Strategy-proofness in the Large*

Eduardo M. Azevedo[†] and Eric Budish[‡]

August 25, 2015

Abstract

We propose a criterion of approximate incentive compatibility, strategy-proofness in the large (SP-L), and argue that it is a useful second-best to exact strategy-proofness (SP) for market design. Conceptually, SP-L requires that an agent who regards a mechanism's "prices" as exogenous to her report – be they traditional prices as in an auction mechanism, or price-like statistics in an assignment or matching mechanism – has a dominant strategy to report truthfully. Mathematically, SP-L weakens SP both by considering incentives in a large-market limit rather than finite economies and by considering incentives from an interim perspective rather than ex-post.

*First public version: October 2011. For helpful discussions we are grateful to Nabil Al-Najjar, Susan Athey, Aaron Bodoh-Creed, Gabriel Carroll, Sylvain Chassang, Jeff Ely, Alex Frankel, Drew Fudenberg, Matt Gentzkow, Jason Hartline, John Hatfield, Richard Holden, Ehud Kalai, Emir Kamenica, Navin Kartik, Fuhito Kojima, Scott Kominers, Jacob Leshno, Jon Levin, Paul Milgrom, Stephen Morris, Roger Myerson, David Parkes, Parag Pathak, Nicola Persico, Andy Postlewaite, Canice Prendergast, Mark Satterthwaite, Ilya Segal, Eran Shmaya, Lars Stole, Rakesh Vohra, Glen Weyl, Mike Whinston, and especially Al Roth. We thank Victor Zhang for excellent research assistance. We are grateful to seminar audiences at Ohio State, the 2011 MFI Conference on Matching and Price Theory, UCLA, Chicago, AMMA 2011, Boston College, the 2011 NBER Conference on Market Design, Duke / UNC, Michigan, Carnegie Mellon / Pittsburgh, Montreal, Berkeley, Northwestern, Rochester, Frontiers of Market Design 2012, ACM EC 2012, Princeton, Maryland, and UPenn.

[†]Wharton, eazevedo@wharton.upenn.edu.

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

1 Introduction

Strategy-proofness (SP), that playing the game truthfully is a dominant strategy, is perhaps the central notion of incentive compatibility in market design. SP is frequently imposed as a design requirement in theoretical analyses, across a broad range of assignment, auction, and matching problems. And, SP has played a central role in several design reforms in practice, including the redesign of school choice mechanisms in several cities, the redesign of the market that matches medical school graduates to residency positions, and efforts to create mechanisms for pairwise kidney exchange (See especially Roth (2008) and Pathak and Sönmez (2008, 2013)). There are several reasons why SP is considered so attractive. First, SP mechanisms are robust: since reporting truthfully is a dominant strategy, equilibrium does not depend on market participants' beliefs about other participants' preferences or information. Second, SP mechanisms are strategically simple: market participants do not have to invest time and effort collecting information about others' preferences or about what equilibrium will be played. Third, with this simplicity comes a measure of fairness: a participant who lacks the information or sophistication to game the mechanism is not disadvantaged relative to sophisticated participants. Fourth, SP mechanisms generate information about true preferences that may be useful to policy makers.[1]

However, SP is restrictive. In a variety of market design contexts, including matching, school choice, course allocation, and combinatorial auctions, impossibility theorems show that SP places meaningful limitations on what other attractive properties a mechanism can hope to satisfy.[2] And, SP is an extremely strong requirement. If there is a single configuration of participants' preferences in which a single participant has a strategic misreport that raises his utility by epsilon, a mechanism is not SP. A natural idea is to look for incentives criteria that are less demanding and less restrictive than SP, while still maintaining some of the

[1] See Wilson (1987) and Bergemann and Morris (2005) on robustness, Fudenberg and Tirole (1991), p. 270 and Roth (2008) on strategic simplicity, Friedman (1991), Pathak and Sönmez (2008) and Abdulkadiroğlu et al. (2006) on fairness and Roth (2008) and Budish et al. (2015) on the advantage of generating preference data.

[2] In matching problems such as the National Resident Matching Program, SP mechanisms are not stable (Roth, 1982). In multi-unit assignment problems such as course allocation, the only SP and ex-post efficient mechanisms are dictatorships (Papai, 2001; Ehlers and Klaus, 2003; Hatfield, 2009), which perform poorly on measures of fairness and ex-ante welfare (Budish and Cantillon, 2012). In school choice problems, which can be interpreted as a hybrid of an assignment and a matching problem (Abdulkadiroğlu and Sönmez, 2003), there is no mechanism that is both SP and ex-post efficient (Abdulkadiroğlu et al., 2009). In combinatorial auction problems such as the FCC spectrum auctions (Milgrom, 2004; Cramton et al., 2006), the only SP and efficient mechanism is Vickrey-Clarke-Groves (Green and Laffont, 1977; Holmstrom, 1979), which has a variety of important drawbacks (Ausubel and Milgrom, 2006). Perhaps the earliest such negative result for SP mechanisms is Hurwicz (1972), which shows that SP is incompatible with implementing a Walrasian equilibrium in an exchange economy.

advantages of SP design.

