Kyoto University,
Graduate School of Economics
Discussion Paper Series

Menu Mechanisms

Andrew Mackenzie

Yu Zhou

Discussion Paper No. E-19-012

Graduate School of Economics
Kyoto University
Yoshida-Hommachi, Sakyo-ku
Kyoto City, 606-8501, Japan

March, 2020

Menu mechanisms

Andrew Mackenzie* and Yu Zhou†‡

This draft: March 23, 2020

Abstract

We investigate *menu mechanisms*: dynamic mechanisms where at each history, an agent selects from a menu of his possible assignments. In comparison to direct mechanisms, menu mechanisms offer better privacy to participants; we formalize this with a novel notion of mechanism informativeness. We consider both ex-post implementation and full implementation, for both subgame perfection and 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 elections, marriage, college admissions, auctions, labor markets, matching with contracts, and object allocation.

Keywords: menu mechanism, privacy, strategy-proofness, robust implementation

JEL Codes: D82, D47, C78

*Department of Microeconomics and Public Economics, Maastricht University, Maastricht, the Netherlands. Email: a.mackenzie@maastrichtuniversity.nl

†Graduate School of Economics, Kyoto University, Kyoto, Japan. Email: zhouyu_0105@hotmail.com

‡Yu Zhou gratefully acknowledges financial support from the Grant-in-aid for Research Activity of the Japan Society for the Promotion of Science (19K13653). We thank Masaki Aoyagi, Inácio Bó, Albin Erlanson, Scott Kominers, Takashi Kunimoto, Jordi Massó, Nozomu Muto, Marek Pycia, Al Roth, Jim Schummer, Alex Teytelboym, and William Thomson; as well as seminar participants at Yokohama National University.

1 Introduction

1.1 Overview

In 2018, the European Union implemented the General Data Protection Regulation, which broadly promotes an individual's right to control his or her personal data. In Chapter 2, Article 5, the law introduces the principle of *data minimization*:

> Personal data shall be adequate, relevant and limited to what is necessary in relation to the purposes for which they are processed.

In this article, we investigate the general application of this principle to a broad variety of problems in mechanism design.

In particular, we consider finite environments with private values and no consumption externalities, and we investigate privacy improvements over existing dominant strategy implementations. 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*: each agent reports all of his private information to a central administrator, who then uses these reports to calculate the desired outcome. Of course, for any reasonable definition of data minimization, direct mechanisms perform disastrously.

We formalize the informativeness of an implementation by considering the perspective of an observer who is (i) interested in the collective private information of the agents, (ii) able to observe all actions taken by the agents, and (iii) working with a natural model about how his observations are generated.[1] The informativeness of the implementation, then, is the partition of the set of type profiles representing what this observer might learn. Instead of seeking a minimally informative implementation for a particular rule, in this article we seek privacy improvements over direct mechanisms—that is, implementations with coarser type profile partitions—for a variety of rules across a variety of environments.

To do so, we introduce *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

- direct menu mechanisms in a variety of environments, including for voting by committees (Barberà, Sonnenschein, and Zhou, 1991) in two-candidate election environments;
- student-proposing deferred acceptance (Gale and Shapley, 1962) in college admissions environments, including male- and female-proposing deferred acceptance in marriage environments;

[1] In the context of partial implementation, the observer assumes that agents conform to the given type-strategy profile, while in the context of full implementation, the observer assumes that agents play some equilibrium given the type profile.

- the salary adjustment process (Crawford and Knoer, 1981; Kelso and Crawford, 1982) in labor market environments, including the English auction in one-object auction environments;

- the cumulative offers process (Hatfield and Milgrom, 2005) in matching with contracts environments; and

- 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 environments.

We illustrate three of these examples in Figure 1. With the exception of direct menu mechanisms, the above menu mechanisms are all clearly less informative than direct mechanisms. This raises the question: are some of these menu mechanisms incentive compatible, and if so, in what way?

Our main results imply that *all* of the menu mechanisms for the above examples are *robustly* incentive compatible; see Section 5 for a detailed discussion of these examples and related ones. In particular, we consider both ex-post implementation and full implementation, for both subgame perfection and *everywhere-dominance*: a strengthening of dominance that covers off-path histories. 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.

Formally, we provide sufficient conditions for a menu mechanism to provide a robust implementation of a rule. We consider 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, while *group strategy-proofness* is the usual strong version requiring that no group of agents can jointly misreport to obtain a Pareto-improvement.

- For the menu mechanism, *non-repeating* requires that an agent can never select the same assignment twice, *public* requires that each agent observes all actions of his peers, and *private* requires that each agent observes no actions of his peers. For private menu mechanisms, *reaction-proofness* requires that whenever one agent can deduce something about another agent's choices, the latter's assignment has already been determined.

- For type-strategy profiles—which for convenience, we refer to as *conventions*—*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 achieved if all agents conform.

Some of the rules we consider are *strategy-proof* but not *group strategy-proof*, and therefore only some of our results apply. Otherwise, each of our applications either satisfies all conditions or violates just one; in these latter cases, there are simple workarounds.[2]

[2]In particular, we consider labor market environments and one-object auction environments that

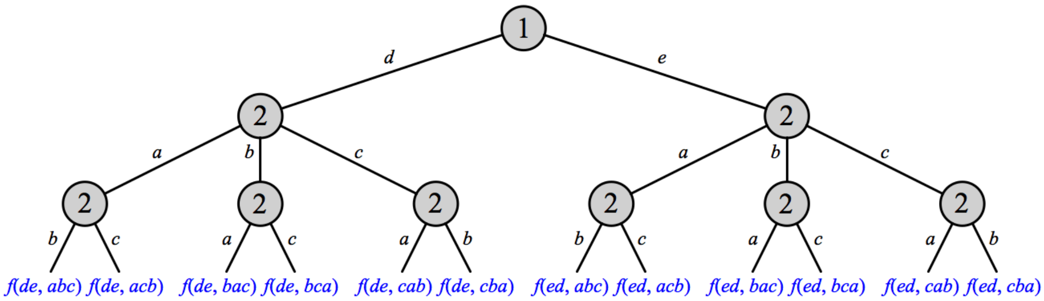

(a) Direct menu mechanism.

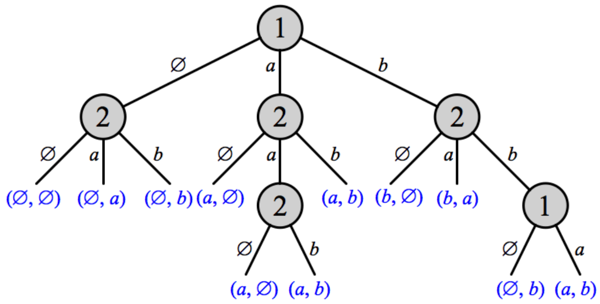

(b) Male-proposing deferred acceptance menu mechanism.

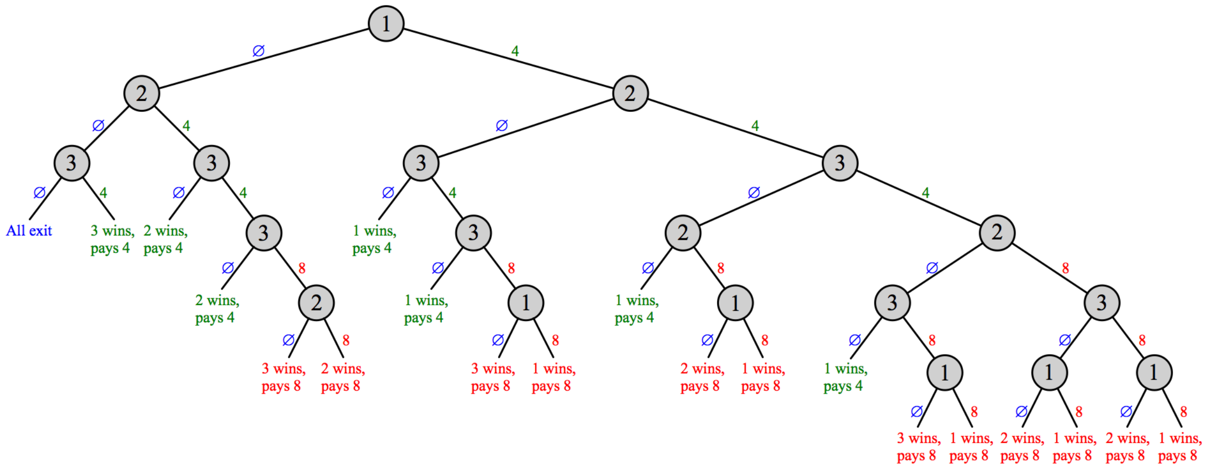

(c) English auction menu mechanism.

