Strategyproofness in the Large as a Desideratum for Market Design*

Eduardo M. Azevedo† and Eric Budish‡

October 14, 2011

Abstract

We distinguish between two ways a mechanism can fail to be strategyproof (SP). A mechanism may have profitable manipulations that *persist with market size*; and, a mechanism may have profitable manipulations that *vanish with market size*. We say that a non-SP mechanism is *strategyproof in the large (SP-L)* if all of its profitable manipulations vanish with market size.

Our main result is as follows. Suppose we are given some mechanism that has Bayes-Nash equilibria but is not SP-L; then, under some commonly satisfied conditions (semi-anonymity, private values, quasi-continuity) we show by construction that there exists another mechanism that is SP-L, and that implements approximately the same outcomes as the original mechanism, with the approximation error vanishing in the large-market limit. Thus, while SP often severely limits what kinds of mechanisms are possible, SP-L is approximately costless, and hence may be a useful second-best. We illustrate with examples from assignment, matching and auctions.

Working Draft. Preliminary and Incomplete.

*For helpful discussions we are grateful to Nabil Al-Najjar, Susan Athey, Aaron Bodoh-Creed, Gabriel Carroll, Jeff Ely, Drew Fudenberg, Jason Hartline, John Hatfield, Richard Holden, Ehud Kalai, Emir Kamenica, Fuhito Kojima, Scott Kominers, Jacob Leshno, Jon Levin, Paul Milgrom, Roger Myerson, David Parkes, Parag Pathak, Nicola Persico, Canice Prendergast, Ilya Segal, Eran Shmaya, Lars Stole, Glen Weyl, and especially Al Roth.

†Harvard University, azevedo@fas.harvard.edu.

‡University of Chicago Booth School of Business, eric.budish@chicagobooth.edu.

1 Introduction

Strategyproofness – i.e., that playing the game truthfully is a dominant strategy – is perhaps the predominant notion of incentive compatibility in practical market design. There are at least five important reasons why strategyproofness is so heavily emphasized relative to other forms of incentive compatibility, such as Bayes-Nash. First, strategyproof mechanisms are detail free for the designer, in the sense of Wilson (1987); the designer need not know anything about participants' preferences or beliefs (cf. Bergemann and Morris (2005)). Second, and relatedly, strategyproof mechanisms are strategically simple for participants; participants need not form beliefs about others' preferences or behavior in order to play the game optimally (Fudenberg and Tirole (1991); Roth (2008)). Third, with this simplicity comes a measure of fairness; agents' outcomes do not depend on their ability to "game the system" (Friedman (1991); Pathak and Sönmez (2008); Abdulkadiroğlu et al. (2006)). Fourth, strategyproof mechanisms generate information about participants' true preferences that may be useful to policy makers (Roth (2008)). Fifth, Bayesian approaches simply have not yet proved tractable for a number of important market design problems.

However, in a wide variety of economic contexts, impossibility theorems indicate that strategyproofness severely limits what kinds of mechanisms are possible. These range from Gibbard (1973) and Satterthwaite's (1975) dictatorship theorem for general social choice problems, to Hurwicz's (1972) impossibility theorem for general equilibrium settings, to the Green and Laffont (1977) VCG theorem for allocation settings with quasi-linear preferences, to Roth's (1982) impossibility theorem for strategyproof stable matching, to Papai's (2001) dictatorship theorem for multi-unit demand assignment problems, to Abdulkadiroğlu et al.'s (2009) impossibility theorem for strategyproof and efficient school assignment.

This creates a conundrum for market designers. Strategyproofness is the only form of incentive compatibility that the literature finds fully satisfying, yet often there are no good strategyproof mechanisms.

This paper proposes a criterion of approximate strategyproofness, and suggests that it may be a useful second-best alternative in environments where strategyproof mechanisms are unattractive. Our criterion is based on a conceptual distinction between two ways a mechanism might fail to be strategyproof. First, a mechanism might have profitable manipulations that *persist with market size*. Second, a mechanism might have profitable manipulations that *vanish with market size*. While both kinds of manipulability are undesirable, we suggest that manipulations that persist with market size

are especially problematic, and are avoidable. If a mechanism only has manipulations that vanish with market size, we will say that it is *strategyproof in the large (SP-L)*.

Whether or not a mechanism is SP-L is simple to check, and generates an intuitively appealing classification of non-strategyproof mechanisms. Many well-known non-strategyproof mechanisms that are thought to work well in practice are SP-L. Examples include the Walrasian mechanism, double auctions, uniform-price auctions, and deferred-acceptance algorithms. Many other non-strategyproof mechanisms that have been shown to have important incentives problems in practice are *not* SP-L, i.e., they are manipulable even in large markets. Examples include the pay-as-bid auction criticized by Friedman (1964, 1991), the Boston mechanism for school choice criticized by Abdulkadiroğlu and Sönmez (2003), the bidding points auction for course allocation criticized by Sönmez and Ünver (2010), and the priority-match algorithm for two-sided matching criticized by Roth (2002). Furthermore, both Friedman's critique of pay-as-bid auctions and Roth's critique of priority-match algorithms explicitly suggested alternative mechanisms that are not strategyproof but that are SP-L: uniform-price auctions and deferred-acceptance algorithms, respectively. This too speaks to the intuitive appeal of the criterion.

Our main result is as follows. Suppose we are given some mechanism that has Bayes-Nash equilibria. Suppose that the mechanism is (semi-)anonymous, which is a common feature of practical market-design settings; that agents have private values, in the sense that they know their own preferences over outcomes without observing other agents' private information; and that the mechanism satisfies a condition called quasi-continuity, which we will describe in more detail below. We show that there necessarily exists another mechanism that is strategyproof in the large, and that implements approximately the same outcomes as the original mechanism, with the approximation error vanishing in the large-market limit. An interpretation of our result is that while restricting attention to SP mechanisms can be very costly in terms of design objectives, restricting attention to SP-L mechanisms is approximately costless. This justifies consideration of SP-L as a second-best alternative to SP.

Our proof is by construction of a specific SP-L mechanism, from a given mechanism that has Bayes-Nash equilibria. The construction works as follows. Agents report their types to our mechanism. Our mechanism then calculates the empirical distribution of these types, and then "activates" the Bayes-Nash equilibrium strategy of the original mechanism associated with this empirical. If agents all report their preferences truthfully, this construction will yield the same outcome as the original

mechanism in the large-market limit, because the empirical distribution of reported types converges to the underlying true distribution. The subtle part of our construction is what happens if some agents systematically misreport their preferences, e.g., they make mistakes. Suppose the true prior is μ, but for some reason the agents other than agent i systematically misreport their preferences, according to distribution m. In a finite market, with sampling error, the empirical distribution of the other agents' reports is say $\hat{m}$. As the market grows large, $\hat{m}$ is converging to m, and also i's influence on the empirical distribution is vanishing. Thus in the limit, our construction will activate the Bayes-Nash equilibrium strategy associated with m. This is the "wrong" prior – but agent i does not care. From his perspective, the other agents are reporting according to m, and then playing the Bayes-Nash equilibrium strategy associated with m, so i too wishes to play the Bayes-Nash equilibrium strategy associated with m. This is exactly what our constructed mechanism does on i's behalf in the limit. Hence, no matter how the other agents play, i wishes to report his own type truthfully in the limit, i.e., the constructed mechanism is SP-L

Our construction resembles a revelation principle construction, in that it takes a mechanism in which agents play the game directly and transforms it into a mechanism in which agents just report their type, and then let the center play optimally on their behalf. However, we emphasize that our mechanism is fundamentally distinct. In a traditional revelation mechanism, the mechanism designer knows the true prior (e.g., μ), and then plays the Bayes-Nash equilibrium strategy associated with this true prior on agents' behalf. It is then a Bayes-Nash equilibrium for agents to report their types truthfully. Our mechanism has two advantages relative to this benchmark. First, it is prior free: neither the agents nor the mechanism need know the underlying distribution of preferences a priori, because the mechanism infers the prior from the empirical. Second, our mechanism provides dominant-strategy incentives in the limit, whereas a traditional revelation mechanism provides just Bayes-Nash incentives even in the limit.

A difficult technical issue throughout the analysis concerns points of discontinuity in a mechanism. The argument sketched above relied implicitly on an assumption of continuity local to m: as the empirical $\hat{m}$ is converging to m, agent i's utility is converging to what he would receive in the Bayes-Nash equilibrium associated with m. However, many familiar mechanisms have points at which agents' outcomes are not locally continuous. As an example, consider the uniform-price auction. Typically, a small change in the distribution of opponents' bids will have only a small effect on agent i's payoff. However, if i is the marginal bidder, a small change could discontinuously

cause i to change from being a winner of the auction to being a loser of the auction. Or, if prices are discrete and demand is exactly equal to supply at some price p, then a small decrease in demand could cause the market clearing price to decrease discontinuously.

Our analysis accommodates discontinuities in two related ways. First, our main result does not require that a mechanism be everywhere continuous, but rather that it satisfy a condition we call *quasi-continuity*. The quasi-continuity condition allows for the kinds of discontinuities that arise in the uniform-price auction. Roughly, the requirement is that discontinuities be "knife edge", in the sense that on either side of a discontinuity is a region where the mechanism is locally continuous. Second is the way we define SP-L itself. A mechanism is strategyproof if, for any profile of the other agents' reports, agent i maximizes his utility by reporting his own preferences truthfully. We say that a mechanism is SP-L if, in the large-market limit, for any *probability distribution* of the other agents' reports, agent i maximizes *expected* utility by reporting his preferences truthfully. When a mechanism is continuous, by a law of large numbers argument, there is no distinction in the limit between expected utility from a probability distribution of reports and realized utility from a specific profile of reports. If a mechanism has discontinuities, however, there can be such a distinction. For instance, in the uniform-price auction, an agent who reports her preferences truthfully might wish ex post to revise her report, in the event that the empirical realization of reports is exactly the knife-edge case where she can have a discontinuous influence on price. We classify the uniform-price auction as SP-L because the likelihood of this event vanishes with market size, for *any* probability distribution over the other agents' reports (cf. Example 1 below).

If we assume a stronger form of continuity, we can get stronger results. Specifically, if we assume that mechanisms are uniformly continuous as defined by Kalai (2004), then we can show that any SP-L mechanism has the property that, in a large enough market, no agent ever gains more than ϵ in any realization by misreporting her preferences. This is a stronger form of ex post robustness than that obtained by Kalai (2004) for Bayes-Nash equilibria, both because the likelihood of having an ϵ deviation is exactly zero in a large enough finite market rather than converging to zero in the limit, and because agents need not know the prior, coordinate on a specific equilibrium, etc.

Related Literature Our paper is related to a large literature that has studied how market size can ease incentive constraints. An early paper in this tradition is Roberts and Postlewaite (1976) on the Walrasian mechanism, which can be seen as a response to Hurwicz's (1972) critique that the Walrasian mechanism is not strategyproof.

Other papers in this tradition include Jackson and Manelli (1997), Kovalenkov (2002), and Al-Najjar and Smorodinsky (2007) on the Walrasian mechanism, Rustichini et al. (1994) on double auctions with private values, Pesendorfer and Swinkels (2000), Cripps and Swinkels (2006), and Reny and Perry (2006) on double auctions with common-value components, Immorlica and Mahdian (2005), Kojima and Pathak (2009), and Lee (2011) on deferred acceptance algorithms, and Kojima and Manea (2010) on the Bogomolnaia and Moulin (2001) probabilistic serial mechanism. Each of these papers provides a defense of a *specific mechanism* based on its incentive properties in large markets. Our paper aims to justify strategyproofness in the large as a *general desideratum* for practical market design. Note that in the context of any of the specific mechanisms named above, our analysis is much less instructive than are previous analyses tailored to the specific mechanism.

Technically, our paper is most closely related to Kalai (2004). Kalai (2004; Theorem 1) shows that Bayes-Nash equilibria are approximately ex post Nash in a class of large continuous and anonymous games.[1] In words, if a large number of agents with private information about their types play some BNE, then ex post – i.e., after seeing each agent's chosen action – agents will have vanishingly little incentive to revise their play. The difference between our results and Kalai's Theorem 1 is that Kalai shows that a given BNE is approximately ex post Nash, whereas we use the BNE of a given mechanism to create a new mechanism that is approximately strategyproof. In our new mechanism players need not have common knowledge of the prior, or of what equilibrium is being played, nor need they be strategically sophisticated in any way. There are several other well-known technical ideas that our paper is related to. First is the revelation principle (Myerson (1979)); see our discussion of how our main result is related to but distinct from the revelation principle in Section 4.2. Second, is the idea that there can be equivalence, in specialized environments, between what is implementable in Bayes-Nash equilibrium and what is implementable in dominant strategies. The revenue equivalence theorem in auction theory is an early example of such a result, since there exist dominant-strategy auctions that maximize revenue. See Manelli and Vincent (2010) for a recent equivalence result, and Gershkov et al. (2011) for a provocative discussion of these issues. Third is the idea of using the empirical distribution of agents' actions to infer the underlying distribution of preferences; see Segal (2003) for a paper that uses this idea in the context of monopoly pricing.

Next, our paper is related to the literature on the role of strategyproofness in practical market design.

[1] Recent work by Azrieli and Shmaya (2011) shows that continuity is the crucial assumption in Kalai (2004), and that anonymity can be relaxed. See also Deb and Kalai (2011) and Carmona and Podczeck (2011) for recent extensions of aspects of Kalai (2004). Recent work by Bodoh-Creed (2010) shows that Kalai-like assumptions imply a close relationship between games with a continuum of players and games with a large finite number of players.

Wilson (1987) famously argued that practical market designs should aim to be detail free, and Bergemann and Morris (2005) formalized the sense in which strategyproof mechanisms are robust in the sense of Wilson. Several recent papers have argued that strategyproofness can be viewed as a design objective and not just as a constraint: papers on this theme include Abdulkadiroğlu et al. (2006), Abdulkadiroğlu et al. (2009), Pathak and Sönmez (2008), and Roth (2008). Our paper contributes to this literature by showing that approximate strategyproofness is approximately costless in large markets, relative to other kinds of incentive compatibility. Also, the distinction we draw between manipulations that persist and manipulations that vanish highlights that many mechanisms in practice are manipulable in a preventable way.

Last, our paper is conceptually related to Parkes et al. (2001), Day and Milgrom (2008), and Pathak and Sönmez (2011), each of which seeks to say something more useful about non-strategyproof mechanisms than simply that they are not strategyproof.[2] Parkes et al. (2001) and Day and Milgrom (2008) propose cardinal measures of a combinatorial auction's manipulability, and seek to design an auction that minimizes manipulability subject to other design objectives. Pathak and Sönmez (2011) propose a method by which to compare non-strategyproof mechanisms based on their vulnerability to manipulations. Mechanism a is said to be more manipulable than Mechanism b if, for any problem instance where b is manipulable by at least one agent, so too is a. This criterion generates a partial order over mechanisms, and helps to explain several recent policy decisions in which school authorities switched from one manipulable mechanism to another. We view our approach as complementary to these prior approaches; see especially our discussion of Pathak and Sönmez (2011) after Example 1. An advantage of our approach is that it yields an explicit design desideratum, namely that mechanisms be strategyproof in the large.

Organization of the paper The rest of this paper is organized as follows. Section 2 describes the environment and some key assumptions. Section 3 defines strategyproof in the large and related concepts, and presents several examples. Section 4 presents the main theoretical result. Section 5 discusses various extensions of the main result. Section 6 concludes. Proofs are in the appendix.

[2] See also Milgrom (2011) Section IV for a general discussion of these issues.

2 Environment

2.1 Preliminaries

There is a finite type space, T, and a finite outcome space, X_0, with $X = \Delta X_0$ denoting the set of lotteries over outcomes. An outcome might be a consumption bundle, a school assignment, a match partner, etc. An agent's type determines her preferences over outcomes; specifically, for each $t_i \in T$ there is a von Neumann-Morgenstern expected utility function $u_{t_i} : X \to [0, 1]$. Preferences are **private values** in the sense that an agent's utility from an outcome depends only on her own type.