This paper proposes a criterion of approximate strategy-proofness called *strategy-proof in the large* (SP-L). SP-L weakens SP in two ways. First, whereas SP requires that truthful reporting is optimal in any size economy, SP-L requires that truthful reporting is optimal only in the limit as the market grows large. Second, whereas SP requires that truthful reporting is optimal against any realization of opponent reports, SP-L requires that truthful reporting is optimal only against any probability distribution of reports. That is, SP-L examines incentives from the interim perspective rather than ex-post. Because of this interim perspective, SP-L is weaker than the traditional notion of approximate strategy-proofness; this weakening is important both conceptually and for our results. At the same time, SP-L is importantly stronger than approximate Bayes-Nash incentive compatibility, because SP-L requires that truthful reporting is best against *any* probability distribution of opponent reports, not just the single probability distribution associated with Bayes-Nash equilibrium play. This strengthening is important because it allows SP-L to approximate, in large markets, the attractive properties such as robustness and strategic simplicity which are the reason why market designers like SP better than Bayes-Nash in the first place.

This combination of the large market limit and the interim perspective is powerful for the following reason: it causes each participant to regard the societal distribution of play as exogenous to his own report (more precisely, the distribution of the societal distribution of play). As will become clear, regarding the societal distribution of play as exogenous to one's own play is a generalization of the idea of regarding prices as exogenous, i.e., of "price taking". In some settings, such as multi-unit auctions or Walrasian exchange, the two concepts are equivalent. In other settings, such as school choice or two-sided matching, regarding the societal distribution of play as exogenous is equivalent to regarding certain price-like summary statistics of the mechanism as exogenous.

SP-L thus draws a distinction between two ways a mechanism can fail to be SP. If a mechanism is manipulable in finite markets, but is no longer manipulable by participants who regard the societal distribution of play as exogenous, the mechanism is SP-L. If a mechanism not only is manipulable in finite markets but remains so even when the societal distribution of play is exogenous – even a price taker, or a taker of price-like statistics, wishes to misreport – then the mechanism is not only not SP, but is not SP-L.

After we present and discuss the formal definition of SP-L, the next part of the paper provides a classification of existing non-SP mechanisms into those that are SP-L and those that are not SP-L. The classification, displayed in Table 1, organizes both the prior theory literature on which non-SP mechanisms have good incentives properties in large markets and the empirical record on when non-SP matters in real-world large markets. In the SP-L column are numerous mechanisms that, while not SP, have been shown theoretically to have approximate incentives for truth telling in large markets. Examples include the Walrasian mechanism (Roberts and Postlewaite, 1976; Jackson and Manelli, 1997), double auctions (Rustichini et al., 1994; Cripps and Swinkels, 2006), multi-unit uniform-price auctions (Swinkels, 2001), the Gale-Shapley deferred acceptance algorithm (Immorlica and Mahdian, 2005; Kojima and Pathak, 2009), and probabilistic serial (Kojima and Manea, 2010). This literature has used a wide variety of definitions of approximate incentive compatibility, as well as a wide variety of analysis techniques. We use a single definition and a single analysis technique (Theorem 1) and find that all of these mechanisms are SP-L.[3] On the other hand, in the non-SP-L column are numerous mechanisms for which there is explicit empirical evidence that real-world market participants strategically misreport their preferences, to the detriment of design objectives such as efficiency or fairness. Examples include multi-unit pay-as-bid auctions (Friedman, 1960, 1991), the Boston mechanism for school choice (Abdulkadiroğlu et al., 2006, 2009), the bidding points auction for course allocation (Sönmez and Ünver, 2010; Budish, 2011), the draft mechanism for course allocation (Budish and Cantillon, 2012), and the priority-match mechanism for two-sided matching (Roth, 2002). This literature has frequently emphasized that the mechanism in question is not SP; our point is that the mechanisms for which there is documentation of important incentives problems in practice not only are not SP, but are not even SP-L. Overall, the classification exercise suggests that the relevant distinction for practice, in markets with a large number of participants, is not "SP vs. not SP", but rather "SP-L vs. not SP-L".