Figure 1: *Examples of menu mechanisms.* **(a)** There are two agents: 1 can consume from $\{d, e\}$ and have any strict ranking, and 2 can consume from $\{a, b, c\}$ and have any strict ranking. The agents sequentially reveal their complete preference rankings. **(b)** There are two men, 1 and 2, and two women, a and b. At each history, a man either (i) proposes to a woman who has not yet rejected him, or (ii) opts to remain single (denoted $\emptyset$). Each woman processes proposals mechanically, tentatively accepting her most preferred suitor while rejecting all others. In this example, a prefers 1 to 2, while b prefers 2 to 1. **(c)** There are three bidders, 1 and 2 and 3, and one object with two possible prices, 4 and 8. At each history, a bidder can either (i) bid at the current price, which ascends over the course of the auction; or (ii) exit (denoted $\emptyset$). If several agents make the same bid, then the one with lowest index is a tentative winner while the others are asked to bid again.

We draw stronger conclusions for private menu mechanisms than public menu mechanisms, and emphasize that these are dramatically different institutions. For example, the public deferred acceptance menu mechanism might describe courtship in a ballroom, while the private one might describe courtship through a dating app. As another example, the public English auction menu mechanism might describe a sale at an auction house, while the private one might describe a sale online (say, on e-Bay). In general, we find it convenient to think of private menu mechanisms as smart phone apps that occasionally notify users that they must select from a given menu. With this in mind, our main results are the following:

Theorem 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.*

Theorem 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 and a full subgame perfect implementation of the rule.*

Theorem 3: *For each rich environment, each strategy-proof rule, each non-repeating and reaction-proof private menu mechanism, and each preferential convention that is compatible with the rule, the private menu mechanism is an ex-post everywhere-dominant implementation of the rule via the convention.*

Theorem 4: *For each rich and strict environment, each group strategy-proof rule, each non-repeating and reaction-proof private menu mechanism, and each preferential convention that is compatible with the rule, the private menu mechanism is both an ex-post everywhere-dominant implementation of the rule and a full everywhere-dominant implementation of the rule.*

As corollaries, our theorems immediately provide novel results for two-candidate elections, marriage, college admissions, auctions with unit demand, labor markets, matching with contracts, and object allocation; see Section 5 for details. Taken together, our results show that like direct mechanisms, menu mechanisms can systematically provide robust implementations, but unlike direct mechanisms, they can require agents to reveal substantially less than their complete private information. For public menu mechanisms, this privacy improvement comes at the cost of moving from dominant strategy implementation to robust implementation in terms of subgame perfection, and this tradeoff may be worthwhile in some situations but not in others. For private menu mechanisms, the privacy improvement comes at no cost in terms of implementation. That said, public menu mechanisms may be more desirable when transparency about the behavior of all agents is required by law; for example, this is required by the Freedom of Information Act for procurement auctions in the United States (Bergemann and Hörner, 2018).

satisfy *strictness* but violate *richness*; for these applications, we can still apply our results indirectly by enriching the environment and then pruning the associated menu mechanism. The menu mechanisms for trading cycles rules that are not hierarchical exchange rules must be slightly modified to be *non-repeating*.

1.2 Literature

Our paper is closely related to the literature on our leading examples, which we discuss in detail in Section 5. Moreover, our paper is related to recent experiments, recent results on special classes of menu mechanisms, implementation theory, privacy in mechanism design, and market design; we discuss each in sequence.

First, in addition to the privacy benefits of menu mechanisms, recent experiments suggest that there may also be simplicity benefits, as measured by the likelihood of subjects conforming to the convention. In particular, the evidence suggests that while obviously strategy-proof mechanisms[3] generally outperform both menu mechanisms and direct mechanisms (Bó and Hakimov, 2020b), menu mechanisms outperform direct mechanisms for both deferred acceptance (Klijn, Pais, and Vorsatz, 2019; Bó and Hakimov, 2020a) and top trading cycles (Bó and Hakimov, 2020b); thus menu mechanisms can jointly provide privacy benefits and simplicity benefits even when obviously strategy-proof implementations are not available. That said, unfortunately these simplicity benefits are not universal: there is no evidence of such an improvement for serial dictatorship (Bó and Hakimov, 2020b).

Second, our results complement recent results in the literature, which we describe using our language. First, Kawase and Bando (2018) prove that for each game associated with a deferred acceptance public menu mechanism, honesty is a subgame perfect equilibrium; this is an implication of our Theorem 1. Interestingly, Kawase and Bando (2018) 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, they violate our *non-repeating* requirement, and therefore our results do not apply. Second, Bó and Hakimov (2019) prove that for deferred acceptance menu mechanisms, honesty is an ordinal perfect equilibrium, and Bó and Hakimov (2020b) extend this result to *pick-an-object mechanisms* for one-sided matching markets; these are menu mechanisms where each agent necessarily consumes the last object he selected. These results are similar to our Theorem 1, but involve a natural subclass of our menu mechanisms and a different solution concept.

Third, our paper is related to two topics in the broader literature on implementation theory: ex-post perfect implementation and double implementation. 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) and voting (Kleiner and Moldovanu, 2017; Gershkov, Moldovanu, and Shi, 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 no two objects are identical and no agent is indifferent between exiting and winning for some price; and (ii) elections with two candidates. Indeed, our paper primarily complements these previous contributions by applying to matching environments. 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.[4]

[3]For rich and strict environments, if a rule has an obviously strategy-proof implementation (Li, 2017), 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.

[4]Double implementation has previously been investigated for full Nash implementation with full

Fourth, we contribute to the literature on privacy in mechanism design with our notion of informativeness. As we discuss later, our privacy ordering for ex-post implementations can be viewed as the ordinary Blackwell ordering (Blackwell, 1951), though this is not the case for full implementations. Moreover, our notion is conceptually related to *informational efficiency*, which involves a formal notion of message space size (Hurwicz, 1960). In the context of auctions, it has previously been argued that one advantage of English auctions over direct mechanisms is that the winner only needs to reveal that his valuation is higher than the second-highest, which is desirable if he expects extensive negotiations to follow or if he does not want to reveal the extent of his technological advantage to competitors (Rothkopf, Teisberg, and Kahn, 1990); this notion is formalized by Milgrom and Segal (2020).[5] With respect to previous notions, the primary advantage of our approach is that it allows us to partially order implementations, and to do so very generally in the context of robust implementation. We remark that other interesting notions of privacy that are less closely related to ours have also been considered.[6]

Finally, our paper is part of the broader literature identifying practical problems with direct mechanisms in market design. As examples, (i) 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, in violation of anti-trust law (Jung et al. versus Association of American Medical Colleges et al., 2002); (ii) direct mechanisms may not be credible because the administrator can deliberately miscalculate outcomes, such as by shill bidding in auctions (Akbarpour and Li, 2020); (iii) the administrator may accidentally miscalculate outcomes, such as when the City of Boston unintentionally subverted the goal of prioritizing nearby students in admissions at each public school, because the assignment algorithm filled the seats reserved for these students first (Dur, Kominers, Pathak, and Sönmez, 2018); and (iv) direct mechanisms may make it unnecessarily difficult for agents who have trouble with contingent reasoning to identify dominant strategies (Li, 2017). To give the complete picture, though, we should mention that alternatives to direct mechanisms have their problems, too. As examples, (i) the National Resident Matching Program is a centralized labor market clearinghouse for medical residents in the United States that was designed in response to market failures via *unraveling*, where hospitals made earlier and earlier exploding offers (Roth and Peranson, 1999), and (ii) even in decentralized labor markets by telephone with extremely short turnaround times, such as the entry-level market for clinical psychologists, there can be bottlenecks that slow the market and promote strategic behavior (Roth and Xing, 1997). Within this literature, we formalize the point that dominant strategy implementation can be preserved when moving from a direct mechanism to a dynamic mechanism that requires agents to reveal considerably less information, such as when moving from a

undominated Nash implementation (Yamato, 1993), as well as for full Nash implementation with full dominant strategy implementation (Saijo, Sjöström, and Yamato, 2007).

[5]In particular, Milgrom and Segal (2020) say that an auction satisfies *unconditional winner privacy* if and only if—in our language—for each agent, and for each pair of type profiles where (i) this agent wins and (ii) his peers' types are the same, the convention assigns both type profiles to the same play. Clearly, if an auction satisfies *unconditional winner privacy*, then any less informative auction does as well.

[6]For example, in the Bayesian tradition, Eliaz, Eilat, and Mu (2019) consider a notion of privacy based on the difference between the planner's prior and posterior. In perception games, agents have privacy concerns that influence their behavior as they care about the beliefs of observers (Gradwohl and Smorodinsky, 2017). See Milgrom and Segal (2020) for a discussion of other privacy notions in the computer science literature, where cryptography technology is taken into account.

centralized clearinghouse to a market organized by a smart phone app.

2 Model

2.1 Preliminaries

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
- 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.