Our interest is in mechanisms that are well defined for various market sizes and various distributions of types, holding fixed T and X_0. The set of possible market sizes is simply $\mathbb{N}$, with $n \in \mathbb{N}$ denoting the number of agents in a particular economy. For each $n \in \mathbb{N}$, let $Y^n \subseteq X^n$ denote the set of allocations that are **feasible**. For instance, if outcomes are consumption bundles then Y^n is the set of allocations that satisfy the relevant production and capacity constraints in the n agent economy. If there is a single social decision to be made, then Y^n is the subset of X^n such that each agent gets exactly the same element of X.

We assume throughout that agents' types are independently and identically distributed (**iid**). Hence the set of possible preference distributions is ΔT, with $\mu \in \Delta T$ denoting the preference distribution (or "prior") in a particular economy.

2.2 Mechanisms

We define a mechanism as follows:

Definition 1. *A **mechanism** $\Gamma = ((\Phi^n)_{\mathbb{N}}, A)$ consists of a finite action space A and a sequence of allocation functions*

$$\Phi^n : A \times A^{n-1} \to X$$

each of which is anonymous and satisfies feasibility.[3] *The tuple (Φ^n, A), for a particular size n, is called an* **n-*mechanism***.

[3]Specifically, letting $\Phi^n(a_i, a_{-i})$ indicate the random bundle received by an agent who plays action $a_i \in A$, when the other agents play $a_{-i} \in A^{n-1}$, feasibility requires that, for all n and all $a = (a_1, \ldots, a_n)$, the allocation $(\Phi^n(a_1, a_{-1}), \ldots, \Phi^n(a_i, a_{-i}), \ldots, \Phi^n(a_n, a_{-n}))$ is in the feasible set Y^n. Anonymity requires that for all n, a_i and a_{-i}, the function $\Phi^n(a_i, a_{-i})$ is invariant to permutations in a_{-i}.

Built into our definition of a mechanism is the assumption that mechanisms are **anonymous**, meaning that each agent's outcome is a common function of her own action and the distribution of all actions.[4] Anonymity rules out that an agent's outcome depends on the precise details of who specifically plays what, and it also rules out that two agents who play the same action get different random bundles. Anonymity is a natural feature of many large-market settings, with examples of anonymous mechanisms including the Walrasian mechanism, most well-known single-object, combinatorial- and double- auction formats, and most of the mechanisms that have been proposed for single- and multi-unit assignment problems. In Section 5.2 we show that all of our results obtain if we relax anonymity to semi-anonymity (Kalai (2004)); semi-anonymity accommodates many additional settings in which there are asymmetries amongst classes of participants, e.g., in certain kinds of two-sided matching markets (cf. Azevedo and Leshno (2011)).

Also built into our definition of a mechanism is the assumption that mechanisms are **detail free** for the designer, in the sense that the function that maps from actions to outcomes does not vary with the prior μ. Of course, how agents choose to play a mechanism may depend on their prior. For instance, in the Bayes-Nash equilibria of the pay-as-bid auction (cf. Example 1 below), how much agents choose to shade their bids will depend on the prior.

2.3 Limit Mechanisms

Suppose there are n agents, and consider a distribution over actions $m \in \Delta A$. Let:

$$\phi^n(a_i, m) = \sum_{a_{-i}} \Phi^n(a_i, a_{-i}) \cdot \Pr(a_{-i}|m) \tag{2.1}$$

where $\Pr(a_{-i}|m)$ denotes the probability that the action vector a_{-i} is realized given $n-1$ iid draws from the action distribution m. The object $\phi^n(a_i, m)$ is a random bundle in X that describes what a generic agent i can expect to receive when he plays action a_i and the other $n-1$ agents play according to m.

We use the function $\phi^n(\cdot)$ to define limit mechanisms.

Definition 2. *The function $\phi^\infty : A \times \Delta A \to X$ is the **limit of mechanism** $((\Phi^n)_{\mathbb{N}}, A)$ if, for all*

[4] Anonymity is sometimes called "equal treatment of equals", after Aristotle's famous dictum (cf. Moulin (1995); Thomson (2011)).

a_i, m:

$$\phi^\infty(a_i, m) = \lim_{n\to\infty} \phi^n(a_i, m)$$

where ϕ^n is as defined in (2.1).

A feature of our method of taking the limit is that each ϕ^n in the sequence converging to ϕ^∞ is random, in the sense that the play of the agent's $n-1$ opponents is stochastic (drawn from distribution m). This is in contrast with, e.g., Debreu and Scarf's (1963) replicator economy, or with the approach pioneered by Aumann (1964) that looks directly at a continuum economy without explicitly modeling finite economies.

Our approach is more convenient than the replicator approach for two reasons. First, it allows each $\phi^n(\cdot,\cdot)$ in the sequence, as well as the limit $\phi^\infty(\cdot,\cdot)$, to be well defined for each possible preference distribution $m \in \Delta T$. This will help us to define persistent versus vanishing deviations (Definition 5) and to define SP-L itself (Definition 6). By contrast, the replicator approach is only well defined in the limit for rational distributions, and in finite markets of size n the outcome is only well defined for distributions that consist of multiples of $\frac{1}{n}$. Second, if a mechanism has a knife-edge point of discontinuity, in our limit landing on the knife's edge becomes vanishingly likely, whereas under the replicator approach landing on the knife's edge could be a certain event.[5]

While most (if not all) practical market design mechanisms we are aware of have limits according to our definition, we note that it is very easy to construct examples of mechanisms that do not. For instance, if a mechanism acts differently depending on whether n is even or odd it will not have a limit. For the remainder of the analysis we limit attention to mechanisms that have limits.

2.4 Standard Equilibrium Concepts

We briefly state the standard concepts of Bayes-Nash equilibrium and strategyproofness.

A strategy is a map $\sigma : T \to \Delta A$ from types to distributions over actions. With slight abuse of notation, we also write $\sigma(\mu)$ for the distribution over actions induced by drawing types iid according to $\mu \in \Delta T$ and then playing according to $\sigma(\cdot)$.

[5]For instance, if a fair coin is tossed 1,000,000 times, the probability that there are exactly 500,000 heads is very small, but a 500,000 fold replication of {heads, tails} will certainly result in exactly 500,000 heads.

Definition 3. *A **μ-Bayes-Nash equilibrium** (μ-BNE) of n-mechanism (Φ^n, A) is a strategy $\sigma^*_\mu(\cdot)$ such that for all $t_i \in T$ and $a'_i \in A$*

$$u_{t_i}[\phi^n(\sigma^*_\mu(t_i), \sigma^*_\mu(\mu))] \geq u_{t_i}[\phi^n(a'_i, \sigma^*_\mu(\mu))].$$

In words, the strategy σ^*_μ is a BNE if each agent's expected utility from playing according to σ^*_μ is higher than that from any other action, given that the other agents' types are distributed according to μ and that they also play according to σ^*_μ. Notice that there is no guarantee that $\sigma^*_\mu(t_i)$ is the best strategy for an agent of type t_i if the other agents play differently, which could occur, e.g., if the other agents make systematic mistakes, or play a different equilibrium, or if their types have a different distribution than μ.[6]

Part of the appeal of strategyproof mechanisms is that these informational requirements are no longer concerns.

Definition 4. *An n-mechanism (Φ^n, A) is **strategyproof** (SP) if $A = T$ and, for all $t_i, t'_i \in T$, and all $t_{-i} \in T^{n-1}$:*

$$u_{t_i}[\Phi^n_i(t_i, t_{-i})] \geq u_{t_i}[\Phi^n_i(t'_i, t_{-i})]$$

In words, a mechanism is strategyproof if the action space is such that agents simply report their types (i.e., $A = T$), and reporting truthfully is a dominant strategy. The definition of strategyproofness can easily be extended to accommodate action spaces that, while not equal to the set of types, nevertheless capture the idea that agents simply report their preferences. For instance, it is often the case in matching applications that the appropriate type space is the set of cardinal preferences, whereas the relevant action space is the set of ordinal preferences (e.g., Abdulkadiroğlu et al. (2011)).[7]

[6] By the standard revelation principle (cf. Fudenberg and Tirole (1991); Section 7.2), for any mechanism with a Bayes-Nash equilibrium in which agents misreport their preferences, there exists a direct-revelation mechanism in which telling the truth is a Bayes-Nash equilibrium. This direct-revelation mechanism, however, is no longer detail free for the designer; the map between types and outcomes will have to vary with the prior. For instance, in the direct-revelation mechanism version of the first-price sealed bid auction, the amount by which the center will shade each type's bid must vary with the prior in order for truthful reporting to be a BNE.

[7] Formally, say that mechanism $((\Phi^n)_{\mathbb{N}}, A)$ is a **preference-reporting mechanism** if the action space A partitions the type space T, and say that a preference-reporting mechanism is strategyproof if it is a dominant strategy to play the action associated with one's type. A direct mechanism in which $A = T$ is just a special case of a preference-reporting mechanism. Any preference-reporting mechanism can be represented as a direct mechanism, by interpreting the report t_i as the action associated with t_i.

3 Strategyproofness in the Large

Suppose that (Φ^n, T) is not strategyproof. This means that there exist t_i, t_i', t_{-i} such that

$$u_{t_i}[\Phi_i^n(t_i', t_{-i})] > u_{t_i}[\Phi_i^n(t_i, t_{-i})] \tag{3.1}$$

i.e., an agent of type t_i profits by misreporting as t_i'. We distinguish between two such kinds of manipulations.

Definition 5. *Suppose, as in (3.1), that $u_{t_i}[\Phi_i^n(t_i', t_{-i})] > u_{t_i}[\Phi_i^n(t_i, t_{-i})]$ for suitable n, t_i, t_i', and t_{-i}. Let $m = \text{emp}[t_{-i}]$ denote the empirical distribution of vector t_{-i} in T.*[8] *This manipulation is said to* ***vanish with market size*** *if $\lim_{n\to\infty} (u_{t_i}[\phi^n(t_i', m)] - u_{t_i}[\phi^n(t_i, m)]) \leq 0$, i.e., if :*

$$u_{t_i}[\phi^\infty(t_i', m)] \leq u_{t_i}[\phi^\infty(t_i, m)] \tag{3.2}$$

Else, the manipulation is said to ***persist with market size****.*

Heuristically, consider a mechanism in which each agent's outcome is a function of her own report, say t_i, and a set of statistics, say p, which are themselves a continuous function of the aggregate distribution of reports.[9] In finite markets each agent influences the aggregate distribution of reports, and hence p, but this influence goes to zero as $n \to \infty$. If a particular manipulation is profitable only because of the agent's ability to influence p – e.g., by reporting as t_i' instead of t_i, agent i can change p to p', and i prefers her outcome from t_i' and p' to that from t_i and p – then it will vanish with market size. If, however, a particular manipulation is profitable even holding fixed p, then the manipulation will persist with market size. The desiderata that we propose as a second-best alternative to strategyproofness is that a mechanism avoid manipulations that persist with market size:

Definition 6. *Mechanism $((\Phi^n)_{\mathbb{N}}, T)$ is* ***strategyproof in the large****, or* ***SP-L****, if, for all t_i, t_i', m:*

$$u_{t_i}[\phi^\infty(t_i, m)] \geq u_{t_i}[\phi^\infty(t_i', m)].$$

[8] Formally, given type space $T = \{t_1, t_2, \ldots, t_{|T|}\}$, and type vector t, emp$[t]$ is a $|T|$-dimensional vector whose j^{th} component is (number of occurences of t_j in t) divided by (number of elements in t).

[9] Formally, there exists a compact set of statistics P, and continuous functions $p : \Delta T \to P$ and $x : T \times P \to X$ such that, for all n, t_i, t_{-i}: $\Phi^n(t_i, t_{-i}) = x(t_i, p(\text{emp}[t_i, t_{-i}])$.

*Equivalently, a mechanism is **SP-L** if all finite-economy manipulations vanish with market size. If a mechanism has manipulations that persist with market size then it is **manipulable in the large**.*

We emphasize that Definitions 5 and 6 apply both to mechanisms that explicitly use prices in their description and those that do not have any explicit prices. Indeed, there are examples of price-based mechanisms that are manipulable in the large (e.g., pay-as-bid auctions), and examples of mechanisms that do not have explicit prices yet are strategyproof in the large (e.g., deferred acceptance algorithms).

We also emphasize how Definitions 5 and 6 treat manipulations that arise at points of discontinuity in a mechanism. If the discontinuity is "knife edge", then, since the probability of landing on the knife's edge vanishes to zero with market size for any iid distribution of opponent reports m, such a manipulation will be said to vanish with market size. The uniform-price auction is an example of a mechanism that has such knife-edge manipulations – an agent with multi-unit demand may wish to reduce her quantity demanded if she is the marginal bidder who sets price – and that we classify as SP-L. In the next section we discuss this and several other examples.

3.1 Examples

Our first example considers uniform-price and pay-as-bid auctions, two mechanisms best known for their use in the allocation of government securities. Neither mechanism is strategyproof. We show that the uniform-price auction is SP-L, while the pay-as-bid auction is not.

Example 1 (Multi-Unit Auctions)**.** There are kn units of a homogeneous good, with $k \in \mathbb{Z}_+$.

To simplify notation, we assume that agents' preferences take the form of linear utility functions, up to a capacity limit. Specifically, each agent i's type t_i consists of a per-unit value v_i and a maximum capacity q_i, with $V = \{1, \ldots, \bar{v}\}$ the set of possible values, $Q = \{1, \ldots, \bar{q}\}$ the set of possible capacity limits, and $T = V \times Q$. We can denote the set of outcomes, X_0, by $X_0 = V \times Q$ as well, by modeling an outcome as consisting of a per-unit payment, bounded above by $\bar{v}$, and a quantity allocated, bounded above by $\bar{q}$.

For both uniform-price and pay-as-bid auctions, agents simply report their types (i.e., $A = T$), and a single cutoff price p^* is calculated as a function of the reports $t = ((v_1, q_1), \ldots, (v_n, q_n))$ as

follows:[10]

$$p^*(t) = \max_{p \in V} \sum_{i=1}^{n} q_i \cdot \mathbf{1}\{v_i \geq p\} \geq kn$$

i.e., p^* is the highest price at which demand weakly exceeds supply. Allocations of the good are equivalent across the two mechanisms: an agent who reports (v_i, q_i) is allocated q_i units if $v_i > p^*$, is allocated 0 units if $v_i < p^*$, and is rationed if $v_i = p^*$. Payments differ across the two mechanisms. In the uniform-price auction, every agent who is allocated units pays the same per-unit price, p^*. In the pay-as-bid auction, every agent who is allocated units pays a per-unit price equal to her own reported value. It is easy to see that the pay-as-bid auction is not strategyproof. More subtly, neither is the uniform-price auction, because in the finite economy an agent may be able to lower the price $p^*(t)$ by reporting to demand fewer units than she actually does (Ausubel and Cramton (2002)).

Now let us consider the limit economy. In the limit, if the measure of agents' reports is $m \in \Delta T$, then for almost all m the cutoff price can be calculated as

$$p^*(m) = \max_{p \in V} \sum_{(v_i, q_i)} m(v_i, q_i) \cdot q_i \cdot \mathbf{1}\{v_i \geq p\} \geq k \tag{3.3}$$

The exception is if the cutoff price p^* that solves (3.3) does so with equality – i.e., there exists a price p^* such that $\sum_{(v_i,q_i)} m(v_i, q_i) \cdot q_i \cdot \mathbf{1}\{v_i \geq p^*\} = k$. In this event, the price in the limit will be p^* with probability one-half and will be $p^* - 1$ with probability one-half. This is due to the stochastic way that we take the limit: as n grows large, the probability that n iid draws from m result in demand strictly greater than supply at p^* is converging to one-half, just as is the probability that n iid draws result in demand strictly less than supply at p^*.

