywhere-dominance is defined for games with imperfect information: at an information set with multiple histories, the player’s continuation strategy must be a best response at *each* of these histories. Indeed, everywhere-dominance requires the player’s strategy to *always* be a best response—regardless of his beliefs about how his peers have played thus far, how his peers will play henceforth, and any other information he might lack about his current history—even if he has already deviated from this strategy (for example, due to error).

Given a solution concept for each game (G, R) , we consider three ways of articulating implementation for G^6 : (i) *ex-post with respect to a convention*, which requires that at each type profile, the convention specifies an equilibrium that leads to the desired outcome; (ii) *full*, which requires that at each type profile, there are equilibria and each of them yields the desired outcome; and (iii) *double*, which we use in this paper to refer to both (i) and (ii) simultaneously:

Definition. Implementations. Fix an environment, a rule, and a solution concept **SC**. For each mechanism G and each convention \mathbb{S} , we say that (G, \mathbb{S}) is an *ex-post SC-implementation of f* if and only if

- for each $\theta \in \Theta$, $\mathcal{X}(\mathbb{S}(\theta)) = f(\theta)$; and
- for each $\theta \in \Theta$, $\mathbb{S}(\theta) \in \mathbf{SC}(G, R(\theta))$.

In this case, we also say that G is an *ex-post SC-implementation of f via \mathbb{S}* . We say that G is a *full SC-implementation of f* if and only if

- for each $\theta \in \Theta$, $\mathbf{SC}(G, R(\theta)) \neq \emptyset$; and
- for each $\theta \in \Theta$ and each $s \in \mathbf{SC}(G, R(\theta))$, $\mathcal{X}(s) = f(\theta)$.

Finally, we say that G is a *double SC-implementation of f* if and only if it is both an ex-post **SC**-implementation of f and a full **SC**-implementation of f .

When we use these terms for particular solution concepts, we sometimes replace **DE** with dominant, **SPE** with perfect or subgame perfect, and **EDE** with everywhere-dominant; for example, we might refer to ex-post perfect implementation or full subgame perfect implementation.

⁶ We remark that another approach to robustness is to pursue interim implementation (or Bayesian implementation when there need not be a common prior) for all beliefs; see for example Penta (2015) and Bó and Hakimov (2020b). This is implied by ex-post implementation.




Fig. 3. *Independence of implementations.* (a) Only ex-post perfect. Consider $N = \{1, 2\}$, $X_1 = X_2 = \{a, b\}$, the types are the strict rankings, the rule specifies that 1 always receives a while 2 receives the assignment preferred by 1, and the mechanism in Fig. 3 (a). (b) Only full subgame perfect. Consider $N = \{1, 2\}$, $X_1 = X_2 = \{a, b, c\}$, the types are the strict rankings, the rule maps each preference profile to the unique subgame perfect outcome of the game given by that profile and the mechanism in Fig. 3 (b), and that same mechanism. There is no ex-post perfect convention because, for example, each convention either specifies 1 should select a when his strict ranking is acb , or it specifies that he should select c when he has this ranking, but it does not specify both.

For our illustrative example, menu mechanism G_3 (Fig. 2 (a)) is both (i) an ex-post everywhere-dominant implementation with respect to the unique *preferential* convention, and (ii) a full ex-post everywhere-dominant implementation; it is therefore a double everywhere-dominant implementation. Before proceeding, we briefly discuss the interpretation of ex-post implementation, the independence of ex-post implementation and full implementation, and the independence of everywhere-dominance and obvious dominance (Li, 2017).

Interpretation of ex-post implementation. For simplicity, we focus on the case of public menu mechanisms; in this case ex-post implementation can be understood entirely in terms of each agent i 's first-order beliefs over $\Theta_{-i} \times S_{-i}$. In particular, agent i does not know the type profile of his peers or the strategy profile of his peers, and he may update his beliefs about this at each history. For ex-post perfection, if at each history i believes that his peers will surely follow the convention in the future (regardless of what has happened in the past), then his conventional strategy is a best response. For ex-post everywhere-dominance, at each history the conventional strategy is a best response to *any* beliefs. In both cases, if the environment is *rich*, the public menu mechanism is *non-repeating*, and the convention is *preferential*, then there are such beliefs that are compatible with (correct) common strong belief in rationality.⁷

Independence of ex-post implementation and full implementation. Ex-post implementation and full implementation are logically independent; we illustrate this for subgame perfect equilibrium in Fig. 3. Indeed, the former requires that each agent's strategy depends only on his own type, while the latter does not; the latter requires that all equilibria are compatible with the rule, while the former does not.

⁷ This can be formalized using *strong Δ -rationalizability* (Battigalli and Siniscalchi, 2003; Battigalli and Siniscalchi, 2007), a generalization of *extensive form rationalizability* (Pearce, 1984)—or equivalently, the iterative removal of *conditionally dominated* strategies (Shimoji and Watson, 1998)—for games with incomplete information. For ex-post perfection, let Δ require (in addition to the standard requirements) that each agent i assigns probability one to the event $\{(\theta_{-i}, S_{-i}(\theta_{-i})) | \theta_{-i} \in \Theta_{-i}\}$ at each history, regardless of his type. By the *Play Lemma* (Appendix D), the hypothesis that one's peers are following the convention can never be falsified, so each history indeed has admissible beliefs. For ex-post everywhere-dominance, Δ need only satisfy the standard requirements. In both cases, the convention survives the iterative deletion procedure in Definition 6 of Battigalli and Siniscalchi (2007).

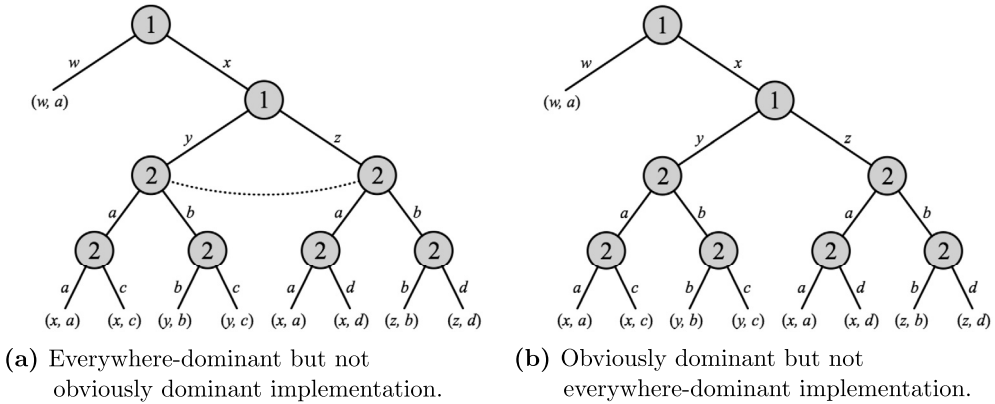


Fig. 4. Independence of everywhere-dominant and obviously dominant implementations. Consider $N = \{1, 2\}$, $X_1 = \{w, x, y, z\}$, $X_2 = \{a, b, c, d\}$, 1 can have rankings in $\{wxyz, yxzw, zxyw\}$, 2 can have any strict ranking, and the rule is given by the preferential convention for both menu mechanisms. (a) When 2 has ranking $dabc$, his preferential strategy is not obviously dominant because at its first information set, the worst case from adhering to the preferential strategy (that is, a) is worse than the best case from deviating (that is, d). (b) When 1 has ranking $wxyz$, his preferential strategy is not everywhere-dominant because when selecting from $\{y, z\}$, the preferential strategy might lead to y while a deviation leads to x and therefore is profitable.

Independence of everywhere-dominance and obvious dominance. As solution concepts, everywhere-dominance and obvious dominance (Li, 2017) both strengthen the standard notion of dominance in the context of dynamic games. Moreover, both have some relationship with menu mechanisms: our main results establish the link for everywhere-dominance, while recent literature has established the link for obvious dominance.⁸ What then is the relationship between these two solution concepts?

For context, a strategy is obviously dominant if and only if at each information set it can reach, the worst case from adhering to the strategy is at least as good as the best case from deviating. In this case, the strategy can be identified as dominant by an agent who has trouble with contingent reasoning, in the sense that he struggles to compare strategies on a case-by-case basis. By contrast, identifying an everywhere-dominant strategy may require contingent reasoning, but such a strategy retains its strong incentives even at information sets that it cannot reach—for example, if the player has previously made a typo. Altogether, then, obvious dominance imposes a stronger requirement on-path and everywhere-dominance imposes a stronger requirement off-path, suggesting that the two concepts are independent, and indeed this is the case.

In fact, the logical independence of these two solution concepts introduces a novel design consideration: concealing information facilitates everywhere-dominance while revealing information facilitates obvious dominance. Indeed, everywhere-dominant implementation is always preserved when thickening information sets (provided that all conventional strategies remain available, which is always the case for preferential strategies in menu mechanisms) because from each history one’s opponents have access to fewer strategies. On the other hand, obviously

⁸ In particular, for rich and strict environments, if a rule has an obviously strategy-proof implementation, then it has one through a millipede mechanism (Pycia and Troyan, 2019), in which case it also has an implementation in weakly dominant strategies through a menu mechanism that is moreover a pick-an-object mechanism (Bó and Hakimov, 2020b). There are millipede mechanisms that are not menu mechanisms, and there are menu mechanisms that are not millipede mechanisms.

strategy-proof implementation is always preserved when information sets are broken (provided that we begin from a mechanism that satisfies perfect recall) because the worst case from adhering to the convention can only improve while the best case from deviating can only worsen (Ashlagi and Gonczarowski, 2018; see also Mackenzie, 2020). In fact, for the example in Fig. 4, there are only two histories that might share an information set, and the choice to pair them or separate them determines whether or not there is only an everywhere-dominant implementation or only an obviously dominant implementation: when we move from the mechanism on the left to the mechanism on the right by making all actions observable, we lose everywhere-dominance but gain obvious dominance.

3. Menu mechanisms for strategy-proof rules

3.1. Result

The classic revelation principle states that for each *strategy-proof* rule, the associated direct mechanism provides an ex-post dominant implementation (Gibbard, 1973; Myerson, 1981). In this section, we consider the full class of these rules, and investigate conditions under which a menu mechanism provides an alternative such implementation that moreover provides robust incentives off the path of play.

To begin, we first investigate the robust incentive compatibility of menu mechanisms for which all actions are observable. In particular, we provide conditions under which a public menu mechanism provides an ex-post perfect implementation:

Proposition 1. *For each rich environment, each strategy-proof rule, each non-repeating public menu mechanism, and each preferential convention that is compatible with the rule, the public menu mechanism is an ex-post perfect implementation of the rule via the convention.*

The formal proof is in Appendix E, and involves lemmas about public menu mechanisms (Appendix D) whose proofs involve lemmas about revealed preference theory (Appendix C). We sketch the arguments below:

Proof sketch. To begin, we take arbitrary $\theta \in \Theta$, $i \in N$, $h \in H_i$, and $s_i \in S_i$, we define a to be the assignment for i when he conforms to the convention and b to be the assignment for i when he deviates—that is, $a \equiv \mathcal{X}_i^h(\mathbb{S}(\theta))$ and $b \equiv \mathcal{X}_i^h(s_i, \mathbb{S}_{-i}(\theta_{-i}))$ —and we seek to prove that $a R_i(\theta_i) b$, dismissing the trivial case where $a = b$. In order to do so, we consider the plays through h where i conforms and where he deviates, $\pi_a \equiv \pi^h(\mathbb{S}(\theta))$ and $\pi_b \equiv \pi^h(s_i, \mathbb{S}_{-i}(\theta_{-i}))$, and seek type profiles for which the convention specifies these plays. Because the environment is *rich* and the convention is *preferential*, this is possible if for each agent, certain choices from the menus of certain histories are together rationalizable, and we therefore apply techniques from revealed preference theory.

In particular, we use the well-known result that choices are rationalizable by a strict preference relation if and only if there is no revealed *cycle* (the [Cycle Lemma](#); see for example Chambers and Echenique, 2016) to prove that for each agent, (i) always staying on a given play is specified by the convention for some type (the [Play Lemma](#)); and (ii) for each history and each continuation strategy specified by the convention, proceeding toward that history and then conforming to the continuation strategy is specified by the convention for some type (the [Continuation Lemma](#)). Both of these proofs rely critically on the fact that the public menu mechanism is *non-repeating*.

Our proof of Proposition 1 also involves a second result from revealed preference theory: if all strict preference relations that rationalize some choices rank a above b , then there is a revealed path from a to b (the Path Lemma).

By the Continuation Lemma, there is $\theta_{-i}^* \in \Theta_{-i}$ at which the convention has each agent $j \in N \setminus \{i\}$ proceed to h and then play according to $S_j(\theta_j)$. By the Play Lemma, i remaining on π_a and i remaining on π_b are both rationalizable. If P_a and P_b rationalize i remaining on π_a and π_b , respectively, then by richness they are given by types θ_i^a and θ_i^b , so as the convention is preferential we have $\pi_a = \pi(S(\theta_i^a, \theta_{-i}^*))$ and $\pi_b = \pi(S(\theta_i^b, \theta_{-i}^*))$. We can therefore apply strategy-proofness to deduce that $a P_a b$ and $b P_b a$. Since P_a and P_b were arbitrary rationalizations, thus by the Path Lemma, i remaining on π_a reveals a path from a to b , and i remaining on π_b reveals a path from b to a .⁹ Finally, we use both revealed paths to prove that the former revealed path is revealed entirely after h , which as the convention is preferential implies that $a R_i(\theta_i) b$, as desired. \square

Before proceeding, we remark that Proposition 1 does not hold if (i) we weaken non-repeating by simply requiring that the choices along each play can be rationalized by some strict preference profile, or (ii) we weaken preferential by simply requiring that the convention always asks each agent to select a most-preferred assignment without necessarily breaking ties consistently; see Fig. 7 in Appendix E.