In a given environment, the agents wish to condition the outcome on their collective private information according to a (social choice) rule. We primarily focus on strategy-proof rules, for which honesty is always a dominant strategy in the associated direct mechanism:

Definition: Fix an environment. A *rule* is a function $f : \Theta \to X$. 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 $f(\theta_i, \theta_{-i}) \; R_i(\theta_i) \; f(\theta'_i, \theta_{-i})$.

We are interested in comparing mechanisms that implement a given rule. A mechanism is simply an extensive game form with players in N and outcomes in X; see Appendix E 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 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.

Table 1: Notation for a generic mechanism.

Name	Notation	Representative element
Set of histories	H	h
Precedence relation over histories	$\precsim$	
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	$\mathcal{X}$	

We formalize this in several ways, using both type-strategy profiles in a mechanism and strategy profiles in its 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: *Strategy profiles, related notation, and conventions.* Fix an environment and a mechanism.

- For each $i \in N$, a *(pure) strategy for* i is a mapping $s_i : H_i \to \mathcal{A}$ 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.
- For each $h \in H$ and each $s \in S$, define $\pi^h(s)$ to be the play that first proceeds from the initial history to h and then proceeds according to s. Moreover, 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 $\mathcal{X}(s)$.
- For each $i \in N$, a *type-strategy for* i is a mapping $\mathbb{S}_i : \Theta_i \to S_i$. A *convention* is a profile of type-strategies $\mathbb{S} = (\mathbb{S}_i)_{i \in N}$.

In a given game, we formalize the plausibility of a strategy profile in four different ways. The first three solution concepts are standard, while the fourth is to our knowledge novel:

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:

- *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 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 $s'_i \in S_i$, $\mathcal{X}^h(s)\ R_i\ \mathcal{X}^h(s'_i, s_{-i})$.
- *everywhere-dominant 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 $s'_i \in S_i$, $\mathcal{X}^h(s_i, s'_{-i})\ R_i\ \mathcal{X}^h(s'_i, s'_{-i})$.

For each of these solution concepts, we consider two notions of implementation: (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; and (ii) full, which requires that at each type profile, there are equilibria and each of them yields the desired outcome:

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

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

We say that G is a *full* $\mathbf{SC}$-*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 $s \in \mathbf{SC}(G, R(\theta))$, $\mathcal{X}(s) = f(\theta)$.

We sometimes say *dominant implementation* for ex-post $\mathbf{DE}$-implementation, *ex-post perfect implementation* for ex-post $\mathbf{SPE}$-implementation, *ex-post everywhere-dominant implementation* for ex-post $\mathbf{EDE}$-implementation, *full subgame perfect implementation* for full $\mathbf{SPE}$-implementation, and *full everywhere-dominant implementation* for full $\mathbf{EDE}$-implementation.

We remark that though ex-post implementation and full implementation are similar, they are in fact logically independent; we illustrate this for $\mathbf{SPE}$ in Figure 2.

2.2 Informativeness

We are interested in comparing mechanisms that implement a given rule on the basis of the privacy that they afford to participants. Informally, 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.[7] In this case, we say that the informativeness of a mechanism describes what this observer can learn. Of course, this depends on the model used by the observer: if the observer believes that the actions and the type profile have no correlation, then the mechanism is not at all informative. We consider two focal models: (i) the observer assumes that the convention $\mathbb{S}$ has been played, and (ii) the observer assumes that some $\mathbf{SC}$-equilibrium has been played.

[7]There are may natural variants that we could also consider; for example, the observer only observes the final outcome.

(a) Only ex-post perfect. **(b)** Only full subgame perfect.

Figure 2: Independence of implementations. *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 Figure 2a. *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 Figure 2b, 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 $a\ P\ c\ P\ b$, or it specifies that he should select c when he has this ranking, but it does not specify both.

Definition: *Informativeness.* Fix an environment and a mechanism. For each convention $\mathbb{S}$,

- for each $\pi \in \Pi$, define the *set of types compatible with* π *given* $\mathbb{S}$ by
$$\Theta^{\mathbb{S}}(\pi) \equiv \{\theta \in \Theta | \pi(\mathbb{S}(\theta)) = \pi\};\ \text{and}$$
- define the *informativeness of* G *given* $\mathbb{S}$ by $\{\Theta^{\mathbb{S}}(\pi)\}_{\pi \in \Pi} \backslash \{\emptyset\}$.

For each solution concept **SC**,

- for each $\pi \in \Pi$, define the *set of types compatible with* π *given* **SC** by
$$\Theta^{\mathbf{SC}}(\pi) \equiv \{\theta \in \Theta |\ \text{there is}\ s \in \mathbf{SC}(G, R(\theta))\ \text{such that}\ \pi(s) = \pi\};\ \text{and}$$
- define the *informativeness of* G *given* **SC** by $\{\Theta^{\mathbf{SC}}(\pi)\}_{\pi \in \Pi} \backslash \{\emptyset\}$.

We compare the relative informativeness of mechanism-model pairs in the obvious way, by the coarseness of their partitions of the type profiles:

Definition: *Relative informativeness.* Fix an environment. Let (G, M) and (G', M') each be a mechanism with an associated model (either a convention or a solution concept). We say that (G, M) *is more informative than* (G', M') if and only if the type profile partition for (G, M) is finer than the type profile partition for (G', M'). In this case, we sometimes say that (G', M') improves upon the privacy of (G, M).

We remark that this can be interpreted as the standard Blackwell order for the convention models, though not for the solution concept models.[8] For a given rule f, the

[8] A Blackwell experiment is a function that associates each unknown *state* with a probability distribution over *signals*. For a convention model, simply consider that type profiles are states, plays are signals, and each state θ surely yields the signal $\pi(\mathbb{S}(\theta))$. In this case, we compare the informativeness of two mechanism-convention pairs using the standard Blackwell order (Blackwell, 1951). That said, this insight does not extend to solution concept models because we make no assumption about how a type profile generates a distribution over its equilibria.

informativeness of an associated mechanism given a convention is bounded both above and below: the partition can be no finer than $\{\{\theta\}|\theta \in \Theta\}$ (which requires each play used by the convention to occur for a unique type profile) and no coarser than $\{f^{-1}(x)\}_{x \in f(X)}$ (which requires each outcome in the range of f to occur on a unique play used by the convention). For *strategy-proof* rules, the direct mechanism is necessarily a maximally-informative implementation:

Observation: For each environment and each *strategy-proof rule*, let G be a direct mechanism and let $\mathbb{S}$ be the convention where each agent always reports his type. Then

- $(G, \mathbb{S})$ is an ex-post **N**-implementation of f;
- $(G, \mathbb{S})$ is an ex-post **DE**-implementation of f; and
- the informativeness of G given $\mathbb{S}$ is $\{\{\theta\}|\theta \in \Theta\}$.

Instead of focusing on a particular rule and seeking a minimally informative implementation, in this article we seek to systematically construct implementations that are less informative than direct mechanisms for a variety of rules across a variety of environments.

2.3 Menu mechanisms

To systematically improve upon the privacy of direct mechanisms, we introduce menu mechanisms, where agents iteratively select from menus of their assignments:

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 of the menu mechanisms we consider are derived naturally from familiar ideas in the literature; recall Figure 1. With the exception of the menu mechanisms derived from direct mechanisms, all of our examples are clearly less informative than direct mechanisms. 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 mechanisms (Rothkopf, Teisberg, and Kahn, 1990). Our main results show that, perhaps surprisingly, all of these examples also provide robust implementations of their rules. The particular kind of implementation depends on whether or not an agent's actions can be observed by his peers:

Definition: 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.[9]

As discussed in the introduction, the public and private versions of a given menu mechanism describe dramatically different institutions.

3 Public menu mechanisms

3.1 Results

In this section, we consider public menu mechanisms. We begin by providing conditions under which a public menu mechanism provides an ex-post perfect implementation. It is well-known that (i) each ex-post perfect implementation is an ex-post Nash implementation, and (ii) for environments with private values, each rule with an ex-post Nash implementation is *strategy-proof* (see, for example, Bergemann and Morris, 2005); thus *strategy-proofness* of the rule is a necessary condition. In addition to this necessary condition, we impose the following:

Definition: *Assumptions for Theorem 1.* Fix an environment, a rule f, a menu mechanism G, and a convention $\mathbb{S}$. We say that

- the 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 $R_i(\theta_i) = P_i$;
- G is *non-repeating* if and only if for each $\pi \in \Pi$, each $i \in N$, and each distinct pair $h, h' \in H_i \cap \pi$, $\alpha^h(\pi) \neq \alpha^{h'}(\pi)$;
- $\mathbb{S}$ is *f-compatible* (or *compatible with the rule*) if and only if for each $\theta \in \Theta$, we have $\mathcal{X}(\mathbb{S}(\theta)) = f(\theta)$; and
- $\mathbb{S}$ is *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 $h \in H_i$, $[\mathbb{S}_i(\theta_i)](h) = \text{argmax}_{\tau_i(\theta_i)}[\text{argmax}_{R_i(\theta_i)} X_i(h)]$.[10]

Richness requires that each agent might have any strict ranking of his assignments, which is satisfied in many matching environments but violated in auction environments (where lower payments must be preferred) and labor market environments (where higher salaries must be preferred); we discuss how to nevertheless apply our result for these particular examples in Section 5.2. The *non-repeating* requirement is that no agent can select the same assignment twice; this is satisfied by all of our examples, but violated by the menu mechanism derived from male-proposing deferred acceptance where women strategically process proposals that are mechanically submitted by men. *Compatibility with the rule* is a basic condition that is necessary for ex-post Nash implementation.