In the limit mechanism, each agent regards price as exogenous to her own report, because they cannot affect m. Notice in particular that our method of taking the limit ignores the vanishingly likely possibility that an individual agent can affect the price; i.e., it ignores the "knife edge" case discussed in Section 2.3.[11] It is thus easy to see that the uniform-price Auction is SP-L whereas

[10] The notation $\mathbf{1}\{statement\}$ denotes the indicator function which returns 1 if the statement is true and 0 if the statement is false.

[11] To illustrate the knife edge deviation, consider the case where there are n units of the good (i.e., $k = 1$), $Q = \{1, 2\}$, and $V = \{L, H\}$, with $H > L$. If exactly 50% of the population consists of H-value agents with demand for two units, and the remainder of agents have value L, then there is a profitable misreport: by reporting to demand one unit instead of two an H-value agent can lower the price from H to L and hence increase his total profit. This manipulation is knife edge because if instead the proportion of H-value agents is $(50 - \epsilon)\%$, then it is strictly more

the pay-as-bid auction is not. In particular, in the pay-as-bid auction an agent of type (v_i, q_i) with $v_i > p^* + 1$ can profitably misreport as $(\hat{v}_i = p^* + 1, \hat{q}_i = q_i)$ to get the same quantity at a strictly lower price than if he reports truthfully.

This example is consistent with Milton Friedman's (1991) observation that "you do not have to be a specialist" to participate in the uniform price auction, because you can just indicate "the maximum amount you are willing to pay for different quantitites ... if you bid a higher price [than the market clearing price], you do not lose as you do under the current [pay-as-bid] method." Friedman seems to be talking about the absence of what we call manipulations that persist, and seems to be less concerned by the vanishing manipulability of the uniform-price auction.[12]

Our next example is the Boston mechanism for school choice, a mechanism that does not explicitly have prices in the description. As mentioned in the introduction, this mechanism was criticized by Abdulkadiroğlu and Sönmez (2003) and Abdulkadiroğlu et al. (2006) for not being strategyproof. We show something stronger, which is that it is not even SP-L.

Example 2. Let X_0 be the set of schools, each with capacity qn, with $q \in (0, 1)$. That is q is the proportion of the overall student body that each school can accommodate.

Agents' types take the form of von-Neumann Morgenstern utility functions over the set of schools, i.e., functions of the form $u_{t_i} : X_0 \to \{0, 1, \ldots, \bar{u}\}$ for large integer $\bar{u}$. The set of actions A is the set of ordinal preferences over X_0, which is a partition of the type space, i.e., this is an example of a preference-reporting mechanism as defined in footnote 7.

The Boston mechanism awards as many students as possible their reported first choice school; then, awards as many students as possible their reported second choice school; etc. To keep the description concise we focus just on the first choices. Let $d_j = \sum_{i=1}^{n} \mathbf{1}\{j \text{ is } a_i\text{'s first choice}\}$ denote the number of students who report that school $j \in X_0$ is their first choice. Such students receive school j with probability $\min(1, \frac{qn}{d_j^1})$. Let $p_j = \min(1, \frac{qn}{d_j^1})$.

profitable to report truthfully and get two units, and if the proportion is $(50 + \epsilon)\%$ then the manipulation neither increases nor decreases profits.

[12] Pathak and Sonmez (2011) provide a complementary perspective on the incentive comparison between the uniform-price and pay-as-bid auctions. Pathak and Sonmez (2011) show that any agent who can profitably manipulate the uniform-price auction in a given finite economy can also profitably manipulate the pay-as-bid auction in that same finite economy. Moreover, the latter manipulation is always larger in utility terms. Thus, Pathak and Sonmez (2011) suggests that the pay-as-bid auction is more manipulable than the uniform-price auction in any given finite economy, whereas our analysis highlights that the pay-as-bid auction's manipulability persists with market size, whereas the uniform-price auction is strategyproof in the large.

The limit is very similar. If the overall measure of agents' reports is $m \in \Delta A$, let m_j denote the measure of students who report that school $j \in X_0$ is their first choice, i.e., $m_j = \sum_{a_i \in A} m(a_i) \cdot (\mathbf{1}\{j \text{ is } a_i\text{'s first choice}\})$. The probability that a student who ranks j first gets it can be calculated as

$$p_j^* = \min(1, \frac{q}{m_j})$$

Notice that in the limit mechanism each agent regards the p_j^*'s as exogenous to their own report. Agent t_i will wish to misreport her first choice school if her first choice is j, but there exists j' where ranking j' first gives her strictly greater expected utility, i.e., $u_{t_i}(j')p_{j'}^* > u_{t_i}(j)p_j^*$. Therefore the mechanism is manipulable in the large.

There are numerous other examples. For single-unit assignment problems such as in Example 2, Hylland and Zeckhauser's (1979) pseudomarket mechanism is an example of a price-based mechanism that is SP-L, while Bogomolnaia and Moulin's (2001) probabilistic serial mechanism is an example of a mechanism that does not explicitly use prices in the original description but that is SP-L (cf. Kojima and Manea (2010)). For multi-unit assignment problems, the mechanisms found in practice are manipulable in the large, specifically the Bidding Points Auction studied by Sönmez and Ünver (2010), and the Draft Mechanism studied by Budish and Cantillon (Forthcoming). Mechanisms recently proposed in theory are SP-L, specifically the Approximate Competitive Equilibrium from Equal Incomes mechanism proposed by Budish (2009), and the multi-unit generalization of Hylland and Zeckhauser's pseudomarket proposed by Budish et al. (2011).

The concepts can also be applied to two-sided matching mechanisms, if we generalize the class of mechanisms considered to be the class of semi-anonymous mechanisms (Kalai (2004)), and not just anonymous mechanisms; cf. Section 5.2. Then, techniques in Kojima and Pathak (2009) or Azevedo and Leshno (2011) can be used to show that Gale and Shapley's deferred acceptance algorithm is SP-L in semi-anonymous environments. It is also easy to see that the priority-match algorithm, criticized by Roth (2002) and others, is manipulable in the large.

The following table summarizes this informal discussion.

Table 1. Which Non-SP Market Designs are SP-L?

Problem	Manipulable in the Large	Strategyproof in the Large
Single-unit Assignment	Boston Mechanism	Prob Serial, HZ Pseudomarket
Multi-unit Assignment	Bidding Points Auction	Approximate CEEI
	HBS Draft Mechanism	Generalized HZ Pseudomarket
Multi-unit Auctions	Pay-as-Bid Auctions	Uniform-Price Auctions
Two-Sided Matching	Priority-Match Algorithm	Deferred Acceptance Algorithm

4 Main Result

Strategyproofness often severely limits what kinds of mechanisms are possible. Our main result identifies a sense in which SP-L does not. The result requires a quasi-continuity assumption, which we present and discuss in Section 4.1. We then present the main result in Section 4.2, along with a proof sketch. Since the proof is by construction, we provide an example construction in Section 4.3, using the Boston mechanism for school choice discussed in Example 2.

4.1 Quasi-Continuity of Equilibria

We first need some new notation. Given a market size n and a distribution $\bar{m} \in \Delta(A^{n-1})$ over action profiles, we may extend $\Phi^n(\cdot,\cdot)$ linearly as:

$$\Phi^n(a,\bar{m}) = \sum_{a_{-i}} \Phi^n(a,a_{-i}) \cdot \bar{m}(a_{-i}).$$

Now consider an $n-1$ vector of types t_{-i}, and a strategy $\sigma : T \to \Delta A$. Together, σ and t_{-i} induce a probability distribution over action profiles, i.e., an element of $\Delta(A^{n-1})$. We will denote this induced distribution as $\sigma(t_{-i})$. We highlight that $\sigma(t_{-i})$ denotes a distribution over A^{n-1}. We will then use the notation

$$\Phi^n(a_i, \sigma(t_{-i}))$$

to describe what happens when i plays a_i, and the other players have types t_{-i} and play according

to σ.[13]

Next, we need to define a limit Bayes-Nash equilibrium. Given a mechanism $((\Phi^n)_{\mathbb{N}}, A)$, with limit $\phi^\infty(\cdot,\cdot)$, the strategy $\sigma^*_\mu(\cdot)$ is a limit μ-BNE if, for all $t_i \in T$ and $a'_i \in A$:

$$u_{t_i}[\phi^\infty(\sigma^*_\mu(t_i), \sigma^*_\mu(\mu))] \geq u_{t_i}[\phi^\infty(a'_i, \sigma^*_\mu(\mu))].$$

Limit equilibria are simply strategy profiles that become arbitrarily close to optimal as the economy grows large. Last, we need the concept of a family of limit equilibria. Given a mechanism $((\Phi^n)_{\mathbb{N}}, A)$ with limit $\phi^\infty(\cdot,\cdot)$, we say that $(\sigma^*_\mu)_{\mu\in\Delta T}$ is a family of limit Bayes-Nash equilibria if, for each $\mu \in \Delta T$, σ^*_μ is a limit μ-BNE.

Our main continuity notion is as follows:

Definition 7. *Consider a mechanism $((\Phi^n)_{\mathbb{N}}, A)$ with limit $\phi^\infty(\cdot,\cdot)$, and a family of limit Bayes-Nash equilibria $(\sigma^*_\mu)_{\mu\in\Delta T}$. The family of equilibria is **quasi-continuous** if, for every $\mu_0 \in \Delta T$ and $\epsilon > 0$, there exists a neighborhood $\mathcal{N}$ of μ_0 such that:*

1. *$\mathcal{N}$ can be decomposed as*
$$\mathcal{N} = \cup_{1\leq k\leq K} \mathcal{A}_k \cup \mathcal{B}$$
where the $\mathcal{A}_k$ are open sets.

2. *$\lim_{n\to\infty} \Pr\{\text{distance}(\text{emp}[t], \mathcal{B}) \leq 1/n | t \in T^n, t \sim iid(\mu_0)\} = 0$, where $t \sim iid(\mu_0)$ denotes a vector of n types t with each component drawn iid according to μ_0.*

3. *There exists n_0 such that for any $n > n_0$, any $\mu, \mu', \text{emp}[t_i, t_{-i}], \text{emp}[t_i, t'_{-i}] \in \mathcal{A}_k$, and any $a_i \in A$ we have:*
$$|\Phi^n(\sigma^*_\mu(t_i), \sigma^*_\mu(t_{-i})) - \Phi^n(\sigma^*_{\mu'}(t_i), \sigma^*_{\mu'}(t'_{-i}))| < \epsilon.$$

*The family of equilibria is **continuous at μ_0** if, for the prior μ_0, Conditions 1 and 3 hold with $K = 1$ and $\mathcal{B} = \emptyset$.*

[13] We highlight that we use the notation
$$\phi^n(a_i, m),$$
where m is a distribution over A to denote the payoff to player i when her opponents' play is independently and identically distributed as m. This is a very different object than
$$\Phi^n(a_i, \sigma(t_{-i})),$$
which is i's payoff when her opponents have types given by the vector t_{-i} and play strategy σ.

To describe Definition 7, first consider the case where the mechanism is continuous at a given prior μ_0. This means, per Condition 3, that agents' outcomes are locally continuous with respect to small changes in the other agents' types, and to the priors they base their strategies on. For instance, in an auction in which bidders shade their values, a small change in the distribution of values might have just a small effect on the amount by which bidders shade their bids.

Quasi-continuity allows for some violations of continuity. Specifically, for each prior μ_0, quasi-continuity requires that a small enough neighborhood $\mathcal{N}$ of μ_0 can be decomposed into two types of sets. First, there is a set $\mathcal{B}$ in which the empirical realization of types is vanishingly likely to land (Condition 2). This set $\mathcal{B}$ can contain points of discontinuity. Second, are sets of the form $\mathcal{A}_k$, where, so long as the empirical realization of types and the prior associated with the strategy that agents use land within the same one of these sets, agents' outcomes are locally continuous in types and strategies (Condition 3). Intuitively, in a knife-edge example such as the uniform-price auction sketched in Section 2.1, think of $\mathcal{B}$ as the set $m = \frac{1}{2}$, $\mathcal{A}_1$ as the set $m < \frac{1}{2}$, and $\mathcal{A}_2$ as the set $m > \frac{1}{2}$.

Many familiar mechanisms have equilibria that are quasi-continuous but not continuous. Examples include the Boston mechanism as modeled by Abdulkadiroğlu et al. (2011), the deferred acceptance algorithm as modeled by Azevedo and Leshno (2011), the Bidding Points Auction modeled by Sönmez and Ünver (2010), and the uniform price auction discussed earlier.[14]

As of the present writing, we are not aware of practical market design mechanisms that do not have a quasi-continuous family of equilibria. Note too that if a mechanism's equilibria are not quasi-continuous, then the analyst's prediction of equilibrium outcomes is highly sensitive to small changes in information about the environment. This itself is arguably an undesirable feature of a mechanism.

4.2 Construction of SP-L Mechanisms

Our main result is the following:

[14]Here are some additional details for readers familiar with these mechanisms. In the Boston mechanism, the potential for discontinuity is when a school reaches capacity exactly at the end of some round. It is possible to construct examples where some student's probability of obtaining a particular school discontinuously changes from 0 to 1 as the quantity available in some round changes from 0 to something strictly positive. In the college admissions problem, for any particular distribution of play, most students' outcomes will be locally continuous, but students who are right at the cutoff for a particular school may have their outcome change discontinuously. The Bidding Points Auction is similar to college admissions, in that most students' outcomes are locally continuous, but students whose bid is right at the cutoff for a particular course may have their outcome change discontinuously. In the uniform price auction, the potential for discontinuity is when a small change in the distribution of play discretely changes the price, e.g. when price discontinuously changes from H to L in the example of Section 2.3.

Theorem 1. *Consider a mechanism $\Gamma = ((\Phi^n)_{n\in\mathbb{N}}, A)$ with a quasicontinuous family of limit equilibria $(\sigma^*_\mu)_{\mu\in\Delta T}$. Then there exists a direct mechanism $\Gamma^D = ((F^n)_{\mathbb{N}}, T)$ with the following properties*

1. *Γ^D is strategyproof in the large.*

2. *If Γ^D is continuous at the true prior μ_0, then in the limit as $n \to \infty$, truthful play of Γ^D and Bayes-Nash equilibrium play of Γ give agents the same utilities. Formally, given $\mu_0 \in \Delta T$ and $\epsilon > 0$, there exists n_0 such that for all $n > n_0$ and all t_i:*

$$|u_{t_i}[f^n(t_i, \mu_0)] - u_{t_i}[\phi^n(\sigma^*_{\mu_0}(t_i), \sigma^*_{\mu_0}(\mu_k))]| < \epsilon,$$

   where $f^n(\cdot)$ is constructed from $F^n(\cdot)$ according to (2.1).

3. *If Γ^D is not continuous at the true prior μ_0, then in the limit as $n \to \infty$, truthful play of Γ^D gives agents the same utilities as a convex combination of equilibrium outcomes under Γ, for priors in a neighborhood of μ_0. Formally, for every $\mu_0 \in \Delta T$ and $\epsilon > 0$, there exist priors μ_k with $|\operatorname{emp}[\mu_k] - \operatorname{emp}[\mu_0]| < \epsilon$, and n_0, such that for all $n > n_0$ there exist weights π^n_k summing to one such that, for all t_i:*

$$|u_{t_i}[f^n(t_i, \mu_0)] - \sum_{k=1,\ldots,K} \pi^n_k \cdot u_{t_i}[\phi^n(\sigma^*_{\mu_k}(t_i), \sigma^*_{\mu_k}(\mu_k))]| < \epsilon.$$

The proof of Theorem 1 is by construction. We provide a sketch as follows (the full details are in Appendix A). Suppose in a market of size n agents report $t = (t_1, \ldots, t_n)$. Let:

$$F^n(t) = \Phi^n(\sigma^*_{\operatorname{emp}(t)}(t)) \tag{4.1}$$

In words, F^n plays action $\sigma^*_{\operatorname{emp}(t)}(t_i)$ for agent i who reports t_i, where $\operatorname{emp}(t)$ is not the true distribution of agents' types μ_0 (which is not known to the mechanism) but rather the *empirical distribution of agents' reports*. We will show that the direct mechanism $\Gamma^D = ((F^n)_{\mathbb{N}}, T)$ is strategyproof in the large and gives agents the same utilities in the limit as the original mechanism.