As discussed in Section 1.2, an agent’s conventional strategy is everywhere-dominant when the ability of his peers to react to deviations is limited. Our first theorem states that if we take the hypotheses for Proposition 1 and then replace the requirement that all actions are observable with reaction-proofness, then the resulting hypotheses guarantee ex-post everywhere-dominant implementation¹⁰:

Theorem 1. *For each rich environment, each strategy-proof rule, each non-repeating and reaction-proof menu mechanism, and each preferential convention that is compatible with the rule, the menu mechanism is an ex-post everywhere-dominant implementation of the rule via the convention.*

Proof. We consider both (i) the non-repeating and reaction-proof menu mechanism G , and (ii) its associated public menu mechanism $G^!$. For each $i \in N$, let S_i denote the set of strategies for i for the former mechanism and let $S_i^!$ denote this set for the latter mechanism. Observe that any strategy for i specified by a preferential convention belongs to $S_i \subseteq S_i^!$, as such a strategy never has i select distinct actions at distinct histories that share a menu. It follows that S is a preferential convention compatible with the rule for both G and $G^!$.

Assume, by way of contradiction, that (G, S) is not an ex-post everywhere-dominant implementation of the rule. Then there are $i \in N$, $\theta_i \in \Theta_i$, $h \in H_i$, and $s_{-i} \in S_{-i}$ such that i has a

⁹ Note that if at history h , selecting a leads toward π_a and selecting b leads toward π_b , then we trivially have both revealed paths; we use the Path Lemma because this need not be the case. For example, consider a deferred acceptance menu mechanism where the menu at h is $\{a, b, x, y\}$, selecting x leads toward π_a with some future choices that eventually result in a , and selecting y leads toward π_b with some future choices that eventually result in b .

¹⁰ In earlier versions of this paper, both Theorem 1 and Theorem 2 required the additional hypothesis that no actions are observable; we thank an anonymous referee and Inácio Bó for correctly conjecturing that this requirement could be dropped by generalizing the earlier version of our reaction-proofness condition.

profitable deviation $s_i'' \in S_i$ from $S_i(\theta_i)$ at h when his type is θ_i and his peers play s_{-i} . To use the notation from the definition of *reaction-proofness*, define $s_i' \equiv S_i(\theta_i)$.

By Proposition 1, (G^1, S) is an ex-post perfect implementation of the rule via S ; thus the restriction of s_{-i} to the given subgame is not specified by the convention for any type profile. Let j be the lowest-index peer of i whose restricted strategy is never specified by the convention. We claim that j has a type θ_j^* such that i has the same profitable deviation $s_i'' \in S_i$ from s_i' at h when his type is θ_i and his peers play $(S_j(\theta_j^*), s_{-i,j})$.

Define the plays $\pi' \equiv \pi^h(s_i', s_{-i})$ and $\pi'' \equiv \pi^h(s_i'', s_{-i})$. Let h_1', h_2', \dots label the histories where j plays after h along π' in order, let h_1'', h_2'', \dots label the histories where j plays after h along π'' in order, and let T denote the set of $t \in \mathbb{N}$ such that there are both h_t' and h_t'' . We consider two cases:

CASE 1: For each $t \in T$, both h_t' and h_t'' share an information set in G . In this case, let $\pi^+ \in \{\pi', \pi''\}$ maximize the number of histories where j plays after h , and let θ_j^* be such that j adheres to π^+ after h according to convention (which exists by the [Play Lemma](#)). Then for each $t \in T$, (i) $s_j(h_t') = s_j(h_t'')$, because $s_j \in S_j$; and (ii) $[S_j(\theta_j^*)](h_t') = [S_j(\theta_j^*)](h_t'')$, because the convention is *preferential*; so necessarily both s_j and $S_j(\theta_j^*)$ agree along π' and π'' after h , so $S_j(\theta_j^*)$ adheres to both π' and π'' after h . Clearly, then, when the peers of i change from s_{-i} to $(S_j(\theta_j^*), s_{-i,j})$, $s_i'' \in S_i$ remains a profitable deviation for i from s_i' as claimed.

CASE 2: There is $t \in T$ such that h_t' and h_t'' are in different information sets in G . In this case, let t^\neq denote the earliest such $t \in T$. By *reaction-proofness*, there is $h^\neq \in \{h_{t^\neq}', h_{t^\neq}''\}$ such that for each pair $\pi_1, \pi_2 \in \Pi$ such that $h^\neq \in \pi_1 \cap \pi_2$, we have $\mathcal{X}_i(\pi_1) = \mathcal{X}_i(\pi_2)$. Let $\pi^\neq \in \{\pi', \pi''\}$ be the play that contains h^\neq , and let $\pi^* \in \{\pi', \pi''\}$ be the other play.

Let $\theta_j^* \in \Theta_j$ be such that j adheres to π^* after h according to convention (which exists by the [Play Lemma](#)). We claim

- $\mathcal{X}_i(\pi^h(s_i', s_{-i})) = \mathcal{X}_i(\pi^h(s_i', S_j(\theta_j^*), s_{-i,j}))$, and
- $\mathcal{X}_i(\pi^h(s_i'', s_{-i})) = \mathcal{X}_i(\pi^h(s_i'', S_j(\theta_j^*), s_{-i,j}))$.

First, assume that $\pi^* = \pi'$. Then $h^\neq = h_{t^\neq}''$. For the first item, the two plays are equivalent by construction of θ_j^* , so we are done. For the second item, if the two plays are distinct, then as any member of S_j that leads from h to h_{t^\neq}' also leads from h to h_{t^\neq}'' , necessarily both $\pi^h(s_i'', s_{-i})$ and $\pi^h(s_i'', S_j(\theta_j^*), s_{-i,j})$ lead from h to $h_{t^\neq}'' = h^\neq$, so $h^\neq \in \pi^h(s_i'', s_{-i}) \cap \pi^h(s_i'', S_j(\theta_j^*), s_{-i,j})$, from which the desired conclusion follows immediately.

Second, assume that $\pi^* = \pi''$; then the desired conclusions follow from a symmetric argument.

Thus in both cases, when the peers of i change from s_{-i} to $(S_j(\theta_j^*), s_{-i,j})$, $s_i'' \in S_i$ remains a profitable deviation for i from s_i' as claimed.

In both Case 1 and Case 2, $s_i'' \in S_i$ remains a profitable deviation for i from $s_i' = S_i(\theta_i)$ as claimed. But then repeating the argument, $s_i'' \in S_i$ remains such a profitable deviation for i from the convention when all of his peers play according to the convention, contradicting that (G^1, S) is an ex-post perfect implementation of the rule. \square




Fig. 5. *Indispensable conditions*. (a) *Rich*. Consider $N = \{1, 2\}$, $X_1 = X_2 = \{a, b, c\}$, $\Theta_1 = \Theta_2 = \{abc, cba\}$, and Fig. 5 (a) with the rule given by the *preferential* convention. (b) *Non-repeating*. Consider $N = \{1\}$, $X_1 = \{a, b, c\}$, the types are the strict rankings, and Fig. 5 (b) with the rule given by the *preferential* convention.

3.2. Logical relations

In this section, we consider both the necessity and the indispensability of our hypotheses for Theorem 1.

Necessity. A hypothesis is logically *necessary (for the conclusion)* if and only if it is an implication of the conclusion. For Theorem 1, *strategy-proofness* is necessary: an ex-post everywhere-dominant implementation is an ex-post Nash implementation, and it is well-known that for environments with private values, each rule with an ex-post Nash implementation is *strategy-proof* (see, for example, Bergemann and Morris, 2005). By definition, *f-compatibility* is necessary. It is easy to construct examples showing that the rest of the assumptions (the *richness*, *non-repeating*, *reaction-proof*, and *preferential* requirements) are not necessary. Indeed, we provide sufficient conditions for implementation, but do not describe all possible implementations.

Tightness. A hypothesis is *indispensable (for the proposition)* if and only if when the proposition is modified by dropping this hypothesis, it becomes false.¹¹ For Theorem 1, it is easy to construct examples showing that the *strategy-proof*, *preferential*, and *f-compatible* requirements are each indispensable. Moreover, menu mechanism G_4 (Fig. 2 (b)) establishes that *reaction-proofness* is indispensable. Finally, Fig. 5 establishes that the *richness* and *non-repeating* requirements are also indispensable. Altogether, all assumptions in Theorem 1 are indispensable, so we say the result is *tight*.

4. Menu mechanisms for group strategy-proof rules

4.1. Result

While Theorem 1 provides conditions that guarantee agents have strong robust incentives to adhere to the given convention, agents may also have strong robust incentives to play other conventions. For example, in Fig. 3 (a), there are four conventions, and each of them specifies an everywhere-dominant strategy equilibrium for each type profile, but only one is compatible

¹¹ To avoid confusion: for hypotheses H_i and conclusion C , consider a proposition with the format (H_1 and H_2 and ...) together imply C . If a hypothesis is necessary, then it is indispensable if and only if it is not implied by the other hypotheses, and therefore tightness (all hypotheses are indispensable) and logical independence of hypotheses are synonymous for axiomatic characterizations (where all hypotheses are necessary). In general, however, it is stronger to state that a hypothesis is indispensable than to state that it is not implied by the other hypotheses.

with the rule. In this section, we investigate when menu mechanisms provide double everywhere-dominant implementation, and thus are further robust to which equilibrium occurs.

As in the previous section, we begin by investigating menu mechanisms for which all actions are observable. In particular, we provide conditions under which a public menu mechanism provides a double subgame perfect implementation:

Proposition 2. *For each rich and strict environment, each group strategy-proof rule, each non-repeating public menu mechanism, and each preferential convention that is compatible with the rule, the public menu mechanism is both an ex-post perfect implementation of the rule via the convention and a full subgame perfect implementation of the rule.*

The formal proof is in Appendix F, and involves the same lemmas as the proof of Proposition 1. We sketch the arguments below:

Proof sketch. We take an arbitrary type profile $\theta \in \Theta$, an arbitrary subgame perfect equilibrium $s^* \in \text{SPE}(G, R(\theta))$, and define $s^\theta \equiv \mathbb{S}(\theta)$. By Proposition 1, we have that $s^\theta \in \text{SPE}(G, R(\theta))$ and $\mathcal{X}(s^\theta) = f(\theta)$. Thus to complete the proof, we need only show that $\mathcal{X}(s^\theta) = \mathcal{X}(s^*)$.

To do so, we first apply a result from the literature to deduce that our rule satisfies *non-bossiness* (Pápai, 2000 and Takamiya, 2001; see Section 4.2). We then proceed by backwards induction, iteratively showing that the two equilibria lead to the same outcome from histories earlier and earlier in the game tree until we conclude that they lead to the same equilibrium from the initial history. The inductive argument involves an agent i contemplating the choice between h^* (as prescribed by s^*) and h^θ (as prescribed by s^θ), where by induction the equilibrium outcomes agree for both h^* and h^θ . To prove that the outcomes agree whether (i) i selects h^* and then all play according to s^* ; or (ii) i selects h^θ and then all play according to s^θ ; we use the [Play Lemma](#), the [Continuation Lemma](#), *strictness*,¹² and *non-bossiness*. \square

Taken together, Proposition 1 and Proposition 2 show that public menu mechanisms provide an interesting alternative to direct mechanisms. Though public menu mechanisms only implement desired outcomes in terms of subgame perfect equilibrium—which means that if agents do not believe that their peers will conform to the convention, or at least pursue their own best interest, then undesirable outcomes can occur—they have the advantage that they do not require all information to be transmitted by the agents. By contrast, direct mechanisms do not require agents to have faith in the rationality of their peers, but they do require the agents to transmit all private information.

Our second theorem states that if we take the hypotheses for Proposition 2 and then replace the requirement that all actions are observable with *reaction-proofness*, then the resulting hypotheses guarantee double everywhere-dominant implementation:

Theorem 2. *For each rich and strict environment, each group strategy-proof rule, each non-repeating and reaction-proof menu mechanism, and each preferential convention that is compatible with the rule, the menu mechanism is both an ex-post everywhere-dominant implementation of the rule via the convention and a full everywhere-dominant implementation of the rule.*

¹² Alternatively, *strictness* can be removed if *non-bossiness* is strengthened to *non-bossiness in welfare/outcome* (Schummer and Velez, 2021), which can be easily verified with our proof.

Proof. By Theorem 1, we have an ex-post everywhere-dominant implementation via \mathbb{S} . Assume, by way of contradiction, we do not have a full everywhere-dominant implementation. Then there are $\theta \in \Theta$ and $s \in \mathbf{EDE}(G, R(\theta))$ such that $\mathcal{X}(\mathbb{S}(\theta)) \neq \mathcal{X}(s)$.

We claim that both $\mathbb{S}(\theta)$ and s are subgame perfect equilibria in the associated public menu mechanism game. Indeed, if not, then for one of these strategy profiles s^* , there is a history h where the player i has a profitable deviation s'_i . By the Play Lemma, there is a type $\theta'_i \in \Theta_i$ such that $\mathbb{S}_i(\theta'_i)$ adheres to $\pi^h(s'_i, s^*_{-i})$. But since \mathbb{S} is preferential, thus $\mathbb{S}_i(\theta'_i)$ is a strategy in the reaction-proof menu mechanism, so i has a profitable deviation from s^* at h , contradicting that s^* is an everywhere-dominant strategy equilibrium in the reaction-proof menu mechanism game.