[9]Formally, for each $i \in N$ and each $h \in H_i$, let $\{h_1, h_2, ..., h_t\}$ denote $\{h' \in H_i | h' \prec h\}$ such that $h_1 \prec h_2 \prec ... \prec h_t$, and define the *experience of i at h* to be the list of menus and selections $\mathcal{E}_i(h) \equiv \big((X_i(h_1), \alpha^{h_1}(h)), (X_i(h_2), \alpha^{h_2}(h)), ..., (X_i(h_t), \alpha^{h_t}(h)), X_i(h)\big)$. 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 $\mathcal{E}_i(h) = \mathcal{E}_i(h')$.

[10]Abusing notation, if $\text{argmax}_{R_i}(X_i')$ is a singleton $\{x\}$, we sometimes let $\text{argmax}_{R_i}(X_i')$ denote x.

Finally, the *preferential* requirement is that the convention specifies that each agent should always pick a most-preferred assignment, breaking ties consistently; this is again satisfied by all of our examples. Our first theorem states that these conditions guarantee a public menu mechanism provides an ex-post perfect implementation:

THEOREM 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 C, and involves lemmas about public menu mechanisms (Appendix B) whose proofs involve lemmas about revealed preference theory (Appendix A). We sketch the arguments below:

Proof sketch. To begin, we take arbitrary $\theta \in \Theta$, $i \in N$, $h \in H_i$, and $s \in S_i$, 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 (i) always staying on a given play is rationalizable (the Play Lemma); and (ii) for any history and any continuation strategy specified by the convention, proceeding toward that history and then conforming to the continuation strategy is rationalizable (the Continuation Lemma). Both of these proofs rely critically on the fact that the public menu mechanism is *non-repeating*. Our proof of Theorem 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 $\theta^*_{-i} \in \Theta_{-i}$ such that each agent $j \in N\backslash\{i\}$ proceeds to h and then plays according to $\mathbb{S}_j(\theta_j)$. 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 $\pi_a = \pi(\mathbb{S}(\theta_i^a, \theta^*_{-i}))$ and $\pi_b = \pi(\mathbb{S}(\theta_i^b, \theta^*_{-i}))$. 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. 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. □

Next, we consider conditions under which a public menu mechanism is not only an ex-post perfect implementation, but moreover a full subgame perfect implementation. In addition to our previous assumptions, we consider the following:

DEFINITION: *Assumptions for Theorem 2 and its corollaries.* Fix an environment and a rule f. For each $i \in N$, each $\theta \in \Theta$, and each $x \in X$, define the *lower counter set of* i *for* x *given* θ, $LCS_i(x|\theta) \equiv \{x' \in X | x \; R_i(\theta_i) \; x'\}$. We say that

- the environment is *strict* if and only if for each $i \in N$ and each $\theta_i \in \Theta_i$, $R_i(\theta_i) \in \mathcal{P}_i$;
- f is *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 $f(\theta'_{N'}, \theta_{N \setminus N'}) \; P_i(\theta_i) \; f(\theta)$;
- f is *Maskin monotonic* if and only if for each pair $\theta, \theta' \in \Theta$ such that for each $i \in N$, $LCS_i(f(\theta)|\theta) \subseteq LCS_i(f(\theta)|\theta')$, 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 $f(\theta) = f(\theta'_i, \theta_{-i})$.

Strictness rules out indifference, which is common in matching settings but uncommon in auction settings. *Group strategy-proofness* requires that no coalition of agents can ever obtain a Pareto-improvement over honesty through a coordinated misrepresentation of their preference. *Maskin monotonicity* is the classic necessary condition for full **NE**-implementation (Maskin, 1999). Finally, *non-bossiness* is precisely the condition that strengthens *strategy-proofness* to *group strategy-proofness* in our environments (Satterthwaite and Sonnenschein, 1981).[11] There are strong logical relationships between these properties for rules:

THEOREM PT (PÁPAI, 2000; TAKAMIYA, 2001):[12] *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.*

Our second theorem provides conditions that guarantee a public menu mechanism is a double implementation—both an ex-post perfect implementation and a full subgame perfect implementation. Due to Theorem PT, there are several ways to state our result; we choose to do so with *group strategy-proofness* and give the other statements as corollaries:

THEOREM 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 and a full subgame perfect implementation of the rule.*

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

[12]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.

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

Proof sketch. We take an arbitrary type profile $\theta \in \Theta$, an arbitrary subgame perfect equilibrium $s^* \in \mathbf{SPE}(G, R(\theta))$, and define $s^\theta \equiv \mathbb{S}(\theta)$. By Theorem 1, we have that $s^\theta \in \mathbf{SPE}(G, R(\theta))$ 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 Theorem PT to work with *non-bossiness*. 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*, and *non-bossiness*. □

3.2 Discussion

Theorem 2 immediately yields several interesting corollaries. To complete their proofs, first observe that for each *rich* and *strict* environment, and for each rule, there is a *non-repeating* public menu mechanism and *preferential* convention that is compatible with the rule. Indeed, simply consider a public menu mechanism derived from the direct mechanism, where the agents sequentially reveal their full preference rankings. It follows immediately from this observation, the classic theorem that *Maskin monotonicity* is necessary for full **NE**-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 public menu mechanism that is both an ex-post perfect implementation of the rule and a full subgame perfect implementation of the rule.*

Corollary 2.2: *For each rich and strict environment, each strategy-proof and non-bossy 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 and a full subgame perfect implementation of the rule.*

We remark that because the latter corollary involves *non-bossiness* and a unique subgame perfect equilibrium outcome, it is conceptually related to Schummer and Velez (2019).[13] Altogether, Corollary 2.1 provides a useful link between the classic literature on Nash implementation and public menu mechanisms, while Corollary 2.2 allows us to

[13] In particular, Schummer and Velez (2019) 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 Swingels (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.

more clearly investigate the logical relationships of the hypotheses and conclusions of our results:

Necessity. A hypothesis is logically *necessary (for the conclusion)* if and only if it is an implication of the conclusion. For Theorem 1, both *strategy-proofness* and *f-compatibility* (of the appropriate objects) are necessary, as discussed earlier. The rest of the assumptions are not: the English auction demonstrates that *richness* is not necessary, the menu mechanism derived from male-proposing deferred acceptance where the agents are the women demonstrates that the *non-repeating* requirement is not necessary (Kawase and Bando, 2018), and it is easy to construct an example showing that the *preferential* requirement is not necessary. Indeed, we provide sufficient conditions for implementation, but do not describe all possible implementations. For Corollary 2.2, it is easy to show that the necessity of each of these assumptions is unchanged, as well as to show that *strictness* is not necessary, but we do not know whether or not *non-bossiness* is necessary; we leave this as an open question:

Conjecture: *For each environment, each rule, each public menu mechanism, and each convention, if the public menu mechanism is both an ex-post perfect implementation of the rule and a full subgame perfect implementation of the rule, then the rule is non-bossy.*

Tightness. A hypothesis is *indispensable (for the theorem)* if and only if when the theorem is modified by dropping this hypothesis, it becomes false.[14] For Theorem 1, it is easy to construct examples showing that (the appropriate object) being *strategy-proof*, *preferential*, and *f-compatible* are each indispensable. In fact, being *rich* and *non-repeating* are also indispensable (Figure 3). For Corollary 2.2, it is easy to show that the indispensability of each of these assumption is unchanged using Figure 3, whose examples still apply. It is also easy to show that being *strict* and *non-bossy* are indispensable. Altogether, all assumptions in both results are indispensable, so we say both results are *tight*.

Taken together, Theorem 1 and Theorem 2 show that public menu mechanisms provide an interesting alternative to direct mechanisms. Public menu mechanisms only require a limited amount of information to be transmitted by the agents, but they only implement desired outcomes in terms of subgame perfect equilibrium, which is closely related to common knowledge of rationality (Aumann, 1995): if agents do not trust that the others will conform to the convention, or at least pursue their own best interest, then undesirable outcomes can occur. By contrast, direct mechanisms require all information to be transmitted by the agents, but they implement outcomes in dominant strategy equilibrium: the agents need not have any faith in the rationality of their peers. One summary is that public menu mechanisms are desirable when agents have optimism about each other but pessimism about the mechanism administrator, while direct mechanisms are desirable when agents have pessimism about each other but optimism about the mechanism

[14]To avoid confusion: for hypotheses H_i and conclusion C, consider a theorem 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.