First, suppose that agents report their preferences truthfully, according to the true prior μ_0. In a finite market of size n there will be sampling error, so the realized empirical will be, say, $\hat{\mu}$. Agent i who reports t_i receives $F^n(t_i, t_{-i}) = \Phi^n(\sigma^*_{\hat{\mu}}(t_i), \sigma^*_{\hat{\mu}}(t_{-i}))$. As the market grows large, the realized

empirical $\hat{\mu}$ converges to the true distribution μ_0, by the law of large numbers. Hence, assuming for now that the original mechanism is continuous at μ_0, agent i's allocation is converging to $\Phi^n(\sigma^*_{\mu_0}(t_i), \sigma^*_{\mu_0}(t_{-i}))$, exactly what he receives under Bayes-Nash equilibrium of the original mechanism. Thus, if all agents report truthfully, our mechanism coincides with the original mechanism in the limit, as required.

Now, suppose that the agents other than i misreport their preferences, according to some distribution $m \in \Delta T$. As before, in a finite market of size n, there will be sampling error, so the realized empirical will be, say, $\hat{m}$. Agent i will thus receive $F^n(t_i, t'_{-i}) = \Phi^n(\sigma^*_{\hat{m}}(t_i), \sigma^*_{\hat{m}}(t'_{-i}))$. As the market grows large, the realized empirical $\hat{m}$ will converge towards m, so, assuming continuity at m, agent i's allocation is converging to $\Phi^n(\sigma^*_m(t_i), \sigma^*_m(t'_{-i}))$. This is what agent i would receive under the original mechanism *in the Bayes-Nash equilibrium corresponding to prior* m. Even though the other agents are systematically misreporting their preferences, our agent i remains happy to tell the truth, because the other agents are acting *as if* their preferences are distributed according to m, and then playing a strategy that is converging to the Bayes-Nash equilibrium corresponding to m. Thus agent i also wants to play the Bayes-Nash equilibrium strategy corresponding to m – which is exactly what happens when he reports his preferences truthfully to $((F^n)_{\mathbb{N}}, T)$.[15] Hence, in the limit, we get dominant-strategy incentives, i.e., our constructed mechanism is SP-L.

The last step of the proof sketch is to describe what happens in the event that the equilibrium of the original mechanism is not continuous at μ_0 – e.g., the uniform-price auction example described in Section 2.3, for the case where $\mu_0 = \frac{1}{2}$. This requires a technical lemma (Lemma 1 in the appendix) which says that, for any arbitrary prior $m \in \Delta T$, the allocation an agent receives under $((F^n)_{\mathbb{N}}, T)$ can be approximated by a convex combination of the allocations he would receive in the limit Bayes-Nash equilibria of $((\Phi^n)_{\mathbb{N}}, A)$, for priors close to m. The key to the proof of the lemma is that, in a large enough market, a single agent cannot appreciably change the probability that the aggregate profile lands in each region $\mathcal{A}_k$, as defined in Definition 7. This allows us to exploit the continuity *within* each region $\mathcal{A}_k$, and the vanishing likelihood that the aggregate profile lands near the discontinuity region $\mathcal{B}$.

It is important to emphasize how our constructed mechanism $((F^n)_{\mathbb{N}}, T)$ differs from a traditional Bayes-Nash direct revelation mechanism (cf. Fudenberg and Tirole (1991); Section 7.2). In a traditional Bayes-Nash DRM, the mechanism needs to know the prior μ_0, i.e., it is not detail free.

[15] Observe that this step of the argument requires the private values assumption. It is important that i does not care per se about the other players' true types.

The mechanism then announces a BNE strategy $\sigma^*_{\mu_0}(\cdot)$, and plays $\sigma^*_{\mu_0}(t_i)$ on behalf of an agent who reports t_i. Our mechanism *infers* a prior from the empirical distribution of agents' play. If agents indeed play truthfully, this inference is exactly correct in the limit, and our detail-free mechanism coincides with the traditional Bayes-Nash DRM. But if the agents other than i misreport, so that the empirical $\hat{m}$ is very different from the prior μ_0, then our mechanism automatically adjusts each agent's play to be the Bayes-Nash equilibrium play in a world where the prior was in fact $\hat{m}$. As a result, an agent who reports his preferences truthfully remains happy to have done so even if the other agents misreport, which is not the case in a traditional Bayes-Nash DRM. To summarize, the two key differences between our mechanism and a traditional Bayes-Nash DRM are: (1) our mechanism is detail free; (2) our mechanism is strategyproof in the large.

4.3 Example of the Construction

We describe our main result in the context of a specific example, the Boston mechanism for school choice (cf. Example 2 above). Abdulkadiroğlu and Sönmez (2003) and Abdulkadiroğlu et al. (2006) criticized the Boston mechanism on the grounds that it is not strategyproof, and proposed that the strategyproof Gale-Shapley deferred acceptance algorithm be used instead. Indeed, the Gale-Shapley algorithm was eventually adopted for use in practice (cf. Roth (2008)). However, a second generation of papers on the Boston mechanism has argued that it has a Bayes-Nash equilibrium that yields greater student welfare than does the dominant strategy equilibrium of the Gale-Shapley procedure (Abdulkadiroğlu et al. (2011); Miralles (2009); Featherstone and Niederle (2008)). Perhaps, these papers argue, the earlier papers were too quick to dismiss the Boston mechanism in favor of strategyproof deferred acceptance.

Of course, the Bayes-Nash equilibria these second-generation papers construct rely on students having common knowledge of the distribution of other students' preferences; on students being able to coordinate on a specific equilibrium; on students being able to make very precise strategic calculations to determine whether to risk asking for a popular school; etc. Our Theorem 1 says that all of this complexity and non-robustness is unnecessary in a large market. Specifically, there must exist yet another mechanism that implements the same outcomes as these desirable Bayes-Nash equilibria of the Boston mechanism, but that is SP-L.

Interestingly, in the simplified version of the Boston Mechanism that we discussed in Example 2 above, the SP-L mechanism that we construct according to (4.1) closely resembles the Hylland and

Zeckhauser (1979) pseudo-market mechanism for single-unit assignment. Specifically, write agent i's type as a vector of von-Neumann Morgenstern utilities, one for each school: $v_{i1}, \ldots, v_{i|X_0|}$. The limit Bayes-Nash equilibria of the Boston mechanism are characterized by a set of probabilities p_j^*, one for each school j, such that when each student i asks for the school that maximizes his expected utility – i.e., maximizes his expectation of $p_j^* \cdot v_{ij}$ – these probabilities are in fact correct. Thus our mechanism works as follows. First, students report their types. Next, the mechanism calculates the Bayes-Nash equilibrium associated with the empirical distribution of the reported types, i.e., calculates the market-clearing probabilities. Last, each student is given the appropriate lottery. This is just like Hylland and Zeckhauser (1979), except that in that paper the p_j^*'s are called "prices".[16]

We emphasize the two advantages of our constructed mechanism as compared to the Bayes-Nash equilibria of the original Boston mechanism. First, students need not estimate the distribution of other students' preferences, the associated market-clearing probabilities, etc. Second, our mechanism is robust to systematic mistakes by the other students, or miscoordination over which equilibrium to play, etc., because in the limit it provides dominant-strategy incentives.

5 Extensions

5.1 Ex post Robustness of SP-L Mechanisms in Large Finite Markets

Kalai (2004) studies the ex post robustness of Bayes-Nash equilibria in large games. Under an equicontinuity assumption that we provide below, he shows that Bayes-Nash equilibria are ex post robust in the following sense: For any $\epsilon > 0$, the probability that any player will have an ex post deviation that yields a gain of more than ϵ converges to 0 exponentially in market size n.[17]

Here we show that, under Kalai's equicontinuity assumption, SP-L mechanisms satisfy a stronger robustness property: For any $\epsilon > 0$, in a large enough market, no player has an ex post deviation

[16] Miralles (2009) contains a nice description of the connection between the Boston mechanism's Bayes-Nash equilibria and Hylland Zeckhauser.

[17] More specifically, Kalai (2004) considers a sequence of (semi-)anonymous games, with an increasing number of players, that satisfies an equicontinuity condition. He defines an ϵ-ex post Nash equilibrium profile as a profile of types and actions such that no player may gain more than ϵ by changing her action. He defines a profile of (possibly mixed) strategies to be an (ϵ, ρ) ex post strategy profile if with probability at least $1 - \rho$ the realized profile of types and strategies is an ϵ ex post Nash equilibrium. Kalai's Theorem 1 shows that, for any sequence of Bayes Nash equilibria σ^n of the games, and any $\epsilon > 0$, there exist constants $\alpha > 0, \beta < 1$ such that σ^n is an $(\epsilon, \alpha\beta^n)$ ex post strategy profile. We refer the interested reader to Kalai (2004) for more details.

which increases her payoff by more than ϵ. This is a stronger ex post robustness property for two reasons. First, in a large enough market, the probability of an ϵ deviation is exactly zero rather than converging to zero. Second, the fact that the ex post gain from any deviation is small does not depend on a player knowing the distribution of opponent types, nor on agents playing in equilibrium.

Formally, we follow Kalai and define an equicontinuous mechanism as follows.

Definition 8. *A mechanism $\Gamma = ((\Phi^n)_{n\in\mathbb{N}}, A)$ is equicontinuous if, for all $\epsilon > 0$ there exists $\delta > 0$ such that for all $n, n', a_{-i} \in A^{n-1}, a'_{-i} \in A^{n'-1}$ with*

$$|\operatorname{emp}[a_{-i}] - \operatorname{emp}[a'_{-i}]| < \delta$$

we have that for all a_i

$$|\Phi^n(a_i, a_{-i}) - \Phi^{n'}(a_i, a'_{-i})| < \epsilon.$$

We then define ϵ-Strategyproofness as follows.

Definition 9. *A direct mechanism $\Gamma = ((\Phi^n)_{n\in\mathbb{N}}, t)$ is (ϵ, n)-strategyproof if for all $t \in T^n, t'_i \in T$*

$$u_{t_i}[\Phi^n(t)] \geq u_{t_i}[\Phi^n(t'_i, t_{-i})] - \epsilon.$$

It is then possible to prove the following result.

Proposition 1. *If $\Gamma = ((\Phi^n)_{n\in\mathbb{N}}, T)$ is SP-L and equicontinuous, then given $\epsilon > 0$ there exists n_0 such that for all $n > n_0$, we have that Γ is (ϵ, n) strategyproof.*

5.2 Semi-Anonymity

Our main analysis considers anonymous mechanisms, where agents' outcomes depend on their own report and the distribution of all reports. The analysis generalizes straightforwardly (though at some notational burden) to the case of semi-anonymous mechanisms, as defined by Kalai (2004).

Assume that agents belong to groups g in a finite set G. Each group has a different set of possible types and actions, so that

$$\begin{aligned} T &= T_{g_1} \cup T_{g_2} \cup \cdots \cup T_{g_G} \\ A &= A_{g_1} \cup A_{g_2} \cup \cdots \cup A_{g_G}. \end{aligned}$$

A semianonymous mechanism is defined as $((\Phi^n)_{n\in\mathbb{N}}, (T_g)_{g\in G}, (A_g)_{g\in G})$.

Consider the case where there is a distribution μ over the set of types T. In this case, the number of agents belonging to each group is random.[18] Assume moreover that strategies are restricted so that an agent cannot misreport her group, that is, if $t_i \in T_g$ then the support of $\sigma(t_i) \in A_g$. Formally, a strategy is defined as a map

$$\sigma : A \to \Delta T$$

such that for all $t_i \in T_g$, the support of $\sigma(t_i)$ is contained in A_g. This framework covers the following type of example.

Example 3. (Two-Sided Matching) This example shows that semi-anonymous mechanisms can cover matching mechanisms in two-sided markets (Roth and Sotomayor 1990).

There are men and women, who differ on a set of characteristics. Groups g index both sex and the characteristics, so that the set of groups is

$$G = \{m_1, m_2, \cdots, m_M\} \cup \{w_1, w_2, \cdots, w_W\}.$$

That is, there are M groups of men and W groups of women. Men and women within each group have the same characteristics, and hence are equally good marriage partners. However, within each group, agents may differ in their preferences over the other groups.

Formally, agent i's type is

$$t_i = (g_{t_i}, u_{t_i}),$$

where g_{t_i} is the agent's characteristics, and u_{t_i} is a strictly positive utility function over the groups of the other sex. The set of outcomes $X_0 = G \cup \emptyset$. That is, each agent only cares about which type of man (woman) she (he) is matched to, or whether she (he) is unmatched. Utilities of each type t_i are given by $u_{t_i}(g)$ if she is matched to someone of the opposite sex. We extend u_{t_i} so that it is 0 if the agent is unmatched or matched to a group of the same sex.

Consider now the direct mechanism where $A = T$. Men and women report a vector of types t, and therefore characteristics. This implies a weak preference ordering of each man over each woman and vice versa. Given these preferences, at least one stable matching exists. If there are multiple stable

[18] The argument below can be extended to the case where in a market of size n there is a fixed number of players n_g within each group.

matchings, choose one uniform randomly. For each type t_i we let

$$\Phi^n(t_i, t_{-i})$$

denote the resulting probability distribution over match partners.

The only difference between the semi-anonymous case and the original anonymous case is that agents of a given group g cannot play actions not in A_g. It is straightforward to incorporate this extra generality within our framework, without affecting any of the results.

This is demonstrated by the following construction. Suppose we are given a semi-anonymous mechanism, as defined in this section. We construct an anonymous mechanism where it is strictly dominated for any agent in group g to play a strategy outside A_g. Denote the original semi-anonymous mechanism by $\bar{\Phi}^n$, with type and action spaces as above, outcome space $\bar{X}_0$, and utility functions $\bar{u}_{t_i}$ such that $\bar{u}_{t_i}(\cdot) > 0$.

We now define the new mechanism Φ such that it is a dominated strategy for an agent in group g to play an action outside A_g. The bundles of the new mechanism are of the form $x_0 = (\bar{x}_0, g) \in \bar{X}_0 \times G$. Let

$$\Phi^n(a_i, a_{-i}) = (\bar{\Phi}(a_i, a_{-i}),\ \text{group of } a_i).$$

If $t_i \in T_g$, then for all $\bar{x}_0 \in \bar{X}_0$, let

$$\begin{aligned} u_{t_i}(\bar{x}_0, g') &= 0 && \text{if } g' \neq g \\ u_{t_i}(\bar{x}_0, g') &= \bar{u}_{t_i}(\bar{x}_0) && \text{if } g' = g. \end{aligned}$$

Thus, it is strictly dominated for an agent of type within group T_g to play an action that is not in A_g. This construction demonstrates that any semi-anonymous mechanism can be embedded in our framework, with the same set of strictly undominated strategies. Formally, we may state this as follows.

Proposition 2. *For any semi-anonymous mechanism $((\Phi^n)_{n\in\mathbb{N}}, (T_g)_{g\in G}, (A_g)_{g\in G})$, there exists an anonymous mehcanism that gives agents with types in T_g playing actions within A_g the same utility as the original mechanism, but where it is a strictly dominated strategy to play strategies outside of A_g.*

Thus, all of our analysis generalizes straightforwardly to the case of semi-anonymous mechanisms as defined by Kalai (2004).