Altogether, then, since $\mathbb{S}(\theta)$ and s are both subgame perfect equilibria in the associated public menu mechanism game, thus by Proposition 2, $\mathcal{X}(\mathbb{S}(\theta)) = \mathcal{X}(s)$, contradicting that $\mathcal{X}(\mathbb{S}(\theta)) \neq \mathcal{X}(s)$. \square

Taken together, Theorem 1 and Theorem 2 show that reaction-proof menu mechanisms provide a particularly interesting alternative to direct mechanisms: while neither requires agents to have faith in the rationality of their peers, the former moreover only requires a limited amount of information to be transmitted by the agents.

4.2. Logical relations

In this section, we consider both the necessity and the indispensability of our hypotheses for Theorem 2. In order to do so, we first separate group strategy-proofness from strategy-proofness using a result from the literature, which fruitfully also allows us to relate our result to the classic literature on Nash implementation. Recall the following two standard conditions for rules:

Definition. *Additional properties for rules.* Fix an environment and a rule f . For each $i \in N$, each $\theta_i \in \Theta_i$, and each $x \in X$, define the lower counter set of i for x given θ_i , $LCS_i(x|\theta_i) \equiv \{x' \in X | x R_i(\theta_i) x'\}$. We say that

- f is Maskin monotonic if and only if for each pair $\theta, \theta' \in \Theta$ such that for each $i \in N$, $LCS_i(f(\theta)|\theta_i) \subseteq LCS_i(f(\theta')|\theta'_i)$, we have $f(\theta) = f(\theta')$; and
- f is non-bossy if and only if for each $\theta \in \Theta$, each $i \in N$, and each $\theta'_i \in \Theta_i$, $f_i(\theta) = f_i(\theta'_i, \theta_{-i})$ implies $f(\theta) = f(\theta'_i, \theta_{-i})$.

Maskin monotonicity is the classic necessary condition for full Nash implementation (Maskin, 1999), while non-bossiness requires that if an agent does not change his own assignment when he changes his report, then he does not change anybody’s assignment (Satterthwaite and Sonnenschein, 1981).¹³ As established in the literature, there are strong logical relationships between these properties:

¹³ We remark that we use the original version of non-bossiness; see Thomson (2016) for a discussion of variants and their normative content.

Theorem PT (Pápai, 2000; Takamiya, 2001).¹⁴ For each rich and strict environment, and for each rule, the following are equivalent:

- the rule is group strategy-proof,
- the rule is strategy-proof and non-bossy, and
- the rule is Maskin monotonic.

This immediately yields two corollaries. To complete their proofs, first observe that for each rich and strict environment, and for each rule, there is a *non-repeating* and *reaction-proof* menu mechanism with a *preferential* convention that is compatible with the rule: simply consider any private menu mechanism that imitates the direct mechanism. It follows directly from this observation, the classic theorem that *Maskin monotonicity* is necessary for full Nash implementation (Maskin, 1999), Theorem 2, and Theorem PT, that:

Corollary 2.1. For each rich and strict environment, if a rule has a full Nash implementation, then there is a menu mechanism that is both an *ex-post everywhere-dominant implementation* of the rule via a *preferential convention* and a *full everywhere-dominant implementation* of the rule.

Corollary 2.2. For each rich and strict environment, each *strategy-proof* and *non-bossy* rule, each *non-repeating* and *reaction-proof* menu mechanism, and each *preferential convention* that is compatible with the rule, the menu mechanism is both an *ex-post everywhere-dominant implementation* of the rule via the convention and a *full everywhere-dominant implementation* of the rule.

Altogether, Corollary 2.1 provides a useful link between the classic literature on Nash implementation and everywhere-dominant implementation via menu mechanisms, while Corollary 2.2¹⁵ allows us to more clearly investigate the logical relationships of our hypotheses and conclusions:

Necessity. For Corollary 2.2, it is easy to see that for each assumption that appears in Theorem 1, necessity is the same across the two results. It is also easy to show that *strictness* is not necessary, and recent work suggests it may be particularly interesting to investigate the relaxation of our domain restrictions in future work.¹⁶ Finally, *non-bossiness* is not necessary (Fig. 6).

¹⁴ Both Pápai (2000) and Takamiya (2001) prove the equivalence of the first and second items, while Takamiya (2001) proves the equivalence of the first and third. Both papers involve models with additional structure, but these particular proofs apply directly to our model.

¹⁵ We remark that because Corollary 2.2 involves *non-bossiness* and a unique subgame perfect equilibrium outcome, it is conceptually related to Schummer and Velez (2021). In particular, Schummer and Velez (2021) consider *sequential equilibria* of imperfect information games where agents sequentially reveal preferences, and then outcomes are given by *strategy-proof* and *non-bossy* rules; they provide sufficient conditions on the prior guaranteeing that all sequential equilibria are welfare-equivalent to truth-telling, which is itself a sequential equilibrium. Interestingly, they observe that their work is also related to Marx and Swinkels (1997), who prove that a version of *non-bossiness* for normal form games guarantees that the order of elimination of weakly dominated strategies does not matter.

¹⁶ In particular, for (a discrete version of) the division problem with single-peaked preferences (Sprumont, 1991), *richness* and *strictness* are violated, but each *sequential allotment rule* (Barberà et al., 1997) has a double everywhere-dominant implementation through a menu mechanism with a *preferential convention* (Arrillaga et al., 2021). Incredibly, these implementations are moreover “everywhere obviously dominant.”




Fig. 6. Double everywhere-dominant implementation without non-bossiness. Consider $N = \{1, 2\}$, $X_1 = X_2 = \{\emptyset, 3, 6\}$, $\Theta_1 = \Theta_2 = \{\emptyset 36, 3\emptyset 6, 36\emptyset\}$, and Fig. 6 with the rule given by the preferential convention. Note that we can interpret this as an English auction with *strictness*: \emptyset is losing and paying nothing, 3 is winning and paying 3, 6 is winning and paying 6, and the admissible valuations are the non-integer reals.

Altogether, while Corollary 2.2 provides sufficient conditions for double everywhere-dominant implementation, there are interesting examples of these implementations that violate our conditions.

Tightness. As with Theorem 1, it is easy to construct examples showing the indispensability of all assumptions for Corollary 2.2 except *reaction-proofness*, *richness*, and the *non-repeating* requirement. Moreover, for these three assumptions, the examples we used for Theorem 1 all satisfy *strictness* and *non-bossiness*, and therefore still apply. Altogether, Corollary 2.2 is tight: we cannot modify the result to cover new double everywhere-dominant implementations by simply dropping some of our requirements.

5. Discussion

We have thus far highlighted how our model, assumptions, and results apply to our simple illustrative example—serial dictatorship for object allocation—but our results are far more general. Even for object allocation, our strongest conclusions apply to non-trivial menu mechanisms (that is, menu mechanisms that do not simply imitate the direct mechanism) for the full class of *group strategy-proof* and *efficient* rules. More generally, our results apply to matching with contracts, labor markets, auctions with unit demand and strictness, school choice, and two-candidate elections. We summarize the key implications of our results for these settings in Table 2, and provide details in Appendix G.

We emphasize that for our two-sided matching applications—doctor to hospital, worker to firm, buyer to seller, and student to school—our model fixes one side as strategic and the other as non-strategic, with the preferences of the non-strategic side reflected in the rule. Moreover, when the matching is many-to-one, the side whose members can only have a single match should be the strategic side. For example, for school choice, the agents are the students, while the preferences of the non-strategic schools and the deferred acceptance algorithm together define the rule.

Importantly, this modeling approach does not actually require us to assume that the other side is never strategic; it simply requires us to assume that the other side is no longer strategic when our analysis begins. For example, consider a school choice environment that is organized by a private menu mechanism based on deferred acceptance, where the students and schools each download an app on an electronic device. Suppose that the schools first strategically report their complete preference information, then the students play a private menu mechanism derived from the algorithm and these reports. From the perspective of the students, we can model this as a

Table 2

Applications of main results. In each row, we specify a class of environments. Moreover, for each of these environments, we specify an example: formally, a class of rules, each with an associated private menu mechanism and *preferential* convention (see Appendix G for details). For each of our hypotheses, a + means that the property is always satisfied while a - means that the property is sometimes violated. In the final column, XPED means that we always have ex-post everywhere-dominant implementation while 2ED means that we always have double everywhere-dominant implementation. Note that while we cannot directly apply either of our main results to labor markets or auctions due to violations of *richness*, we are able to apply Theorem 1 to matching with contracts and then prune the menu mechanism accordingly (see Appendix G for details).

Model	Example	Rich	Strict	Strategy-proof	Group strategy-proof	Non-repeating	Reaction-proof	Preferential	Compatible	Conclusion
Matching with contracts	Cumulative offers process	+	+	+	-	+	+	+	+	XPED
Labor markets	Salary adjustment process	-	+	+	-	+	+	+	+	XPED ^a
Auctions with unit demand and strictness	Crawford-Knoer auction (such as English auction)	-	+	+	-	+	+	+	+	XPED ^a
School choice	Deferred acceptance	+	+	+	-	+	+	+	+	XPED
Object allocation	Trading cycles (such as top trading cycles)	+	+	+	+	+	+	+	+	2ED
Two-candidate elections	Voting by committees	+	+	+	+	+	+	+	+	2ED

^a Conclusion does not follow from direct application of theorem; see caption.

game where (i) nature selects a school preference profile (with an associated rule) and a student preference profile, and then (ii) the students play a private menu mechanism for the selected rule. Even if we model this game using the thickest information sets compatible with perfect recall—so that a student does not even know which rule is being implemented!—everywhere-dominance of the *preferential* convention is preserved from our result for the case where the rule is known. Indeed, each student’s type-strategy only requires him to know his own preferences, and is everywhere-dominant across all private menu mechanisms that nature might select.

We can continue along this line of reasoning to see that everywhere-dominance holds in games with extremely limited information. Continuing with the above example, a fixed environment and rule are associated with multiple private menu mechanisms derived from the deferred acceptance algorithm that differ only in the order in which agents play; everywhere-dominance holds no matter how nature selects among these private menu mechanisms. Moreover, for a given student, everywhere-dominance holds across these private menu mechanisms not only for the current environment, but for all environments that include that student. It follows that everywhere-dominance holds for a student participating in a school choice environment organized by a deferred acceptance app even when he does not know (i) the number of other students in the school choice environment, let alone their identities or preferences, (ii) the order in which he and any other students may play, or (iii) the preferences of the schools. Analogous remarks apply for our other examples, such as online auctions where sellers strategically post their reservation prices and then the private menu mechanism is played by buyers who do not even know the number of other buyers.

Appendix A. Definition of mechanism

In this appendix, we formally define *mechanisms*:

Definition. Fix an environment. A *mechanism* is a tuple $G = (H, \succsim, \mathbb{P}, \mathcal{A}, \alpha, (\mathbb{I}_i)_{i \in N}, \mathcal{X})$ such that the following hold:

- H is the set of *histories* and \succsim is the partial order on H representing *precedence*. We require that (H, \succsim) is a meet-semilattice tree.¹⁷ For each $h \in H$, we let $\sigma(h)$ denote the set of immediate successors of h . A *play* is a maximal chain, which gives a complete description of a sequence of choices; we write π for a play and Π for the set of plays. A *terminal history* is a history with no successor; we write z for a terminal history and Z for the set of terminal histories.
- $\mathbb{P} : H \setminus Z \rightarrow N$ is the *player function*, which associates each non-terminal history with the agent who selects an action at that history. For each $i \in N$, we let $H_i \equiv \{h \in H \setminus Z \mid \mathbb{P}(h) = i\}$ denote the histories that belong to i .
- \mathcal{A} is the set of *actions* and $\alpha : \cup_{H \setminus Z} \sigma(h) \rightarrow \mathcal{A}$ is the *action function*, which at each non-terminal history h associates each immediate successor $h' \in \sigma(h)$ with the action taken to reach it. We require that at any non-terminal history, each available action determines a unique next history: for each $h \in H \setminus Z$ and each pair $h', h'' \in \sigma(h)$, $\alpha(h') \neq \alpha(h'')$. For each non-terminal history h , we let $\mathcal{A}(h) \equiv \{\alpha(h') \mid h' \in \sigma(h)\}$ denote the actions available at h . For each $i \in N$, each $h \in H_i$, and each $\pi \in \Pi$ such that $h \in \pi$, we let $\alpha^h(\pi)$ denote the action taken at h to remain on π . Similarly, for each $i \in N$, each $h \in H_i$, and each $h' \in H$ such that $h < h'$, we let $\alpha^h(h')$ denote the action taken at h to continue toward h' . It is straightforward to show that both $\alpha^h(\pi)$ and $\alpha^h(h')$ are well-defined.
- For each $i \in N$, \mathbb{I}_i is the *information partition for i* , which specifies the *information sets* partitioning H_i . We require that for each pair h, h' in the same information set \mathcal{I}_i , the same actions are available: $\mathcal{A}(h) = \mathcal{A}(h')$. We write $\mathcal{A}(\mathcal{I}_i)$ for the actions $\mathcal{A}(h)$ available at each history $h \in \mathcal{I}_i$. Across all histories in a given information set \mathcal{I}_i , i must behave the same way.
- $\mathcal{X} : \Pi \rightarrow X$ is the *outcome function*, which associates each play with an outcome.

For convenience, whenever we refer to a generic mechanism we implicitly assume all of this notation.

Appendix B. Details about privacy

In this appendix, we introduce a method for comparing the relative informativeness of two implementations, formally observe that menu mechanisms can improve upon the privacy of direct mechanisms, and compare our privacy notion to others in the literature.

¹⁷ These conditions guarantee that choices always determine a unique maximal chain, guarantee that there is a unique *initial history* which precedes all others, and allow an action to be viewed as selecting an immediate successor. For details about these order-theoretic concepts (*meet-semilattice*, *tree*, *successor*, *immediate successor*, and *maximal chain*) in the context of extensive game forms, see Alós-Ferrer and Ritzberger (2016) and Mackenzie (2020).