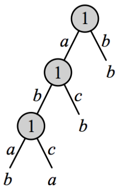

(a) *Non-repeating*: indispensable.

Rule	*efacbd*	*fedcba*
abcdef	(a,a)	(c,c)
badcef	(b,b)	(d,d)
fcedba	(e,e)	(f,f)

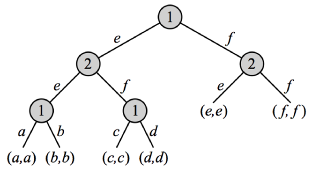

(b) *Rich*: indispensable.

Figure 3: Indispensable conditions. *Non-repeating.* Consider $N = \{1\}$, $X_1 = \{a, b, c\}$, the types are the strict rankings, the rule specifies that b is always selected, and the public menu mechanism in Figure 1a. *Rich.* Consider $N = \{1, 2\}$, $X_1 = X_2 = \{a, b, c, d, e, f\}$, and Figure 1b: Θ_1 is the table rows, Θ_2 is the table columns, the rule is given by the table entries, and the public menu mechanism is given. Each type specifies a strict ranking, most preferred to least preferred from left to right. The *preferential* convention is not ex-post perfect because for type profile $(fcedba, efacbd)$, agent 2 can profitably deviate from the conventional strategy profile by selecting f at the history after 1 selects e.

administrator.

4 Private menu mechanisms

In this section, we consider private menu mechanisms. We begin by providing conditions under which a private menu mechanism provides an ex-post **EDE**-implementation. In fact, in addition to the conditions introduced in the previous section, we only require one new condition:

Definition: Fix an environment. A private menu mechanism is *reaction-proof* if and only if for each distinct pair $i, j \in N$, each pair $s_i, s'_i \in S_i$, each $s_{-i} \in S_{-i}$, and each pair $h, h' \in H_j$ such that

- $h \in \pi(s_i, s_{-i})$ and $h' \in \pi(s'_i, s_{-i})$;
- j plays the same number of times before h and h', and
- $X_j(h) \neq X_j(h')$,

we have

- for each pair $\pi_1, \pi_2 \in \Pi$ such that $h \in \pi_1 \cap \pi_2$, $\mathcal{X}_i(\pi_1) = \mathcal{X}_i(\pi_2)$, and
- for each pair $\pi'_1, \pi'_2 \in \Pi$ such that $h' \in \pi'_1 \cap \pi'_2$, $\mathcal{X}_i(\pi'_1) = \mathcal{X}_i(\pi'_2)$.

Reaction-proofness requires that whenever an agent i signals something to another agent j—in the sense that if j knows that the others are playing s_{-i}, then he can distinguish between two strategies i may have played by looking at his menu—then by the time j can make this distinction, i is already safe from any reaction in the sense that his assignment is already determined. In fact, all of our leading examples are *reaction-proof*. For illustration:

- Serial dictatorship private menu mechanisms are *reaction-proof* because an agent always determines his assignment when he chooses from a menu.

- Deferred acceptance private menu mechanisms are *reaction-proof* because an agent's strategy determines his first experience, his second experience (if there is one), and so on: the next menu (if there is one) is determined by his own strategy.

- Top trading cycles private menu mechanisms are *reaction-proof* because even though an agent can make deductions about the actions of his peers upon observing which assignments are removed from the next menu, these deductions only concern members of previous trading cycles, who have already exited.

Our third theorem states that *reaction-proofness* and our previous conditions for ex-post perfect implementation together guarantee a private menu mechanism provides an ex-post everywhere-dominant implementation:

THEOREM 3: *For each rich environment, each strategy-proof rule, each non-repeating and reaction-proof private menu mechanism, and each preferential convention that is compatible with the rule, the private menu mechanism is an ex-post everywhere-dominant implementation of the rule via the convention.*

Proof. We consider both the private menu mechanism and its associated public menu mechanism. For each $i \in N$, let S_i denote the set of strategies for i in the private menu mechanism (which is contained in the set of strategies for i in the public menu mechanism). Most arguments implicitly involve the private menu mechanism; we explicitly highlight arguments involving the public menu mechanism.

Assume, by way of contradiction, that we do not have an ex-post everywhere-dominant implementation via $\mathbb{S}$. Then there are $\theta \in \Theta$, $i \in N$, $h \in H_i$, and $s_{-i} \in S_{-i}$ such that i has a profitable deviation s'_i from $\mathbb{S}_i(\theta_i)$ at h when his type is θ_i. By Theorem 1, the public menu mechanism is an ex-post perfect implementation of the rule via $\mathbb{S}$; thus the restriction of s_{-i} to the given subgame is not specified by the convention for any type profile. Let j be the least-index peer of i whose restricted strategy is never specified by the convention. We claim that j has a strategy specified by the convention s^*_j such that in response to $(s^*_j, s_{-i,j})$, i has a profitable deviation from $\mathbb{S}_i(\theta_i)$ at h when his type is θ_i.

Define the plays $\pi \equiv \pi^h(\mathbb{S}_i(\theta_i), s_{-i})$ and $\pi' \equiv \pi^h(s'_i, s_{-i})$. We consider two cases:

CASE 1: There is $\theta^*_j \in \Theta_j$ such that j adheres to both π and π' after h according to convention—that is, both (i) for each $h^* \in H_j \cap \pi$ such that $h \prec h^*$, $[\mathbb{S}_j(\theta^*_j)](h^*) = \alpha^{h^*}(\pi)$; and (ii) for each $h^* \in H_j \cap \pi'$ such that $h \prec h^*$, $[\mathbb{S}_j(\theta^*_j)](h^*) = \alpha^{h^*}(\pi')$. In this case, simply define $s^*_j \equiv \mathbb{S}_j(\theta^*_j)$.

CASE 2: There is no $\theta^*_j \in \Theta_j$ such that j adheres to both π and π' after h according to convention. Then let h_1, h_2, ... label the histories where j plays after h along π in order, and let h'_1, h'_2, ... label the histories where j plays after h along π' in order. We claim that there is an earliest $t \in \mathbb{N}$ such that $X_j(h_t) \neq X_j(h'_t)$. Indeed, if not, then 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 we can do by the Play Lemma). Then for each $t \in \mathbb{N}$ such that there are both h_t and h'_t, (i) $s_j(h_t) = s_j(h'_t)$, as s_j is a strategy in the private menu mechanism; and (ii) $[\mathbb{S}_j(\theta^*_j)](h_t) = [\mathbb{S}_j(\theta^*_j)](h'_t)$, because the convention is *preferential*; so necessarily both s_j and $\mathbb{S}_j(\theta_j^*)$ agree along π and π' after h, so θ_j^* is such that j adheres to both π and π' after h according to convention, contradicting that there is no such type. Let $t^{\neq}$ denote this earliest $t \in \mathbb{N}$ such that $X_j(h_t) \neq X_j(h_t')$.

Let $\theta_j^* \in \Theta_j$ be such that j adheres to π after h according to convention (which we can do by the Play Lemma) and define $s_j^* \equiv \mathbb{S}_j(\theta_j^*)$. We claim

- $\mathcal{X}_i(\pi^h(\mathbb{S}_i(\theta_i), s_{-i})) = \mathcal{X}_i(\pi^h(\mathbb{S}_i(\theta_i), s_j^*, s_{-i,j}))$, and
- $\mathcal{X}_i(\pi^h(s_i', s_{-i})) = \mathcal{X}_i(\pi^h(s_i', s_j^*, s_{-i,j}))$.

For the first item, the two plays are equivalent by construction of s_j^*, so we are done. For the second item, if the two plays are distinct, then as any strategy of 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 $\pi^h(s_i', s_j^*, s_{-i,j})$ lead from h to $h'_{t^{\neq}}$. By *reaction-proofness*, the assignment of i is determined after both $h_{t^{\neq}}$ and $h'_{t^{\neq}}$, so we are done.