5.3 Finite-Economy Bayes-Nash Equilibria

Theorem 1 starts from a family of limit equilibria $(\sigma^*_\mu)_{\mu\in\Delta T}$, and constructs a direct SP-L mechanism that implements approximately the same outcome. This construction could also be obtained based on a sequence of families of finite-economy equilibria, $(\sigma^n_\mu)_{\mu\in\Delta T, n\in\mathbb{N}\cup\infty}$. In this section, we provide such a construction. The reason why we focus on limit equilibria for the statement of the main result is that finite-economy equilibria are often analytically less tractable. In multi-unit auctions, for example, closed form solutions for equilibria are often unavailable, and even showing basic properties of equilibria is a difficult problem (Swinkels (2001); Engelbrecht-Wiggans et al. (2006)). Due to the greater cognitive complexity of finite-economy equilibria, an analyst might find limit equilibria a more compelling concept, or at least easier to apply. On the other hand, depending on the application, an analyst might see exact Nash equilibria of finite economies as a more appropriate solution concept. For that reason, the we cover both cases.

In order to modify Theorem 1 to accommodate a sequence of families of finite-economy equilibria, we need the following alternative to Definition 7.

Definition 10. *Consider a mechanism $((\Phi^n)_{\mathbb{N}}, A)$ with limit $\phi^\infty(\cdot,\cdot)$, and a sequence of families of Bayes-Nash equilibria $(\sigma^n_\mu)_{\mu\in\Delta T, n\in\mathbb{N}\cup\infty}$. The sequence of families of equilibria is* ***quasi-continuous*** *if, for every $\mu_0 \in \Delta T$ and $\epsilon > 0$, there exists a neighborhood $\mathcal{N}$ of μ_0 such that:*

1. *$\mathcal{N}$ can be decomposed as*

$$\mathcal{N} = \cup_{1\leq k\leq K} \mathcal{A}_k \cup \mathcal{B}$$

where the $\mathcal{A}_k$ are open sets.

2. *$\lim_{n\to\infty} \Pr\{\text{distance}(\text{emp}[t], \mathcal{B}) \leq 1/n | t \in T^n, t \sim iid(\mu_0)\} = 0$, where $t \sim iid(\mu_0)$ denotes a vector of n types t with each component drawn iid according to μ_0.*

3. *There exists n_0 such that for any $n > n_0$, any $\mu, \mu', \text{emp}[t_i, t_{-i}], \text{emp}[t_i, t'_{-i}] \in \mathcal{A}_k$ we have:*

$$|\Phi^n(\sigma^n_\mu(t_i), \sigma^n_\mu(t_{-i})) - \Phi^n(\sigma^n_{\mu'}(t_i), \sigma^n_{\mu'}(t'_{-i}))| < \epsilon.$$

4. In addition, for all $\mu \in \Delta T, t_i \in T$

$$\lim_{n\to\infty} \phi^n(\sigma^n_\mu(t_i), \sigma^n_\mu(\mu)) = \phi^\infty(\sigma^\infty_\mu(t_i), \sigma^\infty_\mu(\mu)).$$

Under this modified definition, the appropriately stated analogue of Theorem 1 holds. The specific n-mechanisms we use in the construction are

$$F^n(t) = \Phi^n(\sigma^n_{\text{emp}[t]}(t)). \tag{5.1}$$

This is as in Equation (4.1), but using the finite-economy equilibria instead of the limit equilibria. The proof is a variation of the proof of Theorem 1. The Appendix contains the formal statement of the result and discusses the differences between the proofs.

5.4 Complete Information Nash Equilibria

Our construction in Section 4 takes as input a mechanism that has a family of Bayes-Nash equilibria. The same idea can be applied to a mechanism that has a family of complete-information Nash equilibria. In this case, our constructed mechanism takes the same form as (5.1). That is

$$F^n(t) = \Phi^n(\sigma^n_{\text{emp}[t]}(t)). \tag{5.2}$$

The difference is that now $\sigma^n_{\text{emp}[t]}(\cdot)$ is a complete-information Nash equilibrium strategy, in an economy in which the type profile is t (or any permutation thereof). In words, agents report their types to the mechanism, which then computes a symmetric complete information Nash equilibrium strategy in the economy induced by the reports. Note that in general it is *not* a Nash equilibrium for each player to report their preferences truthfully to this mechanism in finite markets. The reason is that, by changing one's report from say t_i to t'_i, one changes the profile of reported types from say t to t', and this in turn changes the strategy that is activated from $\sigma^n_{\text{emp}[t]}(\cdot)$ to $\sigma^n_{\text{emp}[t']}(\cdot)$. Thus, i changing his *report* can have the effect of the mechanism changing j's *action.* As the market grows large, i's influence on emp$[t]$ grows small, so under a continuity assumption analogous to those stated above, mechanism (5.2) is SP-L. An interesting feature of this construction (5.2) is that if agents tell the truth in finite markets, then (5.2) produces outcomes that are *identical* to the outcomes under the Nash equilibria of the original mechanism. By contrast, with Bayes-Nash

equilibria our constructed mechanism only approximates the finite market outcomes.

We illustrate the construction (5.2) with the Generalized Second Price (GSP) auction for search-engine advertising slots, as modeled by Edelman et al. (2007) (EOS). EOS showed that the GSP has complete-information Nash equilibria that coincide with the Vickrey-Clarke-Groves mechanism on the equilibrium path. Interestingly, our constructed mechanism coincides with VCG both on and off the equilibrium path and hence actually provides dominant strategy incentives in finite markets. As emphasized above, however, this is not generally the case; typically we will need to consider a large-market limit for (5.2) to provide exact incentives.

Example 4 (Generalized Second Price Auction)**.** Consider the Generalized Second Price (GSP) auction for search-engine advertising slots, as modeled by EOS. There are k advertising slots, with click-through rates $\alpha_1 > \alpha_2 > \cdots > \alpha_k$. There are $n > k$ bidders, with per-click values of $v_1 > v_2 > \cdots > v_n$. The GSP works as follows. Each bidder i submits a bid b_i; next, the bids are ranked in descending order, with the highest bidder awarded the first advertising slot, the second-highest bidder awarded the second slot, etc.; last, each successful bidder pays a per-click amount equal to the next highest bid after their own. EOS construct a complete information Nash equilibrium of GSP in which bidder 1 bids her value and bidder $i > 1$ bids $\frac{p^{VCG,i-1}}{\alpha_{i-1}}$, where $p^{VCG,i-1}$ is the Vickrey-Clarke-Groves payment of advertiser $i-1$ in the dominant strategy equilibrium where all bidders bid truthfully. Notice that in this equilibrium bidders do not bid their values, and that such manipulations persist with market size (cf. EOS's Remark 3). Additionally, note that it requires each bidder to know enough information about other bidders' values that they can calculate the VCG payments that are an input into their own bids.

Our mechanism constructed according to (5.2) works as follows. First, bidders report their values; call the reported profile $\hat{v} = (\hat{v}_1, \ldots, \hat{v}_n)$, sorted so that $\hat{v}_1 > \hat{v}_2 > \cdots > \hat{v}_n$. Next, execute the EOS complete information Nash equilibrium bids associated with $\hat{v}$; that is, bidder $\hat{v}_1$ bids his value, and bidder $\hat{v}_i$, $i > 1$, bids $\frac{p^{V\hat{C}G,i-1}}{\alpha_{i-1}}$, where $p^{V\hat{C}G,i-1}$ is the VCG payment of bidder $i-1$ in an economy where bidders submit the bids $\hat{v}$. Since the i^{th} highest bidder's total payment is the i^{th} click-through rate α_i multiplied by the $(i+1)^{th}$ bid $\frac{p^{V\hat{C}G,i}}{\alpha_i}$ her total payment is simply $p^{V\hat{C}G,i}$, i.e., our constructed mechanism exactly coincides with VCG.

5.5 Aggregate Uncertainty

Throughout the analysis, we have assumed that agents have an independent prior over opponents' types. This is a restrictive assumption, as in practice we might expect there to be uncertainty in the aggregate about the distribution of types in the economy. In this section we note that additional uncertainty only makes manipulations more difficult to find. If agents were given strictly less information, in the Blackwell sense, then it is still approximately optimal to report truthfully in a SP-L mechanism.

Consider the gain from misreporting for an agent who knows that other agents' actions are distributed according to m. If $((\Phi^n)_{\mathbb{N}}, T)$ is a SP-L mechanism, then this gain must be vanishingly small. That is, given ϵ, t_i, t_i', there exists n_0 such that

$$u_{t_i}[\phi^n(t_i', m)] - u_{t_i}[\phi^n(t_i, m)] \leq \epsilon$$

for all $n \geq n_0$.

Now consider an agent who knows strictly less than agents with any iid beliefs. Following Blackwell, we define a garbling of iid beliefs as a measure $\nu \in \Delta(\Delta A)$. The agent assigns probability $\nu(m)$ that opponents' types are iid according to m.

For an agent with such beliefs, the gain from deviating is

$$\int u_{t_i}[\phi^n(t_i', m)] - u_{t_i}[\phi^n(t_i, m)] d\nu(m).$$

We now show it is approximately optimal for such an agent to report truthfully. Given $\epsilon > 0$, from the definition of SP-L, we know that for each m there must exist $n_0(m)$ such that

$$u_{t_i}[\phi^n(t_i', m)] - u_{t_i}[\phi^n(t_i, m)] < \epsilon/2$$

for all $n \geq n_0(m)$. Now take n_1 such that $n_1 \geq n_0(m)$ for all m in a set $M \subseteq \Delta A$ with measure

$\nu(M)$ at least $1-\epsilon/2$. We then have

$$\begin{aligned}\int u_{t_i}[\phi^n(t_i',m)] - u_{t_i}[\phi^n(t_i,m)]d\nu(m) &= \\ \int_M u_{t_i}[\phi^n(t_i',m)] - u_{t_i}[\phi^n(t_i,m)]d\nu(m) &+ \\ \int_{M^c} u_{t_i}[\phi^n(t_i',m)] - u_{t_i}[\phi^n(t_i,m)]d\nu(m) &< \\ \epsilon/2+\epsilon/2 &= \epsilon.\end{aligned}$$

This result may be formally stated as follows.

Proposition 3. *Consider an SP-L mechanism $\Gamma = ((\Phi^n)_{n\in\mathbb{N}}, T)$. For any garbling ν of iid beliefs over opponents' types, and any $\epsilon > 0$, there exists n_0 such that an agent with beliefs ν in a market of size $n \geq n_0$ cannot gain more than $\epsilon > 0$ by misreporting her type.*

Proposition 3 implies that the SP-L mechanism that we construct to prove Theorem 1 provides approximate incentives in large markets even if agents have aggregate uncertainty.

6 Conclusion

This paper proposes strategyproofness in the large (SP-L) as a second-best alternative to strategyproofness (SP). Our main results show that, while it is well known that SP often severely limits what kinds of mechanisms are possible, there is a sense in which SP-L does not. Specifically, in our class of environments, SP-L is approximately costless to satisfy relative to other forms of incentive compatibility such as Bayes-Nash equilibrium or complete information Nash equilibrium, with the approximation error vanishing to zero in the large-market limit.

We view our results as providing formal justification for focusing on SP-L mechanisms when confronting a new market design problem for which there are no good SP solutions. Additionally, in some environments our proxy-like method of constructing an SP-L mechanism from a given non SP-L mechanism may be of direct use.

We conclude the paper with a few informal arguments in support of SP-L as a desideratum for market design, as well as some caveats.

Empirical Evidence on SP-L There are several empirical studies of mechanisms which are manipulable in the large, and which have been shown to have important incentives problems in practice. These include Jegadeesh (1993) and others on the 1991 pay-as-bid auction scandals, Abdulkadiroğlu et al. (2006, 2009) on the Boston mechanism for school choice, Budish and Cantillon (Forthcoming) on Harvard Business School's course-allocation draft mechanism, Sönmez and Ünver (2010), Krishna and Ünver (2008) and Budish (2010) on the Bidding Points Auction, Edelman and Ostrovsky (2007) on pay-as-bid keyword auctions, Cramton and Katzman (2010) and Merlob et al. (2010) on a proposed Medicare auction for durable equipment, Roth (2002) and others on non-stable matching algorithms such as the priority match, and potentially others. By contrast, to the best of our knowledge, there are no empirical examples of market designs that are SP-L but which have been shown to be harmfully manipulated in large finite markets.

To the extent that this pattern is indeed true, it suggests that perhaps the relevant distinction for practice, in contexts with a large number of participants, is not "SP vs. not SP", but rather "SP-L vs. not SP-L." Or, more conservatively, "SP vs. SP-L vs. not SP-L."

Several Arguments for SP Design are also Arguments for SP-L Design In traditional mechanism design, incentives are viewed as a constraint, not an objective. A number of recent papers in the market design literature have suggested, either formally or informally, that strategyproofness be viewed as an explicit design objective. Many of these arguments can be interpreted as supporting SP-L design as well.

One such argument is that strategyproof mechanisms eliminate any unmodeled costs of calculating an optimal response; e.g., Roth (2008) argues that good markets are "sufficiently simple to participate in" and make it "safe to participate straightforwardly". Any SP-L mechanism has the following property: for any conjecture m about the distribution of opponents' play, and any cost $c > 0$ associated with calculating an optimal response, there exists n_0 such that in markets with $n > n_0$ participants, each agent maximizes her expected utility by simply reporting her preferences truthfully, and avoiding the cost c of strategizing. While Roth (2008) never defines "safe" formally, we note that some mechanisms that are not SP but that are SP-L, such as deferred acceptance, are classified as safe.

A second such argument is that strategyproof mechanisms are fair, in the sense that they do not penalize participants who are strategically unsophisticated (Abdulkadiroğlu et al. (2006); Pathak

and Sönmez (2008)). By an analogous argument to that in the previous paragraph, in a large enough market SP-L mechanisms are approximately fair, in the sense that the cost of being strategically unsophisticated can be bounded above by c.

SP-L Mechanisms Shift Complexity from Humans to Computers In mechanisms such as the pay-as-bid auction described in Example 1, or the Boston mechanism described in Example 2, optimal play requires a participant first to form a belief over how others are going to play the mechanism, and then to calculate what action maximizes their expected utility given how others will play. Our construction in Section 4 shows that, in a large market, this strategic complexity can be shifted from the players to a computer: players just report their types, and then the center computes and executes the equilibrium associated with the empirical distribution of reports.[19] Note, importantly, that the computer need not know the prior, only the empirical; in that sense, this use of computers is consistent with the Wilson (1987) doctrine.

Caveats

We conclude with two important caveats on SP-L. First, there is no simple answer to the question of how large a market is large enough to ignore vanishing deviations.[20] We view the limit representation of a mechanism as a useful if imperfect abstraction for many interesting market design problems, just as the assumption of price-taking behavior is a useful abstraction in other parts of economics. In any specific context, the analyst's case for using an SP-L mechanism can be strengthened with empirical, experimental, or computational evidence suggesting that the gains from misreporting are small and/or rare; see, for instance, Roth and Peranson (1999).

Our second caveat relates to the difficulty of determining and reporting one's type. Many of the mechanisms that we have criticized as being manipulable in the large, and hence strategically complicated for participants, have the virtue that their message spaces are quite simple. For instance, in the Boston mechanism, it may be unrealistic to expect that a student will be able to accurately

[19] A recent literature in computer science has emphasized that finding equilibria can be computationally complex (cf. Roughgarden (2010)). Hence, a reader may worry that our approach is computationally unrealistic. Our basic reply to this concern is that, to the extent that finding an equilibrium is difficult for a computer, we should be worried that finding an equilibrium is difficult for humans as well. Thus, while we agree that computational complexity is an important concern, we submit that it is not a compelling reason to prefer Bayes-Nash mechanisms to SP-L alternatives.