For intuition, suppose that an observer is (i) interested in the collective private information of the agents, and (ii) able to observe all actions taken by the agents.¹⁸ Moreover, suppose that the observer assumes that all agents follow the convention, so that after observing play π he infers that the type profile belongs to $\{\theta \in \Theta | \pi(\mathbb{S}(\theta)) = \pi\}$. In this case, the informativeness of an implementation (G, \mathbb{S}) is given by its partition of Θ , $\{\theta \in \Theta | \pi(\mathbb{S}(\theta)) = \pi\}_{\pi \in \Pi}$.¹⁹

When comparing two implementations, we simply compare the coarseness of their partitions, in which case plays can be interpreted as *messages* that are compared using the relative informativeness order of Segal (2007), and mechanisms can be interpreted as *experiments* that are compared using the Blackwell order (Blackwell, 1951)²⁰:

Definition. Relative informativeness. Fix an environment. Let (G, \mathbb{S}) and (G^*, \mathbb{S}^*) each be a mechanism with an associated convention, let π specify plays in G given strategy profiles in G , and let π^* specify plays in G^* given strategy profiles in G^* . We say that (G, \mathbb{S}) is at least as informative as (G^*, \mathbb{S}^*) if and only if for each $\theta \in \Theta$, we have

$$\{\theta' \in \Theta | \pi(\mathbb{S}(\theta')) = \pi(\mathbb{S}(\theta))\} \subseteq \{\theta' \in \Theta | \pi^*(\mathbb{S}^*(\theta')) = \pi^*(\mathbb{S}^*(\theta))\}.$$

In this case, we sometimes say that (G^*, \mathbb{S}^*) weakly improves upon the privacy of (G, \mathbb{S}) .

Using this notion, we can formally observe that menu mechanisms generally improve upon the privacy of direct mechanisms, and moreover can provide optimal privacy:

Observation. Fix an environment, a rule f , a menu mechanism G , and an f -compatible convention \mathbb{S} . Then

- $\{\theta \in \Theta | \pi(\mathbb{S}(\theta)) = \pi\}_{\pi \in \Pi}$ is weakly coarser than $\{\theta\}_{\theta \in \Theta}$, which is achieved by the direct mechanism with the honesty convention, and
- $\{\theta \in \Theta | \pi(\mathbb{S}(\theta)) = \pi\}_{\pi \in \Pi}$ is weakly finer than $\{f^{-1}(x)\}_{x \in f(X)}$.

Moreover, if we have an object allocation environment, f is serial dictatorship, G is an associated menu mechanism derived from the rule's algorithm, and \mathbb{S} is the *preferential* convention (see Appendix G), then $\{\theta \in \Theta | \pi(\mathbb{S}(\theta)) = \pi\}_{\pi \in \Pi} = \{f^{-1}(x)\}_{x \in f(X)}$. By Corollary G.4, (G, \mathbb{S}) is both an ex-post everywhere-dominant implementation and full everywhere-dominant implementation of f , and by the second item above, (G, \mathbb{S}) is minimally informative among all implementations of f .

Our observation that menu mechanisms generally offer better privacy than direct mechanisms formalizes simple intuition. For example, in an English auction, the winner only needs to reveal that his valuation is higher than the final price, but does not need to reveal more beyond that; indeed, this has been recognized as one of the practical merits of English auctions over direct

¹⁸ There are many natural variants that we could also consider; for example, the observer only observes the final outcome.

¹⁹ When \mathbb{S} is f -compatible, it is well-known that this partitions Θ into " f -monochromatic rectangles" (Kushilevitz and Nisan, 1997).

²⁰ A Blackwell experiment is a function that associates each unknown state with a probability distribution over signals. In our model, type profiles are states, plays are signals, and each state θ surely yields the signal $\pi(\mathbb{S}(\theta))$.

mechanisms (Rothkopf et al., 1990). We leave the general investigation of optimal privacy for future work.

To conclude this appendix, we compare our notion of privacy to several related notions from the literature. First, *communication complexity*, or the worst-case volume of bits that must be exchanged (Yao, 1979; Kushilevitz and Nisan, 1997), has been investigated for (i) calculating an optimal outcome when agents are honest, with or without the assistance of an omniscient oracle such as a Walrasian auctioneer (Segal, 2007), (ii) ex-post Nash implementation and Bayesian-Nash implementation (Fadel and Segal, 2009), and (iii) Nash implementation when each agent knows the entire type profile (Segal, 2010). At a high level, these papers establish that under some assumptions, providing incentives increases communication costs while using dynamic mechanisms reduces communication costs. Conceptually, we make a similar point—we argue that menu mechanisms use less information than direct mechanisms—but we do not formalize this with communication complexity because (i) from a privacy perspective, two bits may not be equal, and (ii) the worst-case analysis may fail to recognize the improvements of menu mechanisms in *most* cases.²¹ Second, we remark that other interesting notions of privacy that are less closely related to ours have also been considered. For example, in the Bayesian tradition, Eliaz et al. (2021) consider a notion of privacy based on the difference between the planner’s prior and posterior. As another example, agents whose preferences are sensitive to privacy concerns have been investigated, both in games and in implementation (Gradwohl and Smorodinsky, 2017; Gradwohl, 2018). Finally, Milgrom and Segal (2020) propose *unconditional winner privacy* in the context of auctions, and also discuss other privacy notions in the computer science literature where cryptography technology is taken into account.

Appendix C. Proofs of Cycle and Path Lemmas

In this appendix, we provide two lemmas about revealed preference theory (the [Cycle Lemma](#) and the [Path Lemma](#)) that are useful for studying public menu mechanisms.

In particular, some of our arguments involve fixing an agent’s choices at some histories—for example, so that the agent always chooses to remain on a given play—and then establishing that the convention specifies these choices for some type. When the environment is *rich* and the convention is *preferential*, this is equivalent to establishing that the choices are rationalizable by a strict preference relation; we therefore begin this appendix by abstracting from other features of our model to focus on menus and choices:

Definition. A *choice space* is a tuple $(A, \mathcal{M}, \mathcal{C})$, where

- A is a finite set of *alternatives*,
- $\mathcal{M} \subseteq 2^A$ is a nonempty collection of *menus*, and
- $\mathcal{C} : \mathcal{M} \rightarrow A$ is a *choice function* such that for each $M \in \mathcal{M}$, $\mathcal{C}(M) \in M$.

We let \mathcal{P}_A denote the set of strict preference relations on A .

For each pair $a, b \in A$, we say that a is *revealed preferred to* b by \mathcal{C} , written $a \triangleright^{\mathcal{C}} b$, if and only if there is $M \in \mathcal{M}$ such that $a, b \in M$ and $\mathcal{C}(M) = a$. A list $(a_1, a_2, \dots, a_k) \in \cup_{k \in \mathbb{N} \setminus \{1\}} A^k$ is

²¹ For example, the worst case in a deferred acceptance menu mechanism may involve each agent revealing his full preference ranking, though in most cases considerably less information is revealed.

a \mathcal{C} -path (from a_1 to a_k) if and only if for each $t \in \{1, 2, \dots, k - 1\}$, $a_t \succ^{\mathcal{C}} a_{t+1}$, and is moreover a \mathcal{C} -cycle (from a_1 to a_1) if and only if $a_k = a_1$.²²

Finally, we say that \mathcal{C} is *rationalizable* (by a strict preference relation) if and only if there is $P \in \mathcal{P}_A$ such that for each $M \in \mathcal{M}$, $\mathcal{C}(M) = \text{argmax}_P M$; in this case, we say that P *rationalizes* \mathcal{C} . We let $\mathcal{P}_{\mathcal{C}} \subseteq \mathcal{P}_A$ denote the set of members of \mathcal{P}_A that rationalize \mathcal{C} .

Our first lemma, which is a standard result, provides a necessary and sufficient condition for rationalizability. When applied to public menu mechanisms (under appropriate conditions), this lemma allows us to establish that certain choices made by an agent at some of his histories are specified by the convention for some type:

Cycle Lemma. *For each choice space $(A, \mathcal{M}, \mathcal{C})$, there are no \mathcal{C} -cycles if and only if \mathcal{C} is rationalizable.*

The **Cycle Lemma** follows easily from Szpilrajn’s Theorem, which states that every strict partial order can be extended to a linear order (Szpilrajn, 1930; see also Chambers and Echenique, 2016); we therefore omit its proof.

Our second lemma applies to public menu mechanisms (under appropriate conditions) when certain choices made by an agent are only specified by the convention for types that rank some assignment a over another assignment b . In this case, the lemma allows us to conclude that these choices yield a path from a to b :

Path Lemma. *Let $(A, \mathcal{M}, \mathcal{C})$ be a choice space such that \mathcal{C} is rationalizable. For each pair $a, b \in A$, there is a \mathcal{C} -path from a to b if and only if for each $P \in \mathcal{P}_{\mathcal{C}}$, we have $a P b$.*

Proof. Let $(A, \mathcal{M}, \mathcal{C})$ satisfy the hypotheses; we prove the parts in sequence.

[\Rightarrow] Assume there is a \mathcal{C} -path from a to b , $(a_1, a_2, \dots, a_k) \in \cup_{\kappa \in \mathbb{N} \setminus \{1\}} A^\kappa$, and let $P \in \mathcal{P}_{\mathcal{C}}$. It follows by definition that for each $t \in \{1, 2, \dots, k - 1\}$, we have $a_t P a_{t+1}$; thus by transitivity, $a = a_1 P a_k = b$, as desired.

[\Leftarrow] We prove the contrapositive; assume there is no \mathcal{C} -path from a to b . Because \mathcal{C} is rationalizable, there is $P \in \mathcal{P}_{\mathcal{C}}$. If $\{a, b\} \in \mathcal{M}$, then since there is no \mathcal{C} -path from a to b , thus we have $\mathcal{C}(\{a, b\}) = b$, so $b P a$ and we are done; thus let us assume $\{a, b\} \notin \mathcal{M}$. Define $\mathcal{M}^* \equiv \mathcal{M} \cup \{a, b\}$ and define $\mathcal{C}^* : \mathcal{M}^* \rightarrow A$ by (i) $\mathcal{C}^*(\{a, b\}) \equiv b$, and (ii) for each $M \in \mathcal{M}$, $\mathcal{C}^*(M) \equiv \mathcal{C}(M)$.

Assume, by way of contradiction, there is a \mathcal{C}^* -cycle $(a_1, a_2, \dots, a_k) \in \cup_{\kappa \in \mathbb{N} \setminus \{1\}} A^\kappa$. Since \mathcal{C} is rationalizable, thus by the **Cycle Lemma** there are no \mathcal{C} -cycles, so in particular (a_1, a_2, \dots, a_k) is not a \mathcal{C} -cycle; thus there is $t \in \{1, 2, \dots, k - 1\}$ such that $a_t = b$ and $a_{t+1} = a$. Using the members of (a_1, a_2, \dots, a_k) , we can construct a \mathcal{C}^* -cycle from a to a , $(a'_1, a'_2, \dots, a'_{k'})$, such that (i) for each $t \in \{2, 3, \dots, k' - 1\}$, $a'_t \neq a$, and (ii) $a'_{k'-1} = b$. For each $t \in \{1, 2, \dots, k' - 2\}$, $a'_t \succ^{\mathcal{C}^*} a'_{t+1}$ and $(a'_t, a'_{t+1}) \neq (b, a)$, so $a'_t \succ^{\mathcal{C}} a'_{t+1}$; but then $(a'_1, a'_2, \dots, a'_{k'-1})$ is a \mathcal{C} -path from a to b , contradicting that there is no \mathcal{C} -path from a to b .

²² To avoid confusion, it is standard in graph theory to require that the members of a path be distinct, and also to require that a cycle has $k \geq 3$ members whose first $k - 1$ members are distinct. These additional requirements have no impact on our arguments, so we omit them for simplicity.

Since there are no C^* -cycles, thus by the **Cycle Lemma** C^* is rationalizable, so there is $P^* \in \mathcal{P}_{C^*}$. By construction of C^* , $b P^* a$. Clearly $\mathcal{P}_{C^*} \subseteq \mathcal{P}_C$, so $P^* \in \mathcal{P}_C$, as desired. \square

Appendix D. Proofs of Play and Continuation Lemmas

In this appendix, we use the **Cycle Lemma** to provide two lemmas about public menu mechanisms (the **Play Lemma** and the **Continuation Lemma**), both of which are used in both proofs of our main results.

The first lemma provides conditions guaranteeing that for each play, each agent has a type for which the convention requires he always remain on the play:

Play Lemma. *Fix a rich environment, a non-repeating public menu mechanism, and a preferential convention. For each $\pi \in \Pi$ and each $i \in N$, there is $\theta_i \in \Theta_i$ such that for each $h \in H_i \cap \pi$, $[\mathbb{S}_i(\theta_i)](h) = \alpha^h(\pi)$.*

Proof. Let $\pi \in \Pi$ and let $i \in N$. To begin, we first show that i reveals no cycles along π . Indeed, define the choice space $(X_i, \mathcal{M}, \mathcal{C})$ by:

- $\mathcal{M} \equiv \{X_i(h) | h \in H_i \cap \pi\}$; and
- for each $h \in H_i \cap \pi$, $\mathcal{C}(X_i(h)) \equiv \alpha^h(\pi)$.

We claim there is no \mathcal{C} -cycle.