In both cases, j has a strategy specified by the convention s_j^* such that in response to $(s_j^*, s_{-i,j})$, i has a profitable deviation from $\mathbb{S}_i(\theta_i)$ at h when his type is θ_i. But then repeating the argument, i has a profitable deviation when all his peers play a restricted strategy profile specified by the convention, contradicting that the public menu mechanism is an ex-post perfect implementation of the rule via the convention. ■

Finally, we consider conditions under which a private menu mechanism is not only an ex-post everywhere-dominant implementation, but moreover a full everywhere-dominant implementation. Our fourth theorem states that *reaction-proofness* and our previous conditions for double implementation with subgame perfection together guarantee this:

THEOREM 4: *For each rich and strict environment, each group strategy-proof rule, each non-repeating and reaction-proof private menu mechanism, and each preferential convention that is compatible with the rule, the private menu mechanism is both an ex-post everywhere-dominant implementation of the rule and a full everywhere-dominant implementation of the rule.*

Proof. By Theorem 3, 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 private menu mechanism, so i has a profitable deviation from s^* at h, contradicting that s^* is an everywhere-dominant equilibrium in the private 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 Theorem 2, $\mathcal{X}(\mathbb{S}(\theta)) = \mathcal{X}(s)$, contradicting that $\mathcal{X}(\mathbb{S}(\theta)) \neq \mathcal{X}(s)$. ■

We are not sure if *reaction-proofness* is necessary or indispensable for these results. That said, there are private menu mechanisms that (i) violate *reaction-proofness*, and

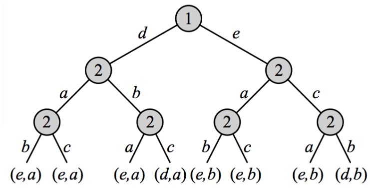

Figure 4: *Ex-post perfect but not ex-post everywhere-dominant.* Consider $N = \{1, 2\}$, $X_1 = \{d, e\}, X_2 = \{a, b, c\}$, the types are the strict rankings, the rule specifies that (i) if 2 ranks a below both b and c, then 1 gets d; else 1 gets e; and (ii) if 1 prefers d to e, then 2 gets a; else 2 gets b. Finally, consider the menu mechanism in Figure 4, which is both public and private. The preferential convention is an ex-post perfect implementation, but not an ex-post everywhere-dominant implementation; in fact, at no type profile does 1 have a dominant strategy.

(ii) are ex-post perfect implementations but not ex-post everywhere-dominant implementations (Figure 4). Taken together, Theorem 3 and Theorem 4 show that private menu mechanisms can provide unambiguous improvements over direct mechanisms according to the criteria considered in this article, preserving dominant strategy implementation (in a particularly robust way that preserves incentives even off-path) while reducing the information transmitted by agents.

5 Applications

We conclude by investigating implications of our main results across some familiar models. In contrast to the introduction, in this section we begin with more general examples to facilitate discussion across examples.

5.1 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, Kominers, and Westkamp, 2019). It is easy to verify that our other conditions are satisfied, and thus we have:

PROPOSITION 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.

It is easy to verify that these menu mechanisms improve upon the privacy of direct mechanisms: after he is hired, a doctor does not need to reveal his ranking of all the contracts he never proposed. 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.

5.2 Labor market

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*, $f^{SAP|\mathcal{C}}$ (Crawford and Knoer, 1981; Kelso and Crawford, 1982), which in fact coincides with the restriction of $f^{COP|\mathcal{C}}$ 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:

PROPOSITION 2: Fix a labor market environment and a profile of choice correspondences $\mathcal{C}$ 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 $f^{SAP|\mathcal{C}}$ via the preferential convention; and
- $G^{private}$ is an ex-post everywhere-dominant implementation of $f^{SAP|\mathcal{C}}$ 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 Proposition 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. ∎

It is easy to verify that these menu mechanisms improve upon the privacy of direct mechanisms: after he is hired, a worker does not need to reveal his reservation salary.

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 multi-item auction with unit demand where the objects are not identical. It is standard in discrete settings to assume that each buyer is indifferent between exiting and receiving the object at one of the prices, but in our variant model this is never the case, which is a natural assumption when the bid increment is large. 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.

If, moreover, there is a single buyer, then this can be interpreted as the auctioneer. 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.[15] Our results show that these insights extend to our variant model with no buyer indifference: the English 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.

5.3 College admissions

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

[15] 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).

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). In this case, $f^{DA|\mathcal{C}}$ is moreover *group strategy-proof* if and only if it is *efficient* if and only if it is *consistent* if and only if $\mathbf{p}$ is *acyclic* (Ergin, 2002). It is easy to verify that our other conditions are satisfied, and thus we have:

Proposition 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 $f^{DA|\mathbf{p}}$.

It is easy to verify that these menu mechanisms improve upon the privacy of direct mechanisms: after he is enrolled, a student does not need to reveal his ranking of the schools he did not apply to. Observe that the first and second item of Proposition 3 can be derived from Proposition 1, as these school choice environments are matching with contracts environments that satisfy the hypotheses of the proposition (Hatfield and Milgrom, 2005; Hatfield, Kominers, and Westkamp, 2019). As discussed earlier, for the first item of Proposition 3, similar insights were recently obtained by Kawase and Bando (2018) and Bó and Hakimov (2019).

Analogues of Proposition 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, 2005; Hafalir, Yenmez, and Yildrim, 2013; Ehlers, Hafalir, Yenmez, and Yildrim, 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).

5.4 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. It is easy to verify that our other conditions are satisfied, and thus we have:

Proposition 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 subgame perfect implementation of $f^{TC|\omega}$.

It is easy to verify that these menu mechanisms improve upon the privacy of direct mechanisms: an agent does not need to reveal how he ranks the objects that he desires less than his assignment (unless he is a broker). Observe that in this model, every *efficient* and *group strategy-proof* rule has a double implementation in terms of everywhere-dominant equilibrium. Finally, we remark that each serial dictatorship rule can clearly be implemented in these ways through a menu mechanism that is minimally informative among all implementations.

5.5 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, Sonnenschein, and Zhou, 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:

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

In this case, menu mechanisms offer no privacy improvements over direct mechanisms; in both cases, the election gathers all voter preference information. 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).

Appendix A: 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
- $\mathcal{C} : \mathcal{M} \to A$ 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 $a \rhd^{\mathcal{C}} b$, if and only if there is $M \in \mathcal{M}$ such that $a, b \in M$ and $\mathcal{C}(M) = a$. A list $(a_1, a_2, ..., a_k) \in$

$\cup_{\kappa\in\mathbb{N}\setminus\{1\}}A^\kappa$ is a $\mathcal{C}$*-path (from* a_1 *to* a_k*)* if and only if for each $t \in \{1, 2, ..., k-1\}$, $a_t \rhd^{\mathcal{C}} a_{t+1}$, and is moreover a $\mathcal{C}$*-cycle (from* a_1 *to* a_1*)* if and only if $a_k = a_1$.[16]

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 is a direct corollary of Proposition 2.7 in 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, $(a_1, a_2, ..., a_k) \in \cup_{\kappa\in\mathbb{N}\setminus\{1\}}A^\kappa$, and let $P \in \mathcal{P}_\mathcal{C}$. It follows by definition that for each $t \in \{1, 2, ..., k-1\}$, 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 $\mathcal{C}^* : \mathcal{M}^* \to A$ 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 $(a_1, a_2, ..., a_k) \in \cup_{\kappa\in\mathbb{N}\setminus\{1\}}A^\kappa$. Since $\mathcal{C}$ is rationalizable, thus by the Cycle Lemma there are no $\mathcal{C}$-cycles, so in particular $(a_1, a_2, ..., a_k)$ is not a $\mathcal{C}$-cycle; thus there is $t \in \{1, 2, ..., k-1\}$ such that $a_t = b$ and $a_{t+1} = a$. Using the members of $(a_1, a_2, ..., a_k)$, we can construct a $\mathcal{C}^*$-cycle from a to a, $(a'_1, a'_2, ..., a'_{k'})$, such that (i) for each $t \in \{2, 3, ..., k'-1\}$, $a'_t \neq a$, and (ii) $a'_{k'-1} = b$. For each $t \in \{1, 2, ..., k'-2\}$, $a'_t \rhd^{\mathcal{C}^*} a'_{t+1}$ and $(a'_t, a'_{t+1}) \neq (b, a)$, so $a'_t \rhd^{\mathcal{C}} a'_{t+1}$; but then $(a'_1, a'_2, ..., a'_{k'-1})$ is a $\mathcal{C}$-path from a to b, contradicting that there is no $\mathcal{C}$-path from a to b.

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

[16]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 at least three members whose first $k-1$ members are distinct. We do not do so here: these additional requirements have no relevance to our proofs, and more importantly we want to remain consistent with Chambers and Echenique (2016) to use one of their results.

Appendix B: 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 $(a_1, a_2, ..., a_k) \in \cup_{\kappa \in \mathbb{N} \setminus \{1\}} X_i^\kappa$. For each $t \in \{1, 2, ..., k-1\}$, $a_t \rhd^{\mathcal{C}} a_{t+1}$, 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 $h_\prec$ which precedes the others. Since $a_1 = a_k$, thus there is $t^* \in \{1, 2, ..., k-1\}$ such that $\mathcal{C}(X_i(h_\prec)) = a_{t^*+1}$. Since $\mathcal{C}(X_i(h_{t^*})) = a_{t^*} \neq a_{t^*+1}$, thus $h_\prec \neq h_{t^*}$, so $h_\prec \prec h_{t^*}$. But then there is $\pi' \in \Pi$ such that $h_\prec, h_{t^*} \in \pi'$ and $\alpha^{h_\prec}(\pi') = \alpha^{h_{t^*}}(\pi') = a_{t^*+1}$, contradicting that G is *non-repeating*.