[20] Indeed, even in analyses of the convergence properties of specific mechanisms, rarely is the analysis sufficient to answer the question of, e.g., "is 1000 participants large?" Convergence is often slow, or includes a large constant term. A notable exception is Rustichini et al. (1994), who are able to show, in the context of a double auction with unit demand and uniformly distributed values, that 6 buyers and sellers is large enough to approximate efficiency to within one percent.

estimate the equilibrium p_j^*'s (cf. Example 2), but it seems realistic to expect that a student could determine which school to ask for as her first choice *given* the p_j^*'s. SP-L mechanisms are strategically simple, modulo the caveat above, but require agents to report a potentially unrealistic amount of information about their preferences: for instance, in the SP-L mechanism we construct based on the Boston mechanism's equilibria (cf. Section 4.3), students' reports of their types consist of their von Neumann-Morgenstern utilities for each possible school, including schools to which they are highly unlikely to be assigned. There is a burgeoning literature in computer science that asks how to design communication systems that allow agents to efficiently report their preferences (cf. Sandholm and Boutilier (2006)), as well as a literature at the intersection of economics and computer science on the design of concise message spaces (cf. Milgrom (2009)). An interesting question that we leave for future research is how to define a useful second-best alternative to strategyproofness when reporting one's type is unrealistic.

A Proof of Theorem 1

As described in (4.1), let

$$F^n(t) = \Phi^n(\sigma_{\text{emp}[t]}(t))$$

and define $f^n(t_i, m)$ from $F^n(t_i, t_{-i})$ as in (2.1). The core of the proof is the following approximation result.

Lemma 1. *Fix a prior μ_0 and $\epsilon > 0$. Let $\mathcal{N}$ be a neighborhood as in Definition 7. Let μ_k be priors $\mu_k \in \mathcal{A}_k$ for each $k = 1, \ldots, K$, with $|\mu_k - \mu_0| < \epsilon$. Then there exists n_0, and positive weights π_k^n with $\sum_{1 \leq k \leq K} \pi_k^n = 1$, such that for all t_i*

$$|f^n(t_i, \mu_0) - \sum_{k=1}^{K} \pi_k^n \cdot z_k(t_i)| < 6\epsilon,$$

where

$$z_k(t_i) = \phi^\infty(\sigma_{\mu_k}(t_i), \sigma_{\mu_k}(\mu_k)).$$

The Lemma states that the bundle received by an agent playing t_i in the mechanism Γ^D can be approximated by a convex combination of the bundles received when playing the original equilibrium within each region $\mathcal{A}_k$. Each $z_k(t_i)$ is defined as the bundle an agent receives when playing t_i when opponents' types and the prior on which equilibrium is selected are each within $\mathcal{A}_k$. The key assertion that the approximation Lemma makes is that the π_k^n do not depend on t_i. That is, irrespective of the type an agent reports, the approximation weights can be taken to be the same. This reflects the fact that a single agent has a very small effect on the probability of the distribution of types falling within each region $\mathcal{A}_k$.

We now prove the Lemma, and then use it to prove Theorem 1.

Proof of Lemma 1.

The proof of the Lemma involves three steps. Throughout the proof we use the shorthand $\hat{\mu} = \text{emp}[t_i, t_{-i}]$. The first step shows that the appoximation formula holds within each region $\mathcal{A}_k$.

Step 1.

There exists n_0 such that, for all $n > n_0$ and $t \in \mathcal{A}_k$ we have

$$|\Phi^n(\sigma_{\hat{\mu}}(t_i), \sigma_{\hat{\mu}}(t_{-i})) - \phi^\infty(\sigma_{\mu_k}(t_i), \sigma_{\mu_k}(\mu_k))| < 4\epsilon.$$

Using the $z_k(t_i)$ notation, this is

$$|\Phi^n(\sigma_{\hat{\mu}}(t_i), \sigma_{\hat{\mu}}(t_{-i}) - z_k(t_i)| < 4\epsilon.$$

Proof. First note that, by Condition 3 of Definition 7 we may take n_1 such that for $n \geq n_1$

$$|\Phi^n(\sigma_{\hat{\mu}}(t_i), \sigma_{\hat{\mu}}(t_{-i})) - \Phi^n(\sigma_{\mu_k}(t_i), \sigma_{\mu_k}(t_{-i}))| < \epsilon. \tag{A.1}$$

Note that the left term $\Phi^n(\sigma_{\hat{\mu}}(t_i), \sigma_{\hat{\mu}}(t_{-i}))$ is the term whose distance to $z_k(t_i)$ we wish to bound. We will do so by showing that $\Phi^n(\sigma_{\mu_k}(t_i), \sigma_{\mu_k}(t_{-i}))$ is close to $\phi^n(\sigma_{\mu_k}(t_i), \sigma_{\mu_k}(\mu_k))$, and then showing that $\phi^n(\sigma_{\mu_k}(t_i), \sigma_{\mu_k}(\mu_k))$ is close to $z_k(t_i)$.

By definition we have that

$$\phi^n(\sigma_{\mu_k}(t_i), \sigma_{\mu_k}(\mu_k)) = \sum_{t'_{-i}} \Pr(t'_{-i}|t'_{-i} \sim \mu_k) \cdot \Phi^n(\sigma_{\mu_k}(t_i), \sigma_{\mu_k}(t'_{-i})). \tag{A.2}$$

We will now bound the distance between $\phi^n(\sigma_{\mu_k}(t_i), \sigma_{\mu_k}(\mu_k))$ and $\Phi^n(\sigma_{\mu_k}(t_i), \sigma_{\mu_k}(t_{-i}))$. For all $t \in \mathcal{A}_k$ we have

$$\begin{aligned}
& |\Phi^n(\sigma_{\mu_k}(t_i), \sigma_{\mu_k}(t_{-i})) - \phi^n(\sigma_{\mu_k}(t_i), \sigma_{\mu_k}(\mu_k))| \\
= \;& |\Phi^n(\sigma_{\mu_k}(t_i), \sigma_{\mu_k}(t_{-i})) - \textstyle\sum_{t_{-i}} \Pr(t'_{-i}|t'_{-i} \sim \mu_k) \cdot \Phi^n(\sigma_{\mu_k}(t_i), \sigma_{\mu_k}(t'_{-i}))| \\
\leq \;& \textstyle\sum_{t'_{-i}} \Pr(t'_{-i}|t'_{-i} \sim \mu_k) \cdot |\Phi^n(\sigma_{\mu_k}(t_i), \sigma_{\mu_k}(t_{-i})) - \Phi^n(\sigma_{\mu_k}(t_i), \sigma_{\mu_k}(t'_{-i}))| \\
= \;& \textstyle\sum_{t'_{-i}:\text{emp}[t_i,t'_{-i}]\in\mathcal{A}_k} \Pr(t'_{-i}|t'_{-i} \sim \mu_k) \cdot |\Phi^n(\sigma_{\mu_k}(t_i), \sigma_{\mu_k}(t_{-i})) - \Phi^n(\sigma_{\mu_k}(t_i), \sigma_{\mu_k}(t'_{-i}))| \\
+ \;& \textstyle\sum_{t'_{-i}:\text{emp}[t_i,t'_{-i}]\notin\mathcal{A}_k} \Pr(t'_{-i}|t'_{-i} \sim \mu_k) \cdot |\Phi^n(\sigma_{\mu_k}(t_i), \sigma_{\mu_k}(t_{-i})) - \Phi^n(\sigma_{\mu_k}(t_i), \sigma_{\mu_k}(t'_{-i}))|.
\end{aligned} \tag{A.3}$$

The first equality follows by substituting the definition of ϕ^n from Equation (A.2). The inequality follows from the triangle inequality and the fact that the probabilities must sum to 1. The last equality simply breaks the sum into two parts, the t'_{-i} such that for which $\text{emp}[t_i, t'_{-i}]$ is in $\mathcal{A}_k$, and the ones for which it is not. Consider now the expression on the right side of the last equality. Note that we may take n_1 such that the first term is bounded by

$$\sum_{t'_{-i}:\text{emp}[t_i,t'_{-i}]\in\mathcal{A}_k} \Pr(t'_{-i}|t'_{-i} \sim \mu_k) \cdot |\Phi^n(\sigma_{\mu_k}(t_i), \sigma_{\mu_k}(t_{-i})) - \Phi^n(\sigma_{\mu_k}(t_i), \sigma_{\mu_k}(t'_{-i}))| < \epsilon,$$

which follows from Condition 3 in Definition 7. As for the second term, by the law of large numbers, we may take n_2 large enough such that the total probability mass that $\text{emp}[t_i, t'_{-i}] \in \mathcal{A}_k$ is greater than $1 - \epsilon$. This bounds the second term by

$$\sum_{t'_{-i}:\text{emp}[t_i,t'_{-i}]\notin\mathcal{A}_k} \Pr(t'_{-i}|t'_{-i} \sim \mu_k) \cdot |\Phi^n(\sigma_{\mu_k}(t_i), \sigma_{\mu_k}(t_{-i})) - \Phi^n(\sigma_{\mu_k}(t_i), \sigma_{\mu_k}(t'_{-i})) < \epsilon,$$

Substituting these bounds in inequality (A.3) then yields

$$|\Phi^n(\sigma_{\mu_k}(t_i), \sigma_{\mu_k}(t_{-i})) - \phi^n(\sigma_{\mu_k}(t_i), \sigma_{\mu_k}(\mu_k))| < \epsilon + \epsilon = 2\epsilon. \tag{A.4}$$

Finally, by the definition of the limit we may take n_3 such that for all $n > n_3$

$$|\phi^n(\sigma_{\mu_k}(t_i), \sigma_{\mu_k}(\mu_k)) - \phi^\infty(\sigma_{\mu_k}(t_i), \sigma_{\mu_k}(\mu_k))| < \epsilon. \tag{A.5}$$

If we take $n_0 = \max\{n_1, n_2, n_3\}$ the Lemma then follows from Inequalities (A.1), (A.4) and (A.5). □

The next step shows that the probability that a vector (t_i, t_{-i}) falls within region $\mathcal{A}_k$, when t_{-i} is distributed randomly, does not vary too much with t_i in large markets. This is a key step in our argument, as it says an individual agent cannot appreciably change the probability that t falls within each $\mathcal{A}_k$, and therefore cannot have a large effect on the aggregate allocation.

Step 2.

There exists n_0 such that, for all $n > n_0$ there exist weights $\pi_1^n, \ldots, \pi_K^n$ such that $\sum_k \pi_k^n = 1$ and

$$|\Pr((t_i, t_{-i}) \in \mathcal{A}_k|t_{-i} \sim \mu_0) - \pi_k^n| < \epsilon/K$$

for all k and all t_i.

Proof. We begin by constructing numbers which will be approximately equal to the π_k^n in the statement of this step. Let

$$\bar{\pi}_k^n = \Pr(t' \in \mathcal{A}_k|t' \in T^n, t' \sim \mu_k)$$

be the probability that a vector of n types drawn independently according to μ_k is in $\mathcal{A}_k$. We will show that for large n these $\bar{\pi}_k^n$ are very close to the probabilities $\Pr\{(t_i, t_{-i}) \in \mathcal{A}_k|t_{-i} \sim \mu_0\}$. For

any type t_i, the difference between the probability of a vector of types falling within region $\mathcal{A}_k$ when i's type is fixed as t_i, versus when i's type is drawn randomly, is

$$\begin{aligned}\Pr((t_i, t'_{-i}) \in \mathcal{A}_k | t'_{-i} \in T^{n-1}, t'_{-i} \sim \mu_0) - \bar{\pi}^n_k \quad = \\ \Pr((t_i, t'_{-i}) \in \mathcal{A}_k, t' \notin \mathcal{A}_k | t' \in T^n, t' \sim \mu_k) \\ - \Pr((t_i, t'_{-i}) \notin \mathcal{A}_k, t' \in \mathcal{A}_k | t' \in T^n, t' \sim \mu_k). \end{aligned} \tag{A.6}$$

This is just the difference between the probability of choosing a vector t' where changing a single type (i's) from t'_i to t_i moves the vector of types from inside $\mathcal{A}_k$ to outside $\mathcal{A}_k$, and the probability of choosing a vector where changing i's type from t'_i to t_i moves the vector from outside $\mathcal{A}_k$ to inside $\mathcal{A}_k$. We now show that the probability of such vectors being drawn can be taken to be very small.

Consider the case where $(t_i, t'_{-i}) \notin \mathcal{A}_k$, but $(t'_i, t'_{-i}) \in \mathcal{A}_k$. One possibility is that $(t_i, t'_{-i}) \notin \mathcal{N}$. By the law of large numbers, we may take n_0 large enough such that for $n > n_0$ the probability of this happening is less than $\epsilon/8$. The other possibility is that $(t_i, t'_{-i}) \in \mathcal{N}$, but $(t_i, t'_{-i}) \notin \mathcal{A}_k$. In that case, the segment $[(t_i, t'_{-i}), t']$ must have a point that lies in $\mathcal{B}$, as we assumed $\mathcal{N}$ to be convex. This means that the distance between t' and $\mathcal{B}$ is at most $1/n$. By Condition 2 of Definition 7, we may take n_0 such that this probability is less than $\epsilon/8$. This argument then yields that

$$\Pr((t_i, t'_{-i}) \notin \mathcal{A}_k, t' \in \mathcal{A}_k | t' \in T^n, t' \sim \mu_k) < \epsilon/8 + \epsilon/8 = \epsilon/4.$$

An analogous argument proves that we may assume that for $n > n_0$

$$\Pr((t_i, t'_{-i}) \notin \mathcal{A}_k, t' \in \mathcal{A}_k | t' \in T^n, t' \sim \mu_k) < \epsilon/4.$$

Substituting these two inequalities in Equation (A.6) yields that

$$|\Pr((t_i, t'_{-i}) \in \mathcal{A}_k | t'_{-i} \in T^{n-1}, t'_{-i} \sim \mu_0) - \bar{\pi}^n_k| < \epsilon/4 + \epsilon/4 = \epsilon/2. \tag{A.7}$$

Note, however, that the $\bar{\pi}^n_k$ do not necessarily sum to 1, as it may the the case that $t' \notin \cup_k \mathcal{A}_k$. To complete the proof, we define

$$\pi^n_k = \bar{\pi}^n_k / \sum_{k'} \bar{\pi}^n_{k'}. \tag{A.8}$$

We have that the probability that $t' \notin \cup_k \mathcal{A}_k$ converges to 0. Therefore, we may take n_0 such that for $n > n_0$

$$|1 - 1/\sum_{k'} \bar{\pi}^n_{k'}| < \epsilon/2. \tag{A.9}$$

Putting this together, we may finish the proof of Step 2.

$$\begin{aligned}
|\Pr((t_i, t'_{-i}) \in \mathcal{A}_k | t'_{-i} \in T^{n-1}, t'_{-i} \sim \mu_0) - \pi^n_k| &\leq \\
|\Pr((t_i, t'_{-i}) \in \mathcal{A}_k | t'_{-i} \in T^{n-1}, t'_{-i} \sim \mu_0) - \bar{\pi}^n_k| + |\pi^n_k - \bar{\pi}^n_k| &< \\
\epsilon/2 + |\pi^n_k - \bar{\pi}^n_k| &= \\
\epsilon/2 + |\bar{\pi}^n_k / \sum_{k'} \bar{\pi}^n_{k'} - \bar{\pi}^n_k| &= \\
\epsilon/2 + |1 - 1/\sum_{k'} \bar{\pi}^n_{k'}| \cdot |\bar{\pi}^n_k| &< \\
\epsilon/2 + \epsilon/2 &< \epsilon.
\end{aligned}$$

The series of steps in the above derivation were as follows. The second line derives from the triangle inequality. The third line uses the bound from Inequality (A.7). The fourth line uses the definition of π^n_k from equation (A.8). Finally, the fifth line is algebra, and the sixth line comes from the bound in inequality (A.9). □

Step 3.