Assume, by way of contradiction, there is a \mathcal{C} -cycle $(a_1, a_2, \dots, a_k) \in \cup_{k \in \mathbb{N} \setminus \{1\}} X_i^k$. For each $t \in \{1, 2, \dots, k - 1\}$, $a_t \succ^{\mathcal{C}} a_{t+1}$, so there is $h_t \in H_i \cap \pi$ such that $a_t, a_{t+1} \in X_i(h_t)$ and $\mathcal{C}(X_i(h_t)) = a_t$. Since $\{h_t\} \subseteq \pi$, thus $\{h_t\}$ has a member $h_{<}$ which precedes the others. We consider the two cases for *non-repeating*:

CASE 1: i has non-repeating choices. Since $a_1 = a_k$, thus there is $t^* \in \{1, 2, \dots, k - 1\}$ such that $\mathcal{C}(X_i(h_{<})) = a_{t^*+1}$. Since $\mathcal{C}(X_i(h_{t^*})) = a_{t^*} \neq a_{t^*+1}$, thus $h_{<} \neq h_{t^*}$, so $h_{<} < h_{t^*}$. But then $a_{t^*+1} \in \{\alpha^{h_{<}}(h_{t^*})\} \cap X_i(h_{t^*})$, contradicting that i has non-repeating choices.

CASE 2: i has non-repeating rejections. Since $a_1 = a_k$, thus there is $t^* \in \{1, 2, \dots, k - 1\}$ such that $a_{t^*+1} \in X_i(h_{<}) \setminus \{\mathcal{C}(X_i(h_{<}))\}$. Since $\mathcal{C}(X_i(h_{t^*+1})) = a_{t^*+1}$, thus $h_{<} \neq h_{t^*+1}$, so $h_{<} < h_{t^*+1}$. But then $a_{t^*+1} \in (X_i(h_{<}) \setminus \{\alpha^{h_{<}}(h_{t^*+1})\}) \cap X_i(h_{t^*+1})$, contradicting that i has non-repeating rejections.

Since there are no \mathcal{C} -cycles, thus by the **Cycle Lemma** \mathcal{C} is rationalizable, so there is $P \in \mathcal{P}_C$. Since the environment is *rich*, thus there is $\theta_i \in \Theta_i$ such that $R_i(\theta_i) = P$. Since $R_i(\theta_i)$ is strict (and therefore requires no tie-breaking), thus for each $h \in H_i \cap \pi$, we have

$$\begin{aligned}
 [\mathbb{S}_i(\theta_i)](h) &= \operatorname{argmax}_{R_i(\theta_i)} X_i(h) && \text{as } \mathbb{S} \text{ is preferential and } R_i(\theta_i) \text{ is strict} \\
 &= \operatorname{argmax}_P X_i(h) && \text{by construction of } \theta_i \\
 &= \mathcal{C}(X_i(h)) && \text{as } P \in \mathcal{P}_C \\
 &= \alpha^h(\pi) && \text{by construction of } \mathcal{C},
 \end{aligned}$$

as desired. \square

The second lemma involves continuation strategies:

Definition. Fix an environment and a public menu mechanism. For each $h \in H$, each $i \in N$, and each $s_i \in S_i$, define $s_i \upharpoonright_h$ to be restriction of s_i to $\{h' \in H_i \mid h' \succsim h\}$; we refer to this as a *continuation strategy (at h for i)*.

In particular, for each history, each agent, and each type, the convention specifies a continuation strategy. The second lemma provides conditions guaranteeing that there is a type that requires the agent to (i) continue toward the given history whenever possible, and (ii) conform to the specified continuation strategy whenever possible:

Continuation Lemma. Fix a rich environment, a non-repeating public menu mechanism, and a preferential convention. For each $h \in H$, each $i \in N$, and each $\theta_i \in \Theta_i$, there is $\theta_i^* \in \Theta_i$ such that

- (i) for each $h' \in H_i$ such that $h' \prec h$, we have $[\mathbb{S}_i(\theta_i^*)](h') = \alpha^{h'}(h)$; and
- (ii) $\mathbb{S}_i(\theta_i^*) \upharpoonright_h = \mathbb{S}_i(\theta_i) \upharpoonright_h$.

Proof. Let $h \in H$, let $i \in N$, and let $\theta_i \in \Theta_i$. If $\{h' \in H_i \mid h' \prec h\} = \emptyset$ then we are done; thus let us assume $\{h' \in H_i \mid h' \prec h\} \neq \emptyset$. Define $H_{\succsim} \subseteq H_i$ and the choice space $(X_i, \mathcal{M}_{\succsim}, \mathcal{C}_{\succsim})$ by:

- $H_{\succsim} \equiv \{h' \in H_i \mid h' \succsim h\}$;
- $\mathcal{M}_{\succsim} \equiv \{X_i(h')\}_{h' \in H_{\succsim}}$; and
- for each $h' \in H_{\succsim}$, $\mathcal{C}_{\succsim}(X_i(h')) \equiv [\mathbb{S}_i(\theta_i)](h')$.

Define $H_{\prec} \subseteq H_i$ and the choice space $(X_i, \mathcal{M}, \mathcal{C})$ by:

- $H_{\prec} \equiv \{h' \in H_i \mid h' \prec h\}$;
- $\mathcal{M} \equiv \{X_i(h')\}_{h' \in H_{\prec} \cup H_{\succsim}}$; and
- for each $h' \in H_{\prec}$, $\mathcal{C}(X_i(h')) \equiv \alpha^{h'}(h)$; and for each $h' \in H_{\succsim}$, $\mathcal{C}(X_i(h')) \equiv [\mathbb{S}_i(\theta_i)](h')$.

Observe that for each $h' \in H_{\succsim}$, $\mathcal{C}(X_i(h')) = \mathcal{C}_{\succsim}(X_i(h'))$.

First, we claim there are no \mathcal{C}_{\succsim} -cycles. Indeed, since \mathbb{S} is *preferential*, thus there is $P_{\theta_i} \in \mathcal{P}_i$ such that for each $h' \in H_i$, $[\mathbb{S}_i(\theta_i)](h') = \text{argmax}_{P_{\theta_i}} X_i(h')$; this P_{θ_i} is easily constructed from $R_i(\theta_i)$ and the tie-breaker $\tau_i(\theta_i)$. Since $P_{\theta_i} \in \mathcal{P}_{\mathcal{C}_{\succsim}}$, thus \mathcal{C}_{\succsim} is rationalizable, so by the **Cycle Lemma** there are no \mathcal{C}_{\succsim} -cycles.

Second, we claim there are no \mathcal{C} -cycles. Indeed, assume, by way of contradiction, there is a \mathcal{C} -cycle $(a_1, a_2, \dots, a_k) \in \cup_{k \in \mathbb{N} \setminus \{1\}} X_i^k$. For each $t \in \{1, 2, \dots, k - 1\}$, $a_t \triangleright^{\mathcal{C}} a_{t+1}$, so there is $h_t \in H_{\prec} \cup H_{\succsim}$ such that $a_t, a_{t+1} \in X_i(h_t)$ and $\mathcal{C}(X_i(h_t)) = a_t$. Because there are no \mathcal{C}_{\succsim} -cycles, thus there is $t^* \in \{1, 2, \dots, k - 1\}$ such that $h_{t^*} \in H_{\prec}$; it follows that $\{h_t\}$ has a member h_{\prec} which precedes the others, and that for each $h' \in \{h_t\} \setminus \{h_{\prec}\}$, $\mathcal{C}(X_i(h_{\prec})) = \alpha^{h_{\prec}}(h')$. From here, the two-case argument about *non-repeating* from the proof of the **Play Lemma** establishes the contradiction.

Since there are no \mathcal{C} -cycles, thus by the **Cycle Lemma** \mathcal{C} is rationalizable, so there is $P \in \mathcal{P}_{\mathcal{C}}$. Since the environment is *rich*, thus there is $\theta_i^* \in \Theta_i$ such that $R_i(\theta_i^*) = P$. For each $h' \in H_{\prec} \cup H_{\succsim}$,

$$\begin{aligned}
 [\mathbb{S}_i(\theta_i^*)](h') &= \operatorname{argmax}_{R_i(\theta_i^*)} X_i(h') && \text{as } \mathbb{S} \text{ is preferential and } R_i(\theta_i^*) \text{ is strict} \\
 &= \operatorname{argmax}_P X_i(h') && \text{by construction of } \theta_i^* \\
 &= \mathcal{C}(X_i(h')) && \text{as } P \in \mathcal{P}_{\mathcal{C}}.
 \end{aligned}$$

Thus by construction of \mathcal{C} ,

(i) for each $h' \in H_i$ such that $h' \prec h$, we have $[\mathbb{S}_i(\theta_i^*)](h') = \alpha^{h'}(h)$; and

(ii) for each $h' \in H_{\succsim}$, $[\mathbb{S}_i(\theta_i^*)](h') = [\mathbb{S}_i(\theta_i)](h')$, so $\mathbb{S}_i(\theta_i^*) \upharpoonright_h = \mathbb{S}_i(\theta_i) \upharpoonright_h$.

Since $h \in H$, $i \in N$, and $\theta_i \in \Theta_i$ were arbitrary, we are done. \square

Appendix E. Proof of Proposition 1

In this appendix, we prove Proposition 1, then observe that the result does not hold if we weaken *non-repeating* in a natural way or weaken *preferential* in a natural way (Fig. 7).

Proposition 1. *For each rich environment, each strategy-proof rule, each non-repeating public menu mechanism, and each preferential convention that is compatible with the rule, the public menu mechanism is an ex-post perfect implementation of the rule via the convention.*

Proof. Let $\theta \in \Theta$, let $i \in N$, let $h \in H_i$, and let $s_i \in S_i$. Define $a, b \in X_i$ by

$$\begin{aligned}
 a &\equiv \mathcal{X}_i^h(\mathbb{S}(\theta)), \text{ and} \\
 b &\equiv \mathcal{X}_i^h(s_i, \mathbb{S}_{-i}(\theta_{-i})).
 \end{aligned}$$

We want to show $a R_i(\theta_i) b$. If $a = b$ then we are done, so assume $a \neq b$.

Define $\pi_a \equiv \pi^h(\mathbb{S}(\theta))$, define $\pi_b \equiv \pi^h(s_i, \mathbb{S}_{-i}(\theta_{-i}))$, and define $H_{\prec} \equiv \{h' \in H_i \mid h' \prec h\}$. Define the choice space $(X_i, \mathcal{M}_a, \mathcal{C}_a)$ by:

- $\mathcal{M}_a \equiv \{X_i(h')\}_{h' \in H_i \cap \pi_a}$; and
- for each $h' \in H_i \cap \pi_a$, $\mathcal{C}_a(X_i(h')) \equiv \alpha^{h'}(\pi_a)$.

Define the choice space $(X_i, \mathcal{M}_b, \mathcal{C}_b)$ by:

- $\mathcal{M}_b \equiv \{X_i(h')\}_{h' \in H_i \cap \pi_b}$; and
- for each $h' \in H_i \cap \pi_b$, $\mathcal{C}_b(X_i(h')) \equiv \alpha^{h'}(\pi_b)$.

By the argument used in the proof of the **Play Lemma**, both \mathcal{C}_a and \mathcal{C}_b are rationalizable.

Let $P_a \in \mathcal{P}_{\mathcal{C}_a}$ and let $P_b \in \mathcal{P}_{\mathcal{C}_b}$. Since the environment is *rich*, thus there are $\theta_i^a, \theta_i^b \in \Theta_i$ such that $R_i(\theta_i^a) = P_a$ and $R_i(\theta_i^b) = P_b$. By the **Continuation Lemma**, for each $j \in N \setminus \{i\}$, there is $\theta_j^* \in \Theta_j$ such that

- (i) for each $h' \in H_j$ such that $h' \prec h$, we have $[\mathbb{S}_j(\theta_j^*)](h') = \alpha^{h'}(h)$; and
- (ii) $\mathbb{S}_j(\theta_j^*) \upharpoonright_h = \mathbb{S}_j(\theta_j) \upharpoonright_h$.

Since \mathbb{S} is *preferential*, thus by construction, $\pi_a = \pi(\mathbb{S}(\theta_i^a, \theta_{-i}^*))$ and $\pi_b = \pi(\mathbb{S}(\theta_i^b, \theta_{-i}^*))$, so by *f-compatibility*, $a = \mathcal{X}_i(\mathbb{S}(\theta_i^a, \theta_{-i}^*)) = f_i(\theta_i^a, \theta_{-i}^*)$ and $b = \mathcal{X}_i(\mathbb{S}(\theta_i^b, \theta_{-i}^*)) = f_i(\theta_i^b, \theta_{-i}^*)$. By *strategy-proofness*,

$$a = f_i(\theta_i^a, \theta_{-i}^*)$$

$$R_i(\theta_i^a) f_i(\theta_i^b, \theta_{-i}^*) = b,$$

so $a P_a b$. By a symmetric argument, $b P_b a$. Since $P_a \in \mathcal{P}_{C_a}$ and $P_b \in \mathcal{P}_{C_b}$ were arbitrary, thus by the **Path Lemma**, there are (i) a C_a -path from a to b , $(a_1, a_2, \dots, a_k) \in \cup_{\kappa \in \mathbb{N} \setminus \{1\}} X_i^\kappa$, and (ii) a C_b -path from b to a , $(b_1, b_2, \dots, b_{k'}) \in \cup_{\kappa \in \mathbb{N} \setminus \{1\}} X_i^\kappa$. For each $t \in \{1, 2, \dots, k - 1\}$, $a_t \succ^{C_a} a_{t+1}$, so there is $h_t^a \in H_i \cap \pi_a$ such that $a_t, a_{t+1} \in X_i(h_t^a)$ and $C_a(X_i(h_t^a)) = a_t$. Similarly, for each $t \in \{1, 2, \dots, k' - 1\}$, $b_t \succ^{C_b} b_{t+1}$, so there is $h_t^b \in H_i \cap \pi_b$ such that $b_t, b_{t+1} \in X_i(h_t^b)$ and $C_b(X_i(h_t^b)) = b_t$.