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 $R_i(\theta_i) = P$. Since $R_i(\theta_i)$ is strict (and therefore requires no tie-breaking), thus for each $h \in H_i \cap \pi$, we have

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

as desired. ■

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 $s_i\restriction_h$ to be restriction of s_i to $\{h' \in H_i | h' \succsim h\}$; 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 $\{h' \in H_i | h' \prec h\} = \emptyset$ then we are done; thus let us assume $\{h' \in H_i | h' \prec h\} \neq \emptyset$. Define $H_{\succsim} \subseteq H_i$ and the choice space $(X_i, \mathcal{M}_{\succsim}, \mathcal{C}_{\succsim})$ by:

- $H_{\succsim} \equiv \{h' \in H_i | h' \succsim h\}$;
- $\mathcal{M}_{\succsim} \equiv \{X_i(h')\}_{h' \in H_{\succsim}}$; 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:

- $H_{\prec} \equiv \{h' \in H_i | h' \prec h\}$;
- $\mathcal{M} \equiv \{X_i(h')\}_{h' \in H_{\prec} \cup H_{\succsim}}$; 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 $(a_1, a_2, ..., a_k) \in \cup_{\kappa \in \mathbb{N}\setminus\{1\}} X_i^{\kappa}$. For each $t \in \{1, 2, ..., k-1\}$, $a_t \rhd^{\mathcal{C}} a_{t+1}$, 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 $t^* \in \{1, 2, ..., k-1\}$ such that $h_{t^*} \in H_{\prec}$; it follows that $\{h_t\}$ has a member $h_{\prec}$ which precedes the others. From here, the argument 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 $h' \in H_{\prec} \cup H_{\succsim}$,

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

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. ■

Appendix C: Proof of Theorem 1

In this appendix, we prove Theorem 1.

Theorem 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

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

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 $H_\prec \equiv \{h' \in H_i | h' \prec h\}$. 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 $j \in N \backslash \{i\}$, 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) $\mathbb{S}_j(\theta_j^*)\restriction_h = \mathbb{S}_j(\theta_j)\restriction_h$.

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*,

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

so $a\ P_a\ b$. By a symmetric argument, $b\ P_b\ a$. Since $P_a \in \mathcal{P}_{\mathcal{C}_a}$ and $P_b \in \mathcal{P}_{\mathcal{C}_b}$ were arbitrary, thus by the Path Lemma, there are (i) a $\mathcal{C}_a$-path from a to b, $(a_1, a_2, ..., a_k) \in \cup_{\kappa \in \mathbb{N} \backslash \{1\}} X_i^\kappa$, and (ii) a $\mathcal{C}_b$-path from b to a, $(b_1, b_2, ..., b_{k'}) \in \cup_{\kappa \in \mathbb{N} \backslash \{1\}} X_i^\kappa$. For each $t \in \{1, 2, ..., k-1\}$, $a_t \rhd^{\mathcal{C}_a} a_{t+1}$, 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 $\mathcal{C}_a(X_i(h_t^a)) = a_t$. Similarly, for each $t \in \{1, 2, ..., k'-1\}$, $b_t \rhd^{\mathcal{C}_b} b_{t+1}$, 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 $\mathcal{C}_b(X_i(h_t^b)) = b_t$.

Assume, by way of contradiction, $\{h_t^a\} \cap H_\prec \neq \emptyset$. Since $\{h_t^a\} \subseteq \pi_a$, thus it has a member $h_\prec^a$ which precedes the others. It must be that $h_\prec^a = h_1^a$; else there is $t \in \{1, 2, ..., k-2\}$ such that $\mathcal{C}_a(X_i(h_\prec^a)) = a_{t+1}$ and $\mathcal{C}(X_i(h_t^a)) = a_t \neq a_{t+1}$, so $h_\prec^a \neq h_t^a$, so $h_\prec^a \prec h_t^a$, so there is $\pi \in \Pi$ such that $h_\prec^a, h_t^a \in \pi$ and $\alpha^{h_\prec^a}(\pi) = \alpha^{h_t^a}(\pi) = a_{t+1}$, contradicting that G is *non-repeating*. By the same argument, since $\{h_t^b\} \subseteq \pi_b$, thus it has a member $h_\prec^b$ which precedes the others, and necessarily $h_\prec^b = h_1^b$. Since $\{h_t^a\} \cap H_\prec \neq \emptyset$, thus $h_1^a \in H_\prec$, 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 $h_1^a \prec h_{k'-1}^b$ or $h_{k'-1}^b \prec h_1^a$. It must be that $h_{k'-1}^b \prec h_1^a$; else there is $\pi' \in \Pi$ such that

$h^a_1, h^b_{k'-1} \in \pi'$ and $\alpha^{h^a_1}(\pi') = \alpha^{h^b_{k'-1}}(\pi') = a$, contradicting that G is *non-repeating*. Then $h^b_1 \prec h^b_{k'-1} \prec h^a_1$. But by a symmetric argument, $h^a_1 \prec h^b_1$, contradicting that $h^b_1 \prec h^a_1$.

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

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

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:

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

Since $\pi_a = \pi^h(\mathbb{S}(\theta))$, thus $P_{\theta_i} \in \mathcal{P}_\mathcal{C}$, so $\mathcal{C}$ is rationalizable. Since $\{h^a_t\} \cap H_\prec = \emptyset$, thus $\{h^a_t\} \subseteq (H_i \cap \pi_a) \setminus H_\prec$, 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 $\mathbb{S}(\theta) \in \textbf{SPE}(G, R(\theta))$. 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 $\mathbb{S}$. ■

Appendix D: Proof of Theorem 2

In this appendix, we prove Theorem 2.

THEOREM 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 and a full subgame perfect implementation of the rule.*

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

Let $\theta \in \Theta$, let $s^* \in \textbf{SPE}(G, R(\theta))$, and define $s^\theta \equiv \mathbb{S}(\theta)$. By Theorem 1, $s^\theta \in \textbf{SPE}(G, R(\theta))$ 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, $\ell(h) \equiv \max_{\pi \in \Pi} |\pi \cap \{h' \in H | h' \succsim h\}|$; 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

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

Since $s^\theta \in \textbf{SPE}(G, R(\theta))$, 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 (i) for each $i \in N$, X_i is finite; and (ii) G is *non-repeating*; 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. ■

Appendix E: Definition of mechanism

In this appendix, we formally define *mechanisms*:

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

- H is the set of *histories* and $\precsim$ is the partial order on H representing *precedence*. We require that $(H, \precsim)$ 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.

- $\mathbb{P} : H \backslash Z \to N$ 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 $H_i \equiv \{h \in H \backslash Z | \mathbb{P}(h) = i\}$ denote the histories that belong to i.

- $\mathcal{A}$ is the set of *actions* and $\alpha : \cup_{H \backslash Z} \sigma(h) \to \mathcal{A}$ 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

[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 (2019).

available action determines a unique next history: for each $h \in H \backslash Z$ and each pair $h', h'' \in \sigma(h)$, $\alpha(h') \neq \alpha(h'')$. For each non-terminal history h, we let $\mathcal{A}(h) \equiv \{\alpha(h') | h' \in \sigma(h)\}$ 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 $h \prec h'$, 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.
- $\mathcal{X} : \Pi \to X$ 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.

References

Abdulkadiroğlu, Atila (2005). "College admissions with affirmative action." *International Journal of Game Theory* 33, 535-549.

Abdulkadiroğlu, Atila and Sönmez, Tayfun (2003). "School Choice: A Mechanism Design Approach." *American Economic Review* 93, 729-747.

Akbarpour, Mohammad and Li, Shengwu (2020). "Credible Mechanisms." *Econometrica* 88, 425-467.

Alós-Ferrer, Carlos and Ritzberger, Klaus (2016). *The Theory of Extensive Form Games.* Berlin, Germany: Springer-Verlag Berlin Heidelberg.

Aumann, Robert (1995). "Backward induction and common knowledge of rationality." *Games and Economic Behavior* 8, 6-19.

Ausubel, Lawrence (2004). "An Efficient Ascending-Bid Auction for Multiple Objects." *American Economic Review* 94, 1452-1475.

Ausubel, Lawrence (2006). "An Efficient Dynamic Auction for Heterogeneous Commodities." *American Economic Review* 96, 602-629.

Bade, Sophie (2020). "Random Serial Dictatorship: The One and Only." *Mathematics of Operations Research* 45, 353-368.

Barberà, Salvador, Sonnenschein, Hugo, and Zhou, Lin (1991). "Voting by Committees." *Econometrica* 59, 595-609.

Bergemann, Dirk and Hörner, Johannes (2018). "Should First-Price Auctions Be Transparent?" *American Economic Journal: Microeconomics* 10, 177-218.