Finally, we apply the results from steps 1 and 2 to prove the Lemma, obtaining the desired approximation formula. That is, there exists n_0 such that for all $n \geq n_0$

$$|f^n(t_i, \mu_0) - \sum_{k=1}^{K} \pi^n_k \cdot z_k(t_i)| < 6\epsilon.$$

Proof. We may write

$$f^n(t_i, \mu_0) - \sum_k \pi^n_k \cdot z_k(t_i) = \sum_{t_{-i}} \Pr(t_{-i} | t_{-i} \sim \mu_0) \cdot \Phi^n(\sigma_{\hat{\mu}}(t_i), \sigma_{\hat{\mu}}(t_{-i})) - \sum_k \pi^n_k \cdot z_k(t_i).$$

This sum can be decomposed depending on whether $\hat{\mu}$ is in each of the $\mathcal{A}_k$ sets or not. We have

$$\begin{aligned} f^n(t_i,\mu_0) \quad -\textstyle\sum_k \pi_k^n \cdot z_k(t_i) = \quad & \sum_k \Big(\sum_{t_{-i}:\hat{\mu}\in\mathcal{A}_k} \Pr(t_{-i}|t_{-i}\sim\mu_0)\cdot\Phi^n(\sigma_{\hat{\mu}}(t_i),\sigma_{\hat{\mu}}(t_{-i})) - \pi_k^n\cdot z_k^n(t_i)\Big) \\ & + \sum_{t_{-i}:\hat{\mu}\notin\cup_k\mathcal{A}_k} \Pr(t_{-i}|t_{-i}\sim\mu_0)\cdot\Phi^n(\sigma_{\hat{\mu}}(t_i),\sigma_{\hat{\mu}}(t_{-i})). \end{aligned} \tag{A.10}$$

We begin by looking at the terms where $\hat{\mu}$ is in one of the $\mathcal{A}_k$. We will show that for each k these terms are small. We have that, for each k,

$$\begin{aligned} & \Big| \sum_{t_{-i}:\hat{\mu}\in\mathcal{A}_k} \Pr(t_{-i}|t_{-i}\sim\mu_0)\cdot\Phi^n(\sigma_{\hat{\mu}}(t_i),\sigma_{\hat{\mu}}(t_{-i})) - \pi_k^n\cdot z_k(t_i)\Big| \\ & \quad \le \sum_{t_{-i}:\hat{\mu}\in\mathcal{A}_k} \Pr(t_{-i}|t_{-i}\sim\mu_0)\cdot|\Phi^n(\sigma_{\hat{\mu}}(t_i),\sigma_{\hat{\mu}}(t_{-i})) - z_k(t_i)| \\ & \qquad + \Big|\sum_{t_{-i}:\hat{\mu}\in\mathcal{A}_k} \Pr(t_{-i}|t_{-i}\sim\mu_0) - \pi_k^n\Big|\cdot|z_k(t_i)| \\ \le & \max_{t_{-i}:\hat{\mu}\in\mathcal{A}_k} |\Phi^n(\sigma_{\hat{\mu}}(t_i),\sigma_{\hat{\mu}}(t_{-i})) - z_k(t_i)| \cdot \sum_{t_{-i}:\hat{\mu}\in\mathcal{A}_k} \Pr(t_{-i}|t_{-i}\sim\mu_0) \\ & \qquad + \Big|\sum_{t_{-i}:\hat{\mu}\in\mathcal{A}_k} \Pr(t_{-i}|t_{-i}\sim\mu_0) - \pi_k^n\Big| \end{aligned} \tag{A.11}$$

The first inequality follows from the triangle inequality. The second inequality bounds each term $|\Phi^n(\sigma_{\hat{\mu}}(t_i),\sigma_{\hat{\mu}}(t_{-i})) - z_k^n(t_i)|$ by the maximum value of these terms, and it bounds $|z_k(t_i)|$ by 1.

Consider now the right side of the last inequality. By step 1, we may take n_0 such that for all $n \ge n_0$,

$$\max_{t_{-i}:\hat{\mu}\in\mathcal{A}_k} |\Phi^n(\sigma_{\hat{\mu}}(t_i),\sigma_{\hat{\mu}}(t_{-i})) - z_k(t_i)| < 4\epsilon.$$

By step 2, for $n \ge n_0$ for a suitable n_0 the second term is bounded by

$$\Big| \sum_{t_{-i}:\hat{\mu}\in\mathcal{A}_k} \Pr(t_{-i}|t_{-i}\sim\mu_0) - \pi_k^n\Big| < \frac{\epsilon}{K}.$$

Substituting these two bounds in inequality (A.11) we have that for all $n \ge n_0$

$$\begin{aligned} & \big|\textstyle\sum_{t_{-i}:\hat{\mu}\in\mathcal{A}_k} \Pr(t_{-i}|t_{-i}\sim\mu_0)\cdot\Phi^n(\sigma_{\hat{\mu}}(t_i),\sigma_{\hat{\mu}}(t_{-i})) - \pi_k^n\cdot z_k^n(t_i)\big| \\ \le & \quad 3\epsilon\cdot\textstyle\sum_{t_{-i}:\hat{\mu}\in\mathcal{A}_k} \Pr(t_{-i}|t_{-i}\sim\mu_0) + \frac{\epsilon}{K}. \end{aligned}$$

Summing over all k we get

$$\sum_k |\sum_{t_{-i}:\hat{\mu}\in\mathcal{A}_k} \Pr(t_{-i}|t_{-i}\sim\mu_0)\cdot\Phi^n(\sigma_{\hat{\mu}}(t_i),\sigma_{\hat{\mu}}(t_{-i})) - \pi_k^n\cdot z_k^n(t_i)| \leq 5\epsilon$$

and then using the triangle inequality we can bring the summation inside to get

$$|\sum_k \sum_{t_{-i}:\hat{\mu}\in\mathcal{A}_k} \Pr(t_{-i}|t_{-i}\sim\mu_0)\cdot\Phi^n(\sigma_{\hat{\mu}}(t_i),\sigma_{\hat{\mu}}(t_{-i})) - \pi_k^n\cdot z_k^n(t_i)| \leq 5\epsilon. \tag{A.12}$$

The argument above bounds the terms in equation (A.10) that correspond to t within the sets $\mathcal{A}_k$. To bound the other term, note that we may take n_0 to be large enough so that for all $n \geq n_0$ the total probability that $t \notin \cup_k \mathcal{A}_k$ is strictly less than ϵ. That is,

$$\sum_{t_{-i}:\hat{\mu}\notin\cup_k\mathcal{A}_k} \Pr(t_{-i}|t_{-i}\sim\mu_0) < \epsilon. \tag{A.13}$$

Plugging in equations (A.12) and (A.13) in equation (A.10) we obtain

$$|f^n(t_i,\mu_0) - \sum_{k=1}^K \pi_k^n\cdot z_k(t_i)| < 6\epsilon,$$

completing the proof of Step 3, and hence the Lemma. □

With the Lemma in hand, it is a simple matter to establish Theorem 1.

Proof. (Theorem 1)

Part 1.

To see that Γ^D is SP-L, consider the gain for type t_i from deviating to $\hat{t}_i$ when opponents play μ_0. That is,

$$u_{t_i}[f^n(\hat{t}_i,\mu_0)] - u_{t_i}[f^n(t_i,\mu_0)].$$

By the approximation Lemma, and the boundedness of u, given $\epsilon > 0$, there exists n_0, π_k^n, μ_k, and

z_k as in the statement of the Lemma such that, for all $n > n_0$:

$$|u_{t_i}[f^n(\hat{t}_i, \mu_0)] - \sum_k \pi_k^n \cdot u_{t_i}[z_k(\hat{t}_i)]| \quad < \quad \epsilon/2 \tag{A.14}$$

$$|u_{t_i}[f^n(t_i, \mu_0)] - \sum_k \pi_k^n \cdot u_{t_i}[z_k(t_i)]| \quad < \quad \epsilon/2. \tag{A.15}$$

Also, by the definition of $z_k(\cdot)$, we have that

$$u_{t_i}[z_k(t_i)] \geq u_{t_i}[z_k(\hat{t}_i)]. \tag{A.16}$$

Therefore, we may bound the gain from deviating for $n > n_0$ by

$$\begin{aligned} u_{t_i}[f^n(\hat{t}_i, \mu_0)] - u_{t_i}[f(t_i, \mu_0)] \quad &\leq \\ \sum_k \pi_k^n \cdot \{u_{t_i}[z_k(\hat{t}_i)] - u_{t_i}[z_k(t_i)]\} & \\ + |u_{t_i}[f^n(\hat{t}_i, \mu_0)] - \sum_k \pi_k^n \cdot u_{t_i}[z_k(\hat{t}_i)]| & \\ + |u_{t_i}[f^n(t_i, \mu_0)] - \sum_k \pi_k^n \cdot u_{t_i}[z_k(t_i)]| \quad &< \\ 0 + \epsilon/2 + \epsilon/2 \quad &= \quad \epsilon. \end{aligned}$$

The first inequality follows from the triangle inequality, and the second inequality from the bounds in inequalities (A.14), (A.15), (A.16).

Part 3.

Part 3 follows from the approximation Lemma. Given $\mu_0 \in \Delta T, \epsilon > 0$, by the Lemma we may take n_0, μ_k such that all $|\mu_k - \mu| < \epsilon$, and for all $n \geq n_0$

$$|f^n(t_i, \mu_0) - \sum_k \pi_k^n \cdot \phi^\infty(\sigma_{\mu_k}(t_i), \sigma_{\mu_k}(\mu_k))| < \epsilon/2. \tag{A.17}$$

By the definition of the limit, we may take n_0 such that for all k, t_i, and $n > n_0$

$$|\phi^\infty(\sigma_{\mu_k}(t_i), \sigma_{\mu_k}(\mu_k)) - \phi^n(\sigma_{\mu_k}(t_i), \sigma_{\mu_k}(\mu_k))| < \epsilon/2. \tag{A.18}$$

By the triangle inequality and inequalities (A.17) and (A.18) we have that

$$|f^n(t_i, \mu_0) - \sum_k \pi_k^n \cdot \phi^n(\sigma_{\mu_k}(t_i), \sigma_{\mu_k}(\mu_k))| < \epsilon/2 + \epsilon/2 = \epsilon. \tag{A.19}$$

The result then follows from the fact that u is bounded above by 1 and takes the expected utility form.

Part 2.

Finally for Part 2, note that we may take $\mathcal{N} = \mathcal{A}_1$ and $\mu_1 = \mu_0$ in the continuous case. Therefore $\pi_k^n = 1$, and equation (A.19) becomes

$$|f^n(t_i, \mu_0) - \phi^n(\sigma_{\mu_0}(t_i), \sigma_{\mu_0}(\mu_0)| < \epsilon.$$

The desired formula then follows from the boundedness of u as above. □

B Other Omitted proofs

B.1 Proof of Proposition 1

Step 1.

Given $\epsilon > 0$, there exists n_0 such that for all $n, n' \geq n_0, t_i, t_{-i}$ we have

$$|\Phi^n(t_i, t_{-i}) - \phi^{n'}(t_i, \operatorname{emp} t_{-i})| < \epsilon.$$

Proof. Let $\hat{\mu} = \operatorname{emp} t_{-i}$. We may write

$$\phi^{n'}(t_i, \hat{\mu}) = \sum_{t'_{-i}} \Pr(t'_{-i} | t'_{-i} \in T^{n'-1}, t'_{-i} \sim \hat{\mu}) \cdot \Phi^{n'}(t_i, t'_{-i}). \tag{B.1}$$

By the definition of equicontinuity, we may take $\delta > 0$ such that for all t'_{-i} with

$$|\operatorname{emp} t'_{-i} - \hat{\mu}| < \delta$$

we have

$$|\Phi^n(t_i, t_{-i}) - \Phi^{n'}(t_i, t'_{-i})| < \epsilon/2. \tag{B.2}$$

Moreover, we may take n_0 such that for all $n \geq n_0$, by the law of large numbers,

$$\sum_{|\operatorname{emp} t'_{-i} - \hat{\mu}| \geq \delta, t'_{-i} \in T^{n'-1}} \Pr(t'_{-i} | t'_{-i} \in T^{n'-1}, t'_{-i} \sim \hat{\mu}) < \epsilon/2. \tag{B.3}$$

Consider now the difference

$$|\Phi^n(t_i, t_{-i}) - \phi^{n'}(t_i, \hat{\mu})|.$$

From Equation (B.1), we have that

$$\begin{aligned} |\Phi^n(t_i, t_{-i}) - \phi^{n'}(t_i, \hat{\mu})| &= \\ |\Phi^n(t_i, t_{-i}) - \sum_{t'_{-i}} \Pr(t'_{-i} | t'_{-i} \in T^{n'-1}, t'_{-i} \sim \hat{\mu}) \cdot \Phi^{n'}(t_i, t'_{-i})|. \end{aligned}$$

By the triangle inequality we have that

$$\begin{aligned} & |\Phi^n(t_i, t_{-i}) - \phi^{n'}(t_i, \hat{\mu})| \\ \leq & \sum_{|\operatorname{emp} t'_{-i} - \hat{\mu}| < \delta, t'_{-i} \in T^{n'-1}} \Pr(t'_{-i} | t'_{-i} \in T^{n'-1}, t'_{-i} \sim \hat{\mu}) \cdot |\Phi^n(t_i, t_{-i}) - \Phi^{n'}(t_i, t'_{-i})| \\ + & \sum_{|\operatorname{emp} t'_{-i} - \hat{\mu}| \geq \delta, t'_{-i} \in T^{n'-1}} \Pr(t'_{-i} | t'_{-i} \in T^{n'-1}, t'_{-i} \sim \hat{\mu}) \cdot |\Phi^n(t_i, t_{-i}) - \Phi^{n'}(t_i, t'_{-i})|. \end{aligned}$$

Plugging in Inequalities (B.2) and (B.3) we have

$$\begin{aligned} |\Phi^n(t_i, t_{-i}) - \phi^{n'}(t_i, \hat{\mu})| &< \\ \epsilon/2 + \epsilon/2 &= \epsilon. \end{aligned}$$

Moreover, note that the above bounds in Inequalities (B.2) and (B.3) may be taken uniform in t_i, t_{-i}. Therefore the overall bound is uniform. This completes this step. □

Step 2.

Given $\epsilon > 0$, there exists n_0 such that for all $n \geq n_0, t_i, t_{-i}$ we have

$$|\Phi^n(t_i, t_{-i}) - \phi^\infty(t_i, \operatorname{emp} t_{-i})| < \epsilon. \tag{B.4}$$

Proof. By Step 1, we may take n_0 such that for all $n, n' \geq n_0, t_i, t_{-i}$ we have

$$|\Phi^n(t_i, t_{-i}) - \phi^{n'}(t_i, \operatorname{emp} t_{-i})| < \epsilon/2.$$

Taking the limit as $n' \to \infty$ we have

$$|\Phi^n(t_i, t_{-i}) - \phi^\infty(t_i, \operatorname{emp} t_{-i})| \leq \epsilon/2 < \epsilon.$$

□

Step 3.

Use Step 2 to complete the proof of the Proposition.

Proof. By Step 2, we may take n_0 large enough such that for all $n \geq n_0, t$:

$$|\Phi^n(t_i, t_{-i}) - \phi^\infty(t_i, \operatorname{emp} t_{-i})| < \epsilon/2. \tag{B.5}$$

Consider now the gain for agent i to, given t_{-i}, perfom an ex post deviation to $\hat{t}_i$. We have

$$\begin{aligned} u_{t_i}[\Phi^n(\hat{t}_i, t_{-i})] - u_{t_i}[\Phi^n(t)] &\leq |u_{t_i}[\Phi^n(\hat{t}_i, t_{-i})] - u_{t_i}[\phi^\infty(\hat{t}_i, \operatorname{emp} t_{-i})]| \\ &+ |u_{t_i}[\phi^\infty(t_i, \operatorname{emp} t_{-i})] - u_{t_i}[\Phi^n(t_i, t_{-i})]| \\ &+ (u_{t_i}[\phi^\infty(\hat{t}_i, \operatorname{emp} t_{-i})] - u_{t_i}[\phi^\infty(t_i, \operatorname{emp} t_{-i})]). \end{aligned}$$

By the boundedness of u we have

$$\begin{aligned} u_{t_i}[\Phi^n(\hat{t}_i, t_{-i})] - u_{t_i}[\Phi^n(t)] &\leq |\Phi^n(\hat{t}_i, t_{-i}) - \phi^\infty(\hat{t}_i, \operatorname{emp} t_{-i})| \\ &+ |\phi^\infty(t_i, \operatorname{emp} t_{-i}) - \Phi^n(t_i, t_{-i})| \\ &+ (u_{t_i}[\phi^\infty(\hat{t}_i, \operatorname{emp} t_{-i})] - u_{t_i}[\phi^\infty(t_i, \operatorname{emp} t_{-i})]). \end{aligned}$$

Plugging in Inequality (B.5), we have

$$\begin{aligned} u_{t_i}[\Phi^n(\hat{t}_i, t_{-i})] - u_{t_i}[\Phi^n(t)] &< \epsilon/2 + \epsilon/2 \\ &+ (u_{t_i}[\phi^\infty(\hat{t}_i, \operatorname{emp} t_{-i})] - u_{t_i}[\phi^\infty(t_i, \operatorname{emp} t_{-i})]). \end{aligned}$$

Since Γ is SP-L, we have that the last term is nonpositive. Therefore,

$$u_{t_i}[\Phi^n(\hat{t}_i, t_{-i})] - u_{t_i}[\Phi^n(t)] < \epsilon.$$

This completes the proof. □

B.2 Finite-Economy Bayes-Nash Equilibria

As described in Section 5.4, the main Theorem could have been stated using exact Nash equilibrium. Following the definition of a quasicontinuous sequence of families of equilibria, the Theorem statement is as follows.