Assume, by way of contradiction, $\{h_t^a\} \cap H_{<} \neq \emptyset$. Since $\{h_t^a\} \subseteq \pi_a$, thus it has a member $h_{<}^a$ which precedes the others. We consider the two cases for *non-repeating*:

CASE 1: i has non-repeating choices. It must be that $h_{<}^a = h_1^a$; else there is $t \in \{1, 2, \dots, k - 1\}$ such that (i) $h_{<}^a < h_t^a$, and (ii) the assignment that is chosen by C_a at $X_i(h_{<}^a)$ is rejected by C_a at $X_i(h_t^a)$, contradicting that i has non-repeating choices. By the same argument, h_1^b precedes the others in $\{h_t^b\}$. Since $\{h_t^a\} \cap H_{<} \neq \emptyset$, thus $h_1^a \in H_{<}$, so $h_1^a \in \pi_b$ and $a = \alpha^{h_1^a}(\pi_b)$. Since $a \neq \alpha^{h_{k'-1}^b}(\pi_b)$, thus $h_1^a \neq h_{k'-1}^b$, so either $h_1^a < h_{k'-1}^b$ or $h_{k'-1}^b < h_1^a$. Since i has non-repeating choices, thus $h_{k'-1}^b < h_1^a$, so $h_1^b \preceq h_{k'-1}^b < h_1^a$. But then $\{h_t^b\} \cap H_{<} \neq \emptyset$, so by a symmetric argument, $h_1^a < h_1^b$, contradicting that $h_1^b < h_1^a$.

CASE 2: i has non-repeating rejections. It must be that $h_{<}^a = h_{k-1}^a$; else there is $t \in \{1, 2, \dots, k - 1\}$ such that (i) $h_{<}^a < h_t^a$, and (ii) an assignment that is rejected by C_a at $X_i(h_{<}^a)$ is chosen by C_a at $X_i(h_t^a)$, contradicting that i has non-repeating rejections. By the same argument, $h_{k'-1}^b$ precedes the others in $\{h_t^b\}$. Since $\{h_t^a\} \cap H_{<} \neq \emptyset$, thus $h_{k-1}^a \in H_{<}$, so $h_{k-1}^a \in \pi_b$ and $b \neq \alpha^{h_{k-1}^a}(\pi_b)$. Since $b = \alpha^{h_1^b}(\pi_b)$, thus $h_{k-1}^a \neq h_1^b$, so either $h_{k-1}^a < h_1^b$ or $h_1^b < h_{k-1}^a$. Since i has non-repeating rejections, thus $h_1^b < h_{k-1}^a$, so $h_{k'-1}^b \preceq h_1^b < h_{k-1}^a$. But then $\{h_t^b\} \cap H_{<} \neq \emptyset$, so by a symmetric argument, $h_{k-1}^a < h_{k'-1}^b$, contradicting that $h_{k'-1}^b < h_{k-1}^a$.

Thus $\{h_t^a\} \cap H_{<} = \emptyset$. Since \mathbb{S} is *preferential*, thus there is $\tau_i(\theta_i) \in \mathcal{P}_i$ such that for each $h' \in H_i$,

$$[\mathbb{S}_i(\theta_i)](h') = \operatorname{argmax}_{\tau_i(\theta_i)} [\operatorname{argmax}_{R_i(\theta_i)} X_i(h')].$$

Define $P_{\theta_i} \in \mathcal{P}_i$ by $a_1 P_{\theta_i} a_2$ if and only if (i) $a_1 P_i(\theta_i) a_2$, or (ii) $a_1 I_i(\theta_i) a_2$ and $a_1 \tau_i(\theta_i) a_2$. Define the choice space $(X_i, \mathcal{M}, \mathcal{C})$ by:

- $\mathcal{M} \equiv \{X_i(h')\}_{h' \in (H_i \cap \pi_a) \setminus H_{<}}$; and
- for each $h' \in (H_i \cap \pi_a) \setminus H_{<}$, $\mathcal{C}(X_i(h')) \equiv \alpha^{h'}(\pi_a)$.

Since $\pi_a = \pi^h(\mathbb{S}(\theta))$, thus $P_{\theta_i} \in \mathcal{P}_{\mathcal{C}}$, so \mathcal{C} is rationalizable. Since $\{h_t^a\} \cap H_{<} = \emptyset$, thus $\{h_t^a\} \subseteq (H_i \cap \pi_a) \setminus H_{<}$, so there is a \mathcal{C} -path from a to b . Since $P_{\theta_i} \in \mathcal{P}_{\mathcal{C}}$, thus by the **Path Lemma**, $a P_{\theta_i} b$. Altogether, then, by construction we have $a R_i(\theta_i) b$, as desired.

To conclude, since $i \in N$, $h \in H_i$, and $s_i \in S_i$ were arbitrary, thus $\mathbb{S}(\theta) \in \mathbf{SPE}(G, R(\theta))$. Since $\theta \in \Theta$ was arbitrary, thus \mathbb{S} satisfies *ex-post perfection*. By *f-compatibility*, G is an *ex-post perfect implementation* of f via \mathbb{S} . \square

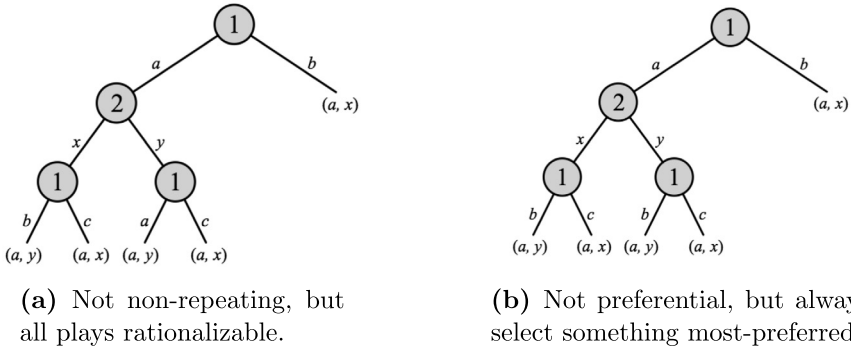


Fig. 7. Weaker hypotheses that do not suffice for Proposition 1. For both examples, consider $N = \{1, 2\}$, $X_1 = \{a, b, c\}$, and $X_2 = \{x, y\}$. (a) Suppose the types are the strict rankings, and consider Fig. 7 (a) with the rule given by the preferential convention. This only violates non-repeating and is not an ex-post perfect implementation, despite satisfying the weaker requirement that the choices along each play can be rationalized by some strict preference profile. (b) Suppose Θ_1 is the strict rankings together with complete indifference and Θ_2 is the strict rankings, and consider Fig. 7 (b) with the rule given by the following convention: (i) each agent always selects a most-preferred assignment, and (ii) when 1 has complete indifference, he selects b at the initial history, selects b after 2 selects x , and selects c after 2 selects y . This only violates preferential and is not an ex-post perfect implementation, despite satisfying the weaker requirement that the convention always asks each agent to always select a most-preferred assignment.

To conclude this appendix, we observe that Proposition 1 does not hold if (i) we weaken non-repeating by simply requiring that the choices along each play can be rationalized by some strict preference profile, or (ii) we weaken preferential by simply requiring that the convention always asks each agent to select a most-preferred assignment without necessarily breaking ties consistently (Fig. 7).

Appendix F. Proof of Proposition 2

In this appendix, we prove Proposition 2.

Proposition 2. For each rich and strict environment, each group strategy-proof rule, each non-repeating public menu mechanism, and each preferential convention that is compatible with the rule, the public menu mechanism is both an ex-post perfect implementation of the rule via the convention and a full subgame perfect implementation of the rule.

Proof. Remove each history where the player has a single action, let G denote the resulting mechanism, and let \mathbb{S} denote the associated restricted convention. Then (G, \mathbb{S}) satisfies the hypotheses. Moreover, since (i) for each $i \in N$, X_i is finite, and (ii) G is non-repeating, thus each play is finite. We prove that the desired conclusion holds for (G, \mathbb{S}) , from which the proposition clearly follows.

By Theorem PT, f is strategy-proof and non-bossy. By Proposition 1, G is an ex-post perfect implementation of the rule.

Let $\theta \in \Theta$, let $s^* \in \text{SPE}(G, R(\theta))$, and define $s^\theta \equiv \mathbb{S}(\theta)$. By Proposition 1, $s^\theta \in \text{SPE}(G, R(\theta))$ and $\mathcal{X}(s^\theta) = f(\theta)$. To prove that $\mathcal{X}(s^*) = \mathcal{X}(s^\theta)$, we use a version of backwards induction, proceeding by induction on history length. In particular, for each $h \in H$, define the length of h , $\ell(h) \equiv \max_{\pi \in \Pi} |\pi \cap \{h' \in H \mid h' \succsim h\}|$; this is the maximum cardinality of a play in the subgame that starts at h .

For the base step, for each $h \in H$ such that $\ell(h) = 1$, we have $h \in Z$, so $\mathcal{X}^h(s^*) = \mathcal{X}^h(s^\theta)$. For the inductive step, assume $L \in \mathbb{N}$ is such that for each $h \in H$ such that $\ell(h) \leq L$, we have $\mathcal{X}^h(s^*) = \mathcal{X}^h(s^\theta)$; and let $h \in H$ such that $\ell(h) = L + 1$. Define $i \equiv \mathbb{P}(h)$, let h^* be the immediate successor of h identified by s^* , and let h^θ be the immediate successor of h identified by s^θ . Since $\ell(h^*) \leq L$ and $\ell(h^\theta) \leq L$, thus

$$\begin{aligned} \mathcal{X}^{h^*}(s^\theta) &= \mathcal{X}^{h^*}(s^*) && \text{by the inductive hypothesis as } \ell(h^*) \leq L \\ &R_i(\theta_i) \mathcal{X}^{h^\theta}(s^*) && \text{as } s^* \in \mathbf{SPE}(G, R(\theta)) \\ &= \mathcal{X}^{h^\theta}(s^\theta) && \text{by the inductive hypothesis as } \ell(h^\theta) \leq L. \end{aligned}$$

Since $s^\theta \in \mathbf{SPE}(G, R(\theta))$, thus $\mathcal{X}^{h^\theta}(s^\theta) R_i(\theta_i) \mathcal{X}^{h^*}(s^\theta)$. Altogether, we have $\mathcal{X}^{h^*}(s^\theta) I_i(\theta_i) \mathcal{X}^{h^\theta}(s^\theta)$.

To conclude the inductive step, we claim $\mathcal{X}^h(s^*) = \mathcal{X}^h(s^\theta)$. Indeed, define $\pi^* \equiv \pi^{h^*}(s^\theta)$ and $\pi^\theta \equiv \pi^{h^\theta}(s^\theta)$. By the **Play Lemma**, there is $\theta_i^* \in \Theta_i$ such that for each $h' \in H_i \cap \pi^*$, $[\mathbb{S}_i(\theta_i^*)](h') = \alpha^{h'}(\pi^*)$. By the **Continuation Lemma**, for each $j \in N$, there is $\theta_j^h \in \Theta_j$ such that

- (i) for each $h' \in H_j$ such that $h' \prec h$, we have $[\mathbb{S}_j(\theta_j^h)](h') = \alpha^{h'}(h)$; and
- (ii) $\mathbb{S}_j(\theta_j^h) \upharpoonright_h = \mathbb{S}_j(\theta_j) \upharpoonright_h$.

By construction, $\pi^* = \pi(\mathbb{S}(\theta_i^*, \theta_{-i}^h))$ and $\pi^\theta = \pi(\mathbb{S}(\theta^h))$, so by *f-compatibility*, $\mathcal{X}^{h^*}(s^\theta) = \mathcal{X}(\mathbb{S}(\theta_i^*, \theta_{-i}^h)) = f(\theta_i^*, \theta_{-i}^h)$ and $\mathcal{X}^{h^\theta}(s^\theta) = \mathcal{X}(\mathbb{S}(\theta^h)) = f(\theta^h)$. Then $f(\theta_i^*, \theta_{-i}^h) I_i(\theta_i) f(\theta^h)$, so by *strictness*, $f_i(\theta_i^*, \theta_{-i}^h) = f_i(\theta^h)$. By *non-bossiness*, $f(\theta_i^*, \theta_{-i}^h) = f(\theta^h)$, so $\mathcal{X}^{h^*}(s^\theta) = \mathcal{X}^{h^\theta}(s^\theta)$. Altogether, then,

$$\begin{aligned} \mathcal{X}^h(s^*) &= \mathcal{X}^{h^*}(s^*) && \text{by definition of } h^* \\ &= \mathcal{X}^{h^*}(s^\theta) && \text{by the inductive hypothesis as } \ell(h^*) \leq L \\ &= \mathcal{X}^{h^\theta}(s^\theta) && \text{by the above argument} \\ &= \mathcal{X}^h(s^\theta) && \text{by definition of } h^\theta, \end{aligned}$$

as desired.