Bergemann, Dirk and Morris, Stephen (2005). "Robust mechanism design." *Econometrica* 73, 1771-1813.

Blackwell, David (1951). "Comparison of Experiments." In: *Proceedings of the Second Berkeley Symposium on Mathematical Statistics and Probability.* Editor: Neyman, Jerzy. Berkeley, California: University of California Press. Pages 93-102.

Bó, Inácio and Hakimov, Rustamdjan (2019). "The iterative deferred acceptance mechanism." Working paper.

Bó, Inácio and Hakimov, Rustamdjan (2020a). "Iterative Versus Standard Deferred Acceptance: Experimental Evidence." *Economic Journal* 130, 356-392.

Bó, Inácio and Hakimov, Rustamdjan (2020b). "Pick-an-object mechanisms." Working paper.

Chambers, Christopher and Echenique, Federico (2016). *Revealed Preference Theory.* Cambridge, United Kingdom: Cambridge University Press.

Crawford, Vincent and Knoer, Elsie (1981). "Job Matching with Heterogeneous Firms and Workers." *Econometrica* 49, 437-450.

Dubins, Lester and Freedman, David (1981). "Machiavelli and the Gale-Shapley Algorithm." *American Mathematical Monthly* 88, 485-494.

Dur, Umut, Kominers, Scott, Pathak, Parag, and Sönmez, Tayfun (2018). "Reserve Design: Unintended Consequences and the Demise of Boston's Walk Zones." *Journal of Political Economy* 126, 2457-2479.

Ehlers, Lars, Hafalir, Isa, Yenmez, M. Bumin, and Yildrim, Muhammed (2014). "School choice with controlled choice constraints: Hard bounds versus soft bounds." *Journal of Economic Theory* 153, 648-683.

Eliaz, Kfir, Eilat, Ran, and Mu, Xiaosheng (2019). "Optimal Privacy-Constrained Mechanisms." Working paper.

Ergin, Haluk (2002). "Efficient resource allocation on the basis of priorities." *Econometrica* 70, 2489-2497.

Gale, David and Shapley, Lloyd (1962). "College admissions and the stability of marriage." *American Mathematical Monthly* 69, 9-15.

Gershkov, Alex, Moldovanu, Benny, and Shi, Xianwen (2017). "Optimal Voting Rules." *The Review of Economic Studies* 84, 688-717.

Gibbard, Allan (1973). "Manipulation of voting schemes: a general result." *Econometrica* 41, 587-601.

Glazer, Jacob and Rubinstein, Ariel (1996). "An Extensive Game as a Guide for Solving a Normal Game." *Journal of Economic Theory* 70, 32-42.

Gradwohl, Ronen and Smorodinsky, Rann (2017). "Perception games and privacy." *Games and Economic Behavior* 104, 293-308.

Hafalir, Isa, Yenmez, M. Bumin, and Yildrim, Muhammed (2013). "Effective affirmative action in school choice." *Theoretical Economics* 8, 325-363.

Hatfield, John and Kojima, Fuhito (2009). "Group incentive compatibility for matching with contracts." *Games and Economic Behavior* 67, 745-749.

Hatfield, John, Kominers, Scott, and Westkamp, Alexander (2019). "Stability, Strategy-Proofness, and Cumulative Offer Mechanisms." Forthcoming, *Review of Economic Studies.*

Hatfield, John and Milgrom, Paul (2005). "Matching with Contracts." *American Economic Review* 95, 913-935.

Hurwicz, Leonid (1960). "Optimality and Informational Efficiency in Resource Allocation Processes." In: *Mathematical Methods in the Social Sciences, 1959: Proceedings of the First Stanford Symposium.* Editors: Arrow, Kenneth, Karlin, Samuel, and Suppes, Patrick. Stanford, California: Stanford University Press.

Jung et al. versus Association of American Medical Colleges et al., 02-CV-00873 (DDC 2002).

Kamada, Yuichiro and Kojima, Fuhito (2015). "Efficient Matching under Distributional Constraints: Theory and Applications." *American Economic Review* 105, 67-99.

Kamada, Yuichiro and Kojima, Fuhito (2018). "Stability and strategy-proofness for matching with constraints: A necessary and sufficient condition." *Theoretical Economics* 13, 761-793.

Kawase, Yasushi and Bando, Keisuke (2018). "Subgame Perfect Equilibria under the Deferred Acceptance Algorithm." Working paper.

Kelso, Alexander and Crawford, Vincent (1982). "Job Matching, Coalition Formation, and Gross Substitutes." *Econometrica* 50, 1483-1504.

Kleiner, Andreas and Moldovanu, Benny (2017). "Content-Based Agendas and Qualified Majorities in Sequential Voting." *American Economic Review* 107, 1477-1506.

Kleiner, Andreas and Moldovanu, Benny (2019). "Abortions, Brexit, and Trees." Working paper.

Klijn, Flip, Pais, Joana, and Vorsatz, Marc (2019). "Statis versus dynamic deferred acceptance in school choice: Theory and experiment." *Games and Economic Behavior* 113, 147-163.

Kojima, Fuhito (2010). "Impossibility of stable and nonbossy matching mechanisms." *Economics Letters* 107, 69-70.

Li, Shengwu (2017). "Obviously Strategy-Proof Mechanisms." *American Economic Review* 107, 3257-3287.

Mackenzie, Andrew (2019). "A revelation principle for obviously strategy-proof implementation." Working paper.

Marx, Leslie and Swinkels, Jeroen (1997). "Order Independence for Iterated Weak Dominance." *Games and Economic Behavior* 18, 219-245.

Maskin, Eric (1979). "Incentive schemes immune to group manipulation." Mimeo.

Maskin, Eric (1999). "Nash Equilibrium and Welfare Optimality." *The Review of Economic Studies* 66, 23-38.

Milgrom, Paul and Segal, Ilya (2020). "Clock Auctions and Radio Spectrum Reallocation." *Journal of Political Economy* 128, 1-31.

Moore, John and Repullo, Rafael (1988). "Subgame perfect implementation." *Econometrica* 56, 1191-1220.

Pápai, Szilvia (2000). "Strategyproof assignment by hierarchical exchange." *Econometrica* 68, 1403-1433.

Pycia, Marek and Troyan, Peter (2019). "A Theory of Simplicity in Games and Mechanism Design." Working paper.

Pycia, Marek and Ünver, Utku (2017). "Incentive compatible allocation and exchange of discrete resources." *Theoretical Economics* 12, 287-329.

Roth, Alvin (1982). "The Economics of Matching: Stability and Incentives." *Mathematics of Operations Research* 7, 617-628.

Roth, Alvin and Peranson, Elliott (1999). "The Redesign of the Matching Market for American Physicians: Some Engineering Aspects of Economic Design." *American Economic Review* 89, 748-780.

Roth, Alvin and Xing, Xiaolin (1997). "Turnaround Time and Bottlenecks in Market Clearing: Decentralized Matching in the Market for Clinical Psychologists." *Journal of Political Economy* 105, 284-329.

Rothkopf, Michael, Teisberg, Thomas, and Kahn, Edward (1990). "Why are Vickrey Auctions Rare?" *Journal of Political Economy* 98, 94-109.

Saijo, Tatsuyoshi, Sjöström, Tomas, and Yamato, Takehiko (2007). "Secure implementation." *Theoretical Economics* 2, 203-229.

Satterthwaite, M. (1975). "Strategy-proofness and Arrow's conditions: existence and correspondence theorems for voting procedures and social welfare functions." *Journal of Economic Theory* 10, 187-217.

Satterthwaite, Mark and Sonnenschein, Hugo (1981). "Strategy-Proof Allocation Mechanisms at Differentiable Points." *The Review of Economic Studies* 48, 587-597.

Schummer, James and Velez, Rodrigo (2019). "Sequential Preference Revelation in Incomplete Information Settings." Forthcoming, *American Economic Journal: Microeconomics.*

Shapley, Lloyd and Scarf, Herbert (1974). "On cores and indivisibility." *Journal of Mathematical Economics* 1, 23-37.

Sun, Ning and Yang, Zaifu (2014). "An Efficient and Incentive Compatible Dynamic Auction for Multiple Complements." *Journal of Political Economy* 112, 422-466.

Svensson, Lars-Gunnar (1999). "Strategy-proof allocation of indivisible goods." *Social Choice and Welfare* 16, 557-567.

Takamiya, Koji (2001). "Coalition strategy-proofness and monotonicity in Shapley-Scarf housing markets." *Mathematical Social Sciences* 41, 201-213.

Thomson, William (2016). "Non-bossiness." *Social Choice and Welfare* 47, 665-696.

Vickrey, William (1961). "Counterspeculation, Auctions, and Competitive Sealed Tenders." *The Journal of Finance* 16, 8-37.

Yamato, Takehiko (1993). "Double implementation in Nash and undominated Nash equilibria." *Journal of Economic Theory* 59, 311-323.