Theorem (Alternative Statement of Theorem 1)**.** *Consider a mechanism $\Gamma = ((\Phi^n)_{n\in\mathbb{N}}, A)$ with a quasicontinuous sequence of families of limit equilibria $(\sigma^n_\mu)_{\mu\in\Delta T, n\in\mathbb{N}\cup\{\infty\}}$. Then there exists a direct mechanism $\Gamma^D = ((F^n)_{\mathbb{N}}, T)$ with the following properties*

1. *Γ^D is strategyproof in the large.*
2. *If Γ^D is continuous at the true prior μ_0, then in the limit as $n \to \infty$, truthful play of Γ^D and Bayes-Nash equilibrium play of Γ give agents the same utilities. Formally, given $\mu_0 \in \Delta T$ and $\epsilon > 0$, there exists n_0 such that for all $n > n_0$ and all t_i:*

$$|u_{t_i}[f^n(t_i, \mu_0)] - u_{t_i}[\phi^n(\sigma^n_{\mu_0}(t_i), \sigma^n_{\mu_0}(\mu_0))]| < \epsilon.$$

3. *If Γ^D is not continuous at μ_0, then in the limit as $n \to \infty$, truthful play of Γ^D gives agents the same utilities as a convex combination of equilibrium outcomes under Γ and priors in a neighborhood of μ_0. Formally, for every $\mu_0 \in \Delta T$, $\epsilon > 0$ there exist priors μ_k with $|\operatorname{emp}[\mu_k] - \operatorname{emp}[\mu_0]| < \epsilon$, and n_0, such that for all $n > n_0$ there are weights π^n_k summing to one such that, for all t_i:*

$$|u_{t_i}[f^n(t_i, \mu_0)] - \sum_{k=1,\dots,K} \pi^n_k \cdot u_{t_i}[\phi^n(\sigma^n_{\mu_k}(t_i), \sigma^n_{\mu_k}(\mu_k))]| < \epsilon.$$

Note that, unlike Theorem 1 in the text, the SP-L mechanism we construct in this case is

$$F^n(t) = \Phi^n(\sigma^n_{\text{emp}(t)}(t))$$

The proof of this Theorem is largely analogous to that of the Theorem for limit Nash equilibria. For conciseness, we will discuss the points where the proofs diverge, and how to adjust the proof of Theorem 1, instead of giving a complete proof. The proof is also based on an approximation Lemma.

The statement of the Lemma differs slightly.

Lemma 2. *Fix a prior μ_0 and $\epsilon > 0$. Let $\mathcal{N}$ be a neighborhood as in Definition 10. Let μ_k be priors $\mu_k \in \mathcal{A}_k$ for each $k = 1, \ldots, K$, with $|\mu_k - \mu_0| < \epsilon$. Then there exists n_0, and positive weights π^n_k with $\sum_{1 \leq k \leq K} \pi^n_k = 1$, such that for all t_i*

$$|f^n(t_i, \mu_0) - \sum_{k=1}^{K} \pi^n_k \cdot z_k(t_i)| < 5\epsilon,$$

where

$$z_k(t_i) = \phi^\infty(\sigma^\infty_{\mu_k}(t_i), \sigma^\infty_{\mu_k}(\mu_k)).$$

The proof of the alternate Lemma is largely similar to the proof of Lemma (1). In fact, the steps are basically the same, but replacing σ by σ^n or σ^∞ as appropriate. The only step of the proof that differs significantly is deriving the analogue of Inequality (A.5). Mutatis mutandis, this inequality would be showing that we may take n_3 large enough such that

$$|\phi^n(\sigma^n_{\mu_k}(t_i), \sigma^n_{\mu_k}(\mu_k)) - \phi^\infty(\sigma^\infty_{\mu_k}(t_i), \sigma^\infty_{\mu_k}(\mu_k))| < \epsilon. \tag{B.6}$$

This is still true. However, it does not follow from the definition of the limit, as in the proof of Lemma (1). Instead, it is a consequence of Condition 4 in the definition of a quasicontinuous sequence of families of equilibria. The rest of the proof follows straightforwardly, with the modifications we described.

References

Abdulkadiroğlu, A. and T. Sönmez, "School choice: A mechanism design approach," *The American Economic Review*, 2003, *93* (3), 729–747.

_ , P. Pathak, A.E. Roth, and T. Sönmez, "Changing the Boston Mechanism: Strategyproofness as Equal Access," *Unpublished manuscript, Harvard University*, 2006.

_ , P.A. Pathak, and A.E. Roth, "Strategy-proofness versus Efficiency in Matching with Indifferences: Redesigning the NYC High School Match," *The American Economic Review*, 2009, *99* (5), 1954–1978.

_ , Y.K. Che, and Y. Yasuda, "Resolving conflicting preferences in school choice: The Boston mechanism reconsidered," *The American Economic Review*, 2011, *101* (1), 399–410.

Al-Najjar, N.I. and R. Smorodinsky, "The efficiency of competitive mechanisms under private information," *Journal of Economic Theory*, 2007, *137* (1), 383–403.

Aumann, R.J., "Existence of competitive equilibria in markets with a continuum of traders," *Econometrica*, 1964, pp. 1–17.

Ausubel, L.M. and P. Cramton, "Demand Reduction and Inefficiency in Multi-Unit Auctions," 2002.

Azevedo, E.M. and J.D. Leshno, "The college admissions problem with a continuum of students," in "Proceedings of the 12th ACM conference on Electronic commerce" ACM 2011, pp. 333–334.

Azrieli, Y. and E. Shmaya, "Lipschitz Games," 2011.

Bergemann, D. and S. Morris, "Robust mechanism design," *Econometrica*, 2005, *73* (6), 1771–1813.

Bodoh-Creed, A., "Approximation of Large Games," Technical Report, Cornell University Mimeo 2010.

Bogomolnaia, A. and H. Moulin, "A new solution to the random assignment problem," *Journal of Economic Theory*, 2001, *100* (2), 295–328.

Budish, E., "The combinatorial assignment problem: Approximate competitive equilibrium from equal incomes," *Harvard University, Unpublished mimeo*, 2009.

_ **and E. Cantillon**, "The Multi-unit Assignment Problem: Theory and Evidence from Course Allocation at Harvard," *American Economic Review*, Forthcoming.

_ **, Y.K. Che, F. Kojima, and P. Milgrom**, "Implementing random assignments: A generalization of the birkhoff-von neumann theorem," 2011.

Carmona, G. and K. Podczeck, "Ex-post stability of Bayes-Nash equilibria of large games," *Games and Economic Behavior*, 2011.

Cramton, P. and B.E. Katzman, "Designed to Fail: The Medicare Auction for Durable Medical Equipment," Technical Report, Working Paper, University of Maryland 2010.

Cripps, M.W. and J.M. Swinkels, "Efficiency of large double auctions," *Econometrica*, 2006, *74* (1), 47–92.

Day, R. and P. Milgrom, "Core-selecting package auctions," *International Journal of Game Theory*, 2008, *36* (3), 393–407.

Deb, J. and E. Kalai, "Stability in Large Bayesian Games with Heterogeneous Players," 2011.

Debreu, G. and H. Scarf, "A Limit Theorem on the Core of an Economy," *International Economic Review*, 1963, *4* (3), 235–246.

Edelman, B. and M. Ostrovsky, "Strategic bidder behavior in sponsored search auctions," *Decision support systems*, 2007, *43* (1), 192–198.

_ **,** _ **, and M. Schwarz**, "Internet Advertising and the Generalized Second-Price Auction: Selling Billions of Dollars Worth of Keywords," *The American Economic Review*, 2007, *97* (1), 242–259.

Engelbrecht-Wiggans, R., J.A. List, and D.H. Reiley, "Demand Reduction in Multi-unit Auctions with Varying Numbers of Bidders: Theory and Field Experiments," *International Economic Review*, 2006, *47* (1), 203–231.

Featherstone, C. and M. Niederle, "Ex ante efficiency in school choice mechanisms: an experimental investigation," 2008.

Friedman, M., "Comment on" Collusion in the Auction Market for Treasury Bills"," *Journal of Political Economy*, 1964, *72*.

_ , "How to sell government securities," *Wall Street Journal*, 1991, p. A8.

Fudenberg, D. and J. Tirole, *Game theory*, MIT Press, 1991.

Gershkov, A., B. Moldovanu, and X. Shi, "Bayesian and Dominant Strategy Implementation Revisited," *Working Papers*, 2011.

Gibbard, A., "Manipulation of voting schemes: a general result," *Econometrica:*, 1973, pp. 587–601.

Green, J. and J.J. Laffont, "Characterization of satisfactory mechanisms for the revelation of preferences for public goods," *Econometrica:*, 1977, pp. 427–438.

Hurwicz, L., "On Informationally Decentralized Systems," *Decision and organization: A volume in honor of Jacob Marschak*, 1972, *12*, 297.

Hylland, A. and R. Zeckhauser, "The efficient allocation of individuals to positions," *The Journal of Political Economy*, 1979, pp. 293–314.

Immorlica, N. and M. Mahdian, "Marriage, honesty, and stability," in "Proceedings of the sixteenth annual ACM-SIAM symposium on Discrete algorithms" Society for Industrial and Applied Mathematics 2005, pp. 53–62.

Jackson, M.O. and A.M. Manelli, "Approximately Competitive Equilibria in Large Finite Economies," *Journal of Economic Theory*, 1997, *77* (2), 354–376.

Jegadeesh, N., "Treasury auction bids and the Salomon squeeze," *Journal of Finance*, 1993, pp. 1403–1419.

Kalai, E., "Large robust games," *Econometrica*, 2004, *72* (6), 1631–1665.

Kojima, F. and M. Manea, "Incentives in the probabilistic serial mechanism," *Journal of Economic Theory*, 2010, *145* (1), 106–123.

_ **and P.A. Pathak**, "Incentives and stability in large two-sided matching markets," *The American Economic Review*, 2009, *99* (3), 608–627.

Kovalenkov, A., "Simple Strategy-Proof Approximately Walrasian Mechanisms," *Journal of Economic Theory*, 2002, *103* (2), 475–487.

Krishna, A. and M.U. Ünver, "Improving the Efficiency of Course Bidding at Business Schools: Field and Laboratory Studies," *Marketing Science*, 2008, *27* (2), 262–282.

Lee, S.M., "Incentive Compatibility of Large Centralized Matching Markets," 2011.

Manelli, A.M. and D.R. Vincent, "Bayesian and Dominant-Strategy Implementation in the Independent Private-Values Model," *Econometrica*, 2010, *78* (6), 1905–1938.

Merlob, B., K. Peters, C.R. Plott, A. Pradhana, and Y. Zhang, "An Evaluation of the Proposed Procurement Auction for the Purchase of Medicare Equipment: Experimental Tests of the Auction Architecture," Technical Report, Working Paper, California Institute of Technology, November 2010.

Milgrom, P., "Assignment messages and exchanges," *American Economic Journal: Microeconomics*, 2009, *1* (2), 95–113.

_ , "Critical Issues in the Practice of Market Design," *Economic Inquiry*, 2011.

Miralles, A., "School choice: The case for the Boston mechanism," *Auctions, Market Mechanisms and Their Applications*, 2009, pp. 58–60.

Moulin, H., *Cooperative Microeconomics: A Game-Theoretic Introduction*, Princeton, NJ: Princeton University Press, 1995.

Myerson, R.B., "Incentive compatibility and the bargaining problem," *Econometrica: journal of the Econometric Society*, 1979, pp. 61–73.

Papai, S., "Strategyproof and Nonbossy Multiple Assignments," *Journal of Public Economic Theory*, 2001, *3* (3), 257–271.

Parkes, D.C., J. Kalagnanam, and M. Eso, "Achieving budget-balance with Vickrey-based payment schemes in exchanges," in "International Joint Conference on Artificial Intelligence," Vol. 17 Citeseer 2001, pp. 1161–1168.

Pathak, P.A. and T. Sönmez, "Leveling the playing field: Sincere and sophisticated players in the Boston mechanism," *The American Economic Review*, 2008, *98* (4), 1636–1652.

_ **and** _ , "School Admissions Reform in Chicago and England: Comparing Mechanisms by Their Vulnerability to Manipulation," Technical Report, National Bureau of Economic Research 2011.

Pesendorfer, W. and J.M. Swinkels, "Efficiency and information aggregation in auctions," *American Economic Review*, 2000, pp. 499–525.

Reny, P.J. and M. Perry, "Toward a strategic foundation for rational expectations equilibrium," *Econometrica*, 2006, *74* (5), 1231–1269.

Roberts, D.J. and A. Postlewaite, "The incentives for price-taking behavior in large exchange economies," *Econometrica:*, 1976, pp. 115–127.

Roth, A.E., "The economics of matching: Stability and incentives," *Mathematics of Operations Research*, 1982, pp. 617–628.

_ , "The economist as engineer: Game theory, experimentation, and computation as tools for design economics," *Econometrica*, 2002, *70* (4), 1341–1378.

_ , "What Have We Learned from Market Design?*," *The Economic Journal*, 2008, *118* (527), 285–310.

_ **and E. Peranson**, "The Redesign of the Matching Market for American Physicians: Some Engineering Aspects of Economic Design," *American Economic Review*, 1999, *89* (4), 748–780.

Roughgarden, T., "Computing equilibria: A computational complexity perspective," *Economic Theory*, 2010, *42* (1), 193–236.

Rustichini, A., M.A. Satterthwaite, and S.R. Williams, "Convergence to efficiency in a simple market with incomplete information," *Econometrica:*, 1994, pp. 1041–1063.

Sandholm, T. and C. Boutilier, "Preference Elicitation in Combinatorial Auctions," in Shoham Cramton and eds. Steinberg, eds., *Combinatorial Auctions*, MIT Press, 2006, chapter 10.

Satterthwaite, M.A., "Strategy-proofness and Arrow's conditions: Existence and correspondence theorems for voting procedures and social welfare functions," *Journal of Economic Theory*, 1975, *10* (2), 187–217.

Segal, I., "Optimal pricing mechanisms with unknown demand," *The American economic review*, 2003, *93* (3), 509–529.

Sönmez, T. and U. Ünver, "Course Bidding at Business Schools," *International Economic Review*, 2010.

Swinkels, J.M., "Efficiency of large private value auctions," *Econometrica*, 2001, *69* (1), 37–68.

Thomson, W., "Fair Allocation Rules," in "Handbook of Social Choice and Welfare," Vol. 2, Elsevier, 2011, chapter 21, pp. 393–506.

Wilson, R., "Game-Theoretic Analyses of Trading Processes„" in "Advances in Economic Theory: Fifth World Congress," Cambridge, UK: Cambridge University Press, 1987, chapter 2, pp. 33–70.