By induction, for each $L \in \mathbb{N}$ and each $h \in H$ such that $\ell(h) = L$, we have $\mathcal{X}^h(s^*) = \mathcal{X}^h(s^\theta)$. Since each play is finite, thus the initial history h_\emptyset is such that $\ell(h_\emptyset) \in \mathbb{N}$, so $\mathcal{X}(s^*) = \mathcal{X}(s^\theta) = f(\theta)$. Since $s^\theta \in \mathbf{SPE}(G, R(\theta))$, since $s^* \in \mathbf{SPE}(G, R(\theta))$ was arbitrary, and since $\theta \in \Theta$ was arbitrary, we are done. \square

Appendix G. Details about applications

In this appendix, we provide formal details about the implications of our general results (Proposition 1, Proposition 2, Theorem 1, and Theorem 2) for familiar settings. In particular, for each application, we specify the additional structure that we impose on an environment to fit the application, specify the class of rules (each with an associated menu mechanism and convention) that we consider for each of these environments, formally state the associated corollaries, and briefly highlight relevant ideas from the literature. To simplify the discussion, we focus on private menu mechanisms as our examples of *reaction-proof* mechanisms throughout this appendix. The corollaries about everywhere-dominance are summarized in the main text in Table 2.

G.1. Matching with contracts

First, we consider matching with contracts (Hatfield and Milgrom, 2005). Suppose that there is a finite set of doctors and a finite set of hospitals, including a special *outside option*. Moreover, there is a finite set of contracts—each of which has an associated doctor, an associated hospital, and some terms—such that each doctor has a unique contract with the outside option. Let us say that an environment is a *matching with contracts environment* if and only if (i) the agents are the doctors; (ii) for each agent, the possible assignments are his contracts; (iii) an outcome is any assignment profile; and (iv) for each agent, the possible types are the strict rankings of his assignments.

For each hospital, a choice correspondence associates each set of its contracts with a subset of those contracts, and each profile of choice correspondences \mathcal{C} determines a (*doctor-proposing*) *cumulative offers process rule*, $f^{COP|\mathcal{C}}$ (Hatfield and Milgrom, 2005). This rule associates each type profile with the outcome of the cumulative offers process, which is described in terms of doctors iteratively proposing contracts to hospitals, and it is straightforward to adapt this process into a menu mechanism. Moreover, for each \mathcal{C} such that (i) the outside option always chooses all contracts, and (ii) each hospital’s choice correspondence satisfies ‘observable substitutability,’ ‘observable size monotonicity,’ and ‘non-manipulability via contractual terms,’ the rule $f^{COP|\mathcal{C}}$ is *strategy-proof* (Hatfield et al., 2021). It is easy to verify that our other conditions are satisfied, and thus we have:

Corollary G.1. *Fix a matching with contracts environment and a profile of choice correspondences \mathcal{C} such that (i) the outside option always chooses all contracts, and (ii) each hospital’s choice correspondence satisfies ‘observable substitutability,’ ‘observable size monotonicity,’ and ‘non-manipulability via contractual terms.’ Let G^{public} be a public menu mechanism for the associated cumulative offers process, and let $G^{private}$ be the associated private menu mechanism. Then*

- G^{public} is an ex-post perfect implementation of $f^{COP|\mathcal{C}}$ via the preferential convention; and
- $G^{private}$ is an ex-post everywhere-dominant implementation of $f^{COP|\mathcal{C}}$ via the preferential convention.

In general, cumulative offers process rules are not *group strategy-proof*, which follows from Kojima (2010). We remark that while the literature has identified conditions that guarantee a cumulative offers process rule is *weakly group strategy-proof* (Hatfield and Kojima, 2009), this is not enough for us to apply our results on full implementation.

G.2. Labor markets and auctions

Second, we consider labor markets with salaries, which can be viewed as matching with contracts environments that are modified to have restricted preference domains (Crawford and Knoer, 1981; Kelso and Crawford, 1982). Suppose that there is a finite set of workers and a finite set of firms, including a special *outside option*. Moreover, there is a finite set of salaries, and the set of contracts is the product of workers and firms and salaries except that only the lowest salary is ever available for the outside option. Let us say that an environment is a *labor market environment* if and only if (i) the agents are the workers; (ii) for each agent, the possible assignments are his contracts; (iii) an outcome is any assignment profile; and (iv) for each agent, the possible

types are the strict rankings of his assignments such that a higher salary is always preferred to a lower salary for a given firm.

Each profile of choice correspondences for firms determines a (*worker-proposing*) *salary adjustment process rule*, $f^{SAP|C}$ (Crawford and Knoer, 1981; Kelso and Crawford, 1982), which in fact coincides with the restriction of $f^{COP|C}$ to monotonic types. As with our last example, the salary adjustment process naturally yields an associated menu mechanism, but now we cannot directly apply our result because labor market environments violate *richness*. Nevertheless, we prove that our results still have implications for this setting because we can enrich the type space, apply our result for matching with contracts, and then prune the menu mechanism from the richer environment:

Corollary G.2. *Fix a labor market environment and a profile of choice correspondences C such that (i) the outside option takes all contracts, and (ii) each firm’s choice correspondence satisfies ‘observable substitutability,’ ‘observable size monotonicity,’ and ‘non-manipulability via contractual terms.’ Let G^{public} be a public menu mechanism for the associated salary adjustment process, and let $G^{private}$ be the associated private menu mechanism. Then*

- G^{public} is an ex-post perfect implementation of $f^{SAP|C}$ via the preferential convention; and
- $G^{private}$ is an ex-post everywhere-dominant implementation of $f^{SAP|C}$ via the preferential convention.

Proof. We prove both parts with the same arguments. First, modify the labor market environment so that each agent can have all strict rankings and apply Corollary G.1 to obtain the desired implementation. Second, remove all types that violate monotonicity; we still have the desired implementation. Finally, modify the mechanism by pruning off all histories that are unused by the convention; we still have an implementation, and the result is the desired menu mechanism. □

An important special case is when each firm hires at most one worker, first prioritizes having an employee, then prioritizes minimizing the sole employee’s salary, and finally considers the employee’s identity. In this case, we can reinterpret workers as buyers with unit demand, reinterpret firms as sellers with unit supply, and reinterpret higher salaries as lower prices, resulting in a buyer-proposing Crawford-Knoer auction with *strictness* (Crawford and Knoer, 1981; Demange et al., 1986). It is standard in discrete settings to assume that each buyer is necessarily indifferent between exiting and receiving the object at one of the prices, but in our variant model we assume this is never the case; our *strictness* assumption holds if each admissible valuation is generic in the sense that it belongs to an open ball whose members all induce the same ranking of assignments. Note that under the labor market interpretation, the implementation is a descending salary procedure, while under the auction interpretation, it is an ascending price procedure.

To conclude, we consider the special case where there is a single seller with one object. In the standard model, it is well-known that the Vickrey rule is *strategy-proof* (Vickrey, 1961), and that it is implemented by the English auction, which can clearly be written as a menu mechanism.²³ Our results show that these insights extend to our variant model with *strictness*: the English

²³ In fact, ex-post perfect implementation has been investigated for a variety of more complex auction environments (Ausubel, 2004; Ausubel, 2006; Sun and Yang, 2014; Drexl and Kleiner, 2015).

auction implements a variant of the Vickrey rule, where the winner pays the highest possible price that is no greater than the second-highest valuation (see Fig. 6).²⁴

G.3. School choice and marriage

Third, we consider school choice (Gale and Shapley, 1962). Suppose there is a finite set of students and a finite set of schools, including a special *outside option*. Moreover, each school has a quota (or capacity), where the quota of the outside option is infinite. Let us say that an environment is a *school choice environment* if and only if (i) the agents are the students; (ii) for each agent, the possible assignments are the schools; (iii) an outcome is any assignment profile such that no school is assigned to more students than its quota; and (iv) for each agent, the possible types are the strict rankings of his assignments.

Each profile of choice correspondences for schools \mathcal{C} determines a (*student-proposing*) *deferred acceptance rule*, $f^{DA|\mathcal{C}}$ (Gale and Shapley, 1962). As with our previous examples, there is an associated menu mechanism. For each school, a priority is a strict ranking of students that implicitly deems all students acceptable. If \mathcal{C} is responsive to a profile of priorities \mathbf{p} and respects the quotas, then $f^{DA|\mathcal{C}}$ is *strategy-proof* (Dubins and Freedman, 1981; Roth, 1982). Moreover, the following are equivalent: (i) $f^{DA|\mathcal{C}}$ is *group strategy-proof*, (ii) $f^{DA|\mathcal{C}}$ is *efficient*, (iii) $f^{DA|\mathcal{C}}$ is *consistent*, and (iv) \mathbf{p} is *acyclic* (Ergin, 2002). It is easy to verify that our other conditions are satisfied, and thus we have:

Corollary G.3. *Fix a school choice environment, a profile of priorities \mathbf{p} , and a profile of choice correspondences \mathcal{C} that is responsive to \mathbf{p} and respects the quotas. Let G^{public} be a public menu mechanism for the associated deferred acceptance algorithm, and let $G^{private}$ be the associated private menu mechanism. Then*

- G^{public} is an ex-post perfect implementation of $f^{DA|\mathbf{p}}$ via the preferential convention;
- $G^{private}$ is an ex-post everywhere-dominant implementation of $f^{DA|\mathbf{p}}$ via the preferential convention; and
- if $f^{DA|\mathbf{p}}$ is group strategy-proof, or $f^{DA|\mathbf{p}}$ is efficient, or $f^{DA|\mathbf{p}}$ is consistent, or \mathbf{p} is acyclic, then moreover (i) G^{public} is a full subgame perfect implementation of $f^{DA|\mathbf{p}}$, and (ii) $G^{private}$ is a full everywhere-dominant implementation of $f^{DA|\mathbf{p}}$.

Observe that the first and second item of Corollary G.3 can be derived from Corollary G.1, as these school choice environments are matching with contracts environments that satisfy the given hypotheses (Hatfield and Milgrom, 2005; Hatfield et al., 2021). As discussed earlier, for the first item of Corollary G.3, similar insights were recently obtained by Kawase and Bando (2021) and Bó and Hakimov (2019).

Analogues of Corollary G.3 hold in many interesting variants of the model; we mention three. First, deferred acceptance remains *strategy-proof* even when schools individually face affirmative action constraints, including the special case where each school simply reserves a certain number of seats for minority students (Abdulkadiroğlu and Sönmez, 2003; Abdulkadiroğlu,

²⁴ The menu mechanism in Fig. 6 is simpler than the pruned cumulative offers process mechanism described in the proof of Corollary G.1. In particular, the latter has a history where 2 selects \emptyset or 3 immediately after 1 selects 3; this history can be removed while preserving implementation.

2005; Hafalir et al., 2013; Ehlers et al., 2014). Second, deferred acceptance remains *strategy-proof* even when schools collectively face distributional constraints, which has been applied to the matching of doctors to hospitals in Japan while respecting the Japanese regional caps (Kamada and Kojima, 2015; Kamada and Kojima, 2018). Finally, in the special case where each school has a quota of one, we can view students as men and schools as women (or alternatively, view students as women and schools as men); this is the classical marriage problem (Gale and Shapley, 1962).

G.4. Object allocation

Fourth, we consider object allocation without money. Suppose there is a finite set of agents and a finite set of objects, with at least many objects as agents. Let us say that an environment is an *object allocation environment* if and only if (i) for each agent, the possible assignments are the objects; (ii) an outcome is any assignment profile where no object is assigned to more than one agent; and (iii) for each agent, the possible types are the strict rankings of his assignments.

Each consistent control rights structure ω determines an associated *trading cycles rule*, $f^{TC|\omega}$, and the class of trading cycles rules is precisely the class of *efficient* and *group strategy-proof* rules (Pycia and Ünver, 2017; Bade, 2020). This class includes the hierarchical exchange rules (Pápai, 2000), which in turn include both (i) Gale’s top trading cycles (reported in Shapley and Scarf, 1974); and (ii) serial dictatorship (see, for example, Svensson, 1999).

Each trading cycles rule is defined by a trading cycles algorithm that has an associated menu mechanism, and for hierarchical exchange rules these menu mechanisms are *non-repeating*, but in general they are not due to the presence of brokers. In the trading cycles algorithm, a broker cannot select his own object until late in the procedure. For our menu mechanisms, we modify this to recover our *non-repeating* property in a manner that has no impact on the algorithm: if a broker selects his own object before the algorithm wants him to, then he must immediately select another object. To ensure that the private menu mechanisms are *reaction-proof*, the agents should play in a simple order; see Example 3. It is easy to verify that our other conditions are satisfied, and thus we have:

Corollary G.4. *Fix an object allocation environment and a consistent control rights structure ω . Let G^{public} be a non-repeating public menu mechanism adapted for the associated trading cycles algorithm, and let $G^{private}$ be the associated private menu mechanism. Then*

- G^{public} is both (i) an ex-post perfect implementation of $f^{TC|\omega}$ via the preferential convention, and (ii) a full subgame perfect implementation of $f^{TC|\omega}$; and
- $G^{private}$ is both (i) an ex-post everywhere-dominant implementation of $f^{TC|\omega}$ via the preferential convention, and (ii) a full everywhere-dominant implementation of $f^{TC|\omega}$.

Observe that in this model, every *efficient* and *group strategy-proof* rule has a double implementation in terms of everywhere-dominant strategy equilibrium.

G.5. Two-candidate elections

Finally, we consider two-candidate elections. Suppose there is a finite set of voters and two candidates. Let us say that an environment is a *two-candidate election environment* if and only if (i) for each agent, the possible assignments are the candidates; (ii) an outcome is any assignment

profile where all agents are assigned the same candidate; and (iii) for each agent, the possible types are the strict rankings of his assignments.

Each committee C determines a voting by committees rule f^C , and the class of voting by committees rules is precisely the class of *strategy-proof* and *onto* rules (Barberà et al., 1991). It is easy to see that *strategy-proofness* is equivalent to *group strategy-proofness* for these environments, and to verify that our other conditions are satisfied; thus we have:

Corollary G.5. *Fix a two-candidate election environment and a committee C . Let G^{public} be a public menu mechanism for the associated direct mechanism, and let $G^{private}$ be the associated private menu mechanism. Then*

- G^{public} is both (i) an ex-post perfect implementation of f^C via the preferential convention, and (ii) a full subgame perfect implementation of f^C ; and
- $G^{private}$ is both (i) an ex-post everywhere-dominant implementation of f^C via the preferential convention, and (ii) a full everywhere-dominant implementation of f^C .

We remark that in this setting, because players have strict preferences over all outcomes, implementation in subgame perfect equilibrium is equivalent to implementation in guided iteratively undominated strategies, which has been proposed as a simple implementation where players receive assistance in making their calculations (Glazer and Rubinstein, 1996). When there are three or more candidates, *richness* and *strictness* imply that the only *strategy-proof* rules that respect consensus are the dictator rules (Gibbard, 1973; Satterthwaite, 1975).

References

- Abdulkadiroğlu, Atila, 2005. College admissions with affirmative action. *Int. J. Game Theory* 33, 535–549.
- Abdulkadiroğlu, Atila, Sönmez, Tayfun, 2003. School choice: a mechanism design approach. *Am. Econ. Rev.* 93, 729–747.
- Akbarpour, Mohammad, Li, Shengwu, 2020. Credible mechanisms. *Econometrica* 88, 425–467.
- Alós-Ferrer, Carlos, Ritzberger, Klaus, 2016. *The Theory of Extensive Form Games*. Springer-Verlag Berlin Heidelberg, Berlin, Germany.
- Arribillaga, R. Pablo, Massó, Jordi, Neme, Alejandro, 2021. All Sequential Allotment Rules Are Obviously Strategy-Proof. Working paper.
- Ashlagi, Itai, Gonczarowski, Yannai, 2018. Stable matching mechanisms are not obviously strategy-proof. *J. Econ. Theory* 177, 405–425.
- Ausubel, Lawrence, 2004. An efficient ascending-bid auction for multiple objects. *Am. Econ. Rev.* 94, 1452–1475.
- Ausubel, Lawrence, 2006. An efficient dynamic auction for heterogeneous commodities. *Am. Econ. Rev.* 96, 602–629.
- Bade, Sophie, 2020. Random serial dictatorship: the one and only. *Math. Oper. Res.* 45, 353–368.
- Barberà, Salvador, Jackson, Matthew, Neme, Alejandro, 1997. Strategy-proof allotment rules. *Games Econ. Behav.* 18, 1–21.
- Barberà, Salvador, Sonnenschein, Hugo, Zhou, Lin, 1991. Voting by committees. *Econometrica* 59, 595–609.
- Battigalli, Pierpaolo, Siniscalchi, Marciano, 2003. Rationalization and incomplete information. *Adv. Theor. Econ.* 3, 3.
- Battigalli, Pierpaolo, Siniscalchi, Marciano, 2007. Interactive epistemology in games with payoff uncertainty. *Res. Econ.* 61, 165–184.
- Bergemann, Dirk, Morris, Stephen, 2005. Robust mechanism design. *Econometrica* 73, 1771–1813.
- Blackwell, David, 1951. Comparison of experiments. In: Neyman, Jerzy (Ed.), *Proceedings of the Second Berkeley Symposium on Mathematical Statistics and Probability*. University of California Press, Berkeley, California, pp. 93–102.
- Bó, Inácio, Hakimov, Rustamdjan, 2019. The iterative deferred acceptance mechanism. Working paper.
- Bó, Inácio, Hakimov, Rustamdjan, 2020a. Iterative versus standard deferred acceptance: experimental evidence. *Econ. J.* 130, 356–392.
- Bó, Inácio, Hakimov, Rustamdjan, 2020b. Pick-an-object mechanisms. Working paper.

- Chambers, Christopher, Echenique, Federico, 2016. *Revealed Preference Theory*. Cambridge University Press, Cambridge, United Kingdom.
- Crawford, Vincent, Knoer, Elsie, 1981. Job matching with heterogeneous firms and workers. *Econometrica* 49, 437–450.
- Demange, Gabrielle, Gale, David, Sotomayor, Marilda, 1986. Multi-item auctions. *J. Polit. Econ.* 94, 863–872.
- Drexl, Moritz, Kleiner, Andreas, 2015. An Efficient Dynamic Auction for General Economies with Indivisibilities. Working paper.
- Dubins, Lester, Freedman, David, 1981. Machiavelli and the Gale-Shapley algorithm. *Am. Math. Mon.* 88, 485–494.
- Dur, Umut, Kominers, Scott, Pathak, Parag, Sönmez, Tayfun, 2018. Reserve design: unintended consequences and the demise of Boston's walk zones. *J. Polit. Econ.* 126, 2457–2479.
- Ehlers, Lars, Hafalir, Isa, Yenmez, M. Bumin, Yildirim, Muhammed, 2014. School choice with controlled choice constraints: hard bounds versus soft bounds. *J. Econ. Theory* 153, 648–683.
- Eliaz, Kfir, Eilat, Ran, Mu, Xiaosheng, 2021. Bayesian privacy. *Theor. Econ.* 16, 1557–1603.
- Ergin, Haluk, 2002. Efficient resource allocation on the basis of priorities. *Econometrica* 70, 2489–2497.
- Escudé, Matteo, Sinander, Ludvig, 2020. Strictly strategy-proof auctions. *Math. Soc. Sci.* 107, 13–16.
- Fadel, Ronald, Segal, Ilya, 2009. The communication cost of selfishness. *J. Econ. Theory* 144, 1895–1920.
- Gale, David, Shapley, Lloyd, 1962. College admissions and the stability of marriage. *Am. Math. Mon.* 69, 9–15.
- Gershkov, Alex, Moldovanu, Benny, Shi, Xianwen, 2017. Optimal voting rules. *Rev. Econ. Stud.* 84, 688–717.
- Gibbard, Allan, 1973. Manipulation of voting schemes: a general result. *Econometrica* 41, 587–601.
- Glazer, Jacob, Rubinstein, Ariel, 1996. An extensive game as a guide for solving a normal game. *J. Econ. Theory* 70, 32–42.
- Gradwohl, Ronen, 2018. Privacy in implementation. *Soc. Choice Welf.* 50, 547–580.
- Gradwohl, Ronen, Smorodinsky, Rann, 2017. Perception games and privacy. *Games Econ. Behav.* 104, 293–308.
- Hafalir, Isa, Yenmez, M. Bumin, Yildirim, Muhammed, 2013. Effective affirmative action in school choice. *Theor. Econ.* 8, 325–363.
- Hagiwara, Makoto, 2020. Double Implementation in Dominant Strategy Equilibria and Ex Post Equilibria with Private Values. Working paper.
- Hakimov, Rustamdjan, Raghavan, Madhav, 2020. Transparency in Centralized Allocation: Theory and Experiment. Working paper.
- Hatfield, John, Kojima, Fuhito, 2009. Group incentive compatibility for matching with contracts. *Games Econ. Behav.* 67, 745–749.
- Hatfield, John, Kominers, Scott, Westkamp, Alexander, 2021. Stability, strategy-proofness, and cumulative offer mechanisms. *Rev. Econ. Stud.* 88, 1457–1502.
- Hatfield, John, Milgrom, Paul, 2005. Matching with contracts. *Am. Econ. Rev.* 95, 913–935.
- Jung v. Association of American Medical Colleges, 02-CV-00873 (DDC 2002).
- Kamada, Yuichiro, Kojima, Fuhito, 2015. Efficient matching under distributional constraints: theory and applications. *Am. Econ. Rev.* 105, 67–99.
- Kamada, Yuichiro, Kojima, Fuhito, 2018. Stability and strategy-proofness for matching with constraints: a necessary and sufficient condition. *Theor. Econ.* 13, 761–793.
- Kawase, Yasushi, Bando, Keisuke, 2021. Subgame perfect equilibria under the deferred acceptance algorithm. *Int. J. Game Theory* 50, 503–546.
- Kelso, Alexander, Crawford, Vincent, 1982. Job matching, coalition formation, and gross substitutes. *Econometrica* 50, 1483–1504.
- Kleiner, Andreas, Moldovanu, Benny, 2017. Content-based agendas and qualified majorities in sequential voting. *Am. Econ. Rev.* 107, 1477–1506.
- Kleiner, Andreas, Moldovanu, Benny, 2019. Abortions, Brexit, and Trees. Working paper.
- Klijn, Flip, Pais, Joana, Vorsatz, Marc, 2019. Static versus dynamic deferred acceptance in school choice: theory and experiment. *Games Econ. Behav.* 113, 147–163.
- Kojima, Fuhito, 2010. Impossibility of stable and nonbossy matching mechanisms. *Econ. Lett.* 107, 69–70.
- Kushilevitz, Eyal, Nisan, Noam, 1997. *Communication Complexity*. Cambridge University Press, Cambridge, United Kingdom.
- Li, Shengwu, 2017. Obviously strategy-proof mechanisms. *Am. Econ. Rev.* 107, 3257–3287.
- Mackenzie, Andrew, 2020. A revelation principle for obviously strategy-proof implementation. *Games Econ. Behav.* 124, 512–533.
- Marx, Leslie, Swinkels, Jeroen, 1997. Order independence for iterated weak dominance. *Games Econ. Behav.* 18, 219–245.
- Maskin, Eric, 1979. Incentive schemes immune to group manipulation. *Mimeo*.

- Maskin, Eric, 1999. Nash equilibrium and welfare optimality. *Rev. Econ. Stud.* 66, 23–38.
- Milgrom, Paul, Segal, Ilya, 2020. Clock auctions and radio spectrum reallocation. *J. Polit. Econ.* 128, 1–31.
- Moore, John, Repullo, Rafael, 1988. Subgame perfect implementation. *Econometrica* 56, 1191–1220.
- Myerson, Roger, 1981. Optimal auction design. *Math. Oper. Res.* 6, 58–73.
- Pápai, Szilvia, 2000. Strategyproof assignment by hierarchical exchange. *Econometrica* 68, 1403–1433.
- Pearce, David, 1984. Rationalizable strategic behavior and the problem of perfection. *Econometrica* 52, 1029–1050.
- Penta, Antonio, 2015. Robust dynamic implementation. *J. Econ. Theory* 160, 280–316.
- Pycia, Marek, Trojan, Peter, 2019. A Theory of Simplicity in Games and Mechanism Design. Working paper.
- Pycia, Marek, Ünver, Utku, 2017. Incentive compatible allocation and exchange of discrete resources. *Theor. Econ.* 12, 287–329.
- Roth, Alvin, 1982. The economics of matching: stability and incentives. *Math. Oper. Res.* 7, 617–628.
- Roth, Alvin, Xing, Xiaolin, 1997. Turnaround time and bottlenecks in market clearing: decentralized matching in the market for clinical psychologists. *J. Polit. Econ.* 105, 284–329.
- Rothkopf, Michael, Teisberg, Thomas, Kahn, Edward, 1990. Why are Vickrey auctions rare? *J. Polit. Econ.* 98, 94–109.
- Saijo, Tatsuyoshi, Sjöström, Tomas, Yamato, Takehiko, 2007. Secure implementation. *Theor. Econ.* 2, 203–229.
- Saijo, Tatsuyoshi, Tatamitani, Yoshikatsu, Yamato, Takehiko, 1996. Toward natural implementation. *Int. Econ. Rev.* 37, 949–980.
- Satterthwaite, Mark, 1975. Strategy-proofness and Arrow's conditions: existence and correspondence theorems for voting procedures and social welfare functions. *J. Econ. Theory* 10, 187–217.
- Satterthwaite, Mark, Sonnenschein, Hugo, 1981. Strategy-proof allocation mechanisms at differentiable points. *Rev. Econ. Stud.* 48, 587–597.
- Schummer, James, Velez, Rodrigo, 2021. Sequential preference revelation in incomplete information settings. *Am. Econ. J. Microecon.* 13, 116–147.
- Segal, Ilya, 2007. The communication requirements of social choice rules and supporting budget sets. *J. Econ. Theory* 136, 341–378.
- Segal, Ilya, 2010. Nash implementation with little communication. *Theor. Econ.* 5, 51–71.
- Shapley, Lloyd, Scarf, Herbert, 1974. On cores and indivisibility. *J. Math. Econ.* 1, 23–37.
- Shimomi, Makoto, Watson, Joel, 1998. Conditional dominance, rationalizability, and game forms. *J. Econ. Theory* 83, 161–195.
- Sprumont, Yves, 1991. The division problem with single-peaked preferences: a characterization of the uniform allocation rule. *Econometrica* 59, 509–519.
- Sun, Ning, Yang, Zifu, 2014. An efficient and incentive compatible dynamic auction for multiple complements. *J. Polit. Econ.* 112, 422–466.
- Svensson, Lars-Gunnar, 1999. Strategy-proof allocation of indivisible goods. *Soc. Choice Welf.* 16, 557–567.
- Szpilrajn, Edward, 1930. Sur l'extension de l'ordre partiel [in French]. *Fundam. Math.* 16, 386–389.
- Takamiya, Koji, 2001. Coalition strategy-proofness and monotonicity in Shapley-Scarf housing markets. *Math. Soc. Sci.* 41, 201–213.
- Thomson, William, 2016. Non-bossiness. *Soc. Choice Welf.* 47, 665–696.
- Vickrey, William, 1961. Counterspeculation, auctions, and competitive sealed tenders. *J. Finance* 16, 8–37.
- Yamato, Takehiko, 1993. Double implementation in Nash and undominated Nash equilibria. *J. Econ. Theory* 59, 311–323.
- Yao, Andrew, 1979. Some complexity questions related to distributive computing (preliminary report). In: *STOC '79: Proceedings of the Eleventh Annual ACM Symposium on Theory of Computing*. Association for Computing Machinery, Atlanta, Georgia.