The last part of the paper provides conditions under which, in large markets, SP-L is no more restrictive than Bayes-Nash incentive compatibility. The result is similar in spirit to the classic revelation principle (Myerson, 1979) (for more details on the relationship see Section 5.2). Suppose we are given a mechanism that is not SP-L but that has Bayes-Nash equilibria, for any common-knowledge i.i.d. prior beliefs about the distribution of types, and that the equilibria are continuous, in a sense made precise in the text. We then construct a mechanism that is SP-L and implements approximately the same outcomes as the original mechanism. The construction uses proxy agents. Participants report their types

[3]Note as well that the traditional ex-post notion of approximate strategy-proofness is too strong to obtain the classification. For instance, the uniform-price auction is SP-L but is not approximately strategy-proof in an ex-post sense; even in a large economy it is always possible to construct a knife-edge situation where a single player, by shading her demand, can have a large discontinuous influence on the market-clearing price.

Table 1: SP-L and non SP-L mechanisms for some canonical market design problems

Problem	Manipulable in the Large	SP-L
Multi-Unit Auctions	Pay as Bid	Uniform Price
Single-Unit Assignment	Boston Mechanism	Probabilistic Serial HZ Pseudomarket
Multi-Unit Assignment	Bidding Points Auction HBS Draft	Approximate CEEI Generalized HZ
Matching	Priority Match	Deferred Acceptance
Other		Walrasian Mechanism Double Auctions

Notes: See Supplementary Appendix D for a detailed description of each mechanism in the table as well as a proof of the mechanism's classification as either SP-L or manipulable in the large. Abbreviations: HBS = Harvard Business School; HZ = Hylland and Zeckhauser; CEEI = competitive equilibrium from equal incomes.

to our mechanism, which computes the empirical distribution of types, and then plays the original mechanism on each participant's behalf using the Bayes-Nash equilibrium strategy associated with the empirical distribution of reports. This converts a mechanism with Bayes-Nash equilibria, which depend on the prior, into an SP-L mechanism that implements the same outcome in the large-market limit.

We then apply the construction theorem to the Boston mechanism for school assignment.[4] This both illustrates the content of the theorem and offers a new perspective on an ongoing debate in the market design literature. The first papers on the Boston mechanism criticized it for not being SP and suggested that an SP mechanism be used instead (Abdulkadiroğlu and Sönmez, 2003; Abdulkadiroğlu et al., 2006). A second generation of papers then argued that the Boston mechanism, while not SP, has Bayes-Nash equilibria which generate greater welfare for students than do SP alternatives (Abdulkadiroğlu et al., 2011; Miralles, 2009; Featherstone and Niederle, 2011; Agarwal and Somaini, 2014; Hwang, 2014). However, these equilibria require students to have common knowledge of the preference distribution, coordinate on a specific equilibrium, and make sophisticated strategic calculations. Our the-

[4]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. See Section 6 for a formal description of the mechanism. See Pathak and Sönmez (2013) for a list of school-districts in the US and England that use or formerly used this mechanism. Note that the mechanism is no longer used in Boston itself, which switched to a strategy-proof mechanism in 2006.

orem says that all of this complexity and non-robustness is unnecessary in a large market. Specifically, we show that the Boston mechanism satisfies our continuity condition, and then apply our proxy construction to implement the same outcomes as these desirable Bayes-Nash equilibria of the Boston mechanism, but in a way that is SP-L.[5]

Overall, our analysis suggests that in large market settings SP-L approximates the advantages of SP design while being significantly less restrictive. Our hope is that market designers will view SP-L as a practical alternative to SP in settings with a meaningful number of participants and in which SP mechanisms perform poorly. A recent example of this approach to market design is the new course allocation mechanism implemented at The Wharton School at the University of Pennsylvania in Fall 2013. (Budish and Kessler, 2015; Budish et al., 2015).

The rest of the paper is organized as follows. Section 2 defines the environment. Section 3 defines and discusses SP-L. Section 4 presents the classification of non-SP mechanisms. Section 5 presents the result on constructing SP-L mechanisms from Bayes-Nash mechanisms. Section 6 applies the construction to the Boston mechanism. Section 7 discusses technical extensions and related literature. Section 8 concludes. Proofs and other supporting materials are in the appendix.

2 Environment

We work with an abstract mechanism design environment in which mechanisms assign outcomes to agents based on the set of agents' reports. There is a finite set of **(payoff) types** T and a finite set of **outcomes (or consumption bundles)** X_0**.** The outcome space describes the outcome possibilities for an individual agent. For example, in an auction the elements in X_0 specify both the objects an agent receives and the payment she makes. In school assignment, X_0 is the set of schools to which a student can be assigned. An agent's type determines her preferences over outcomes. For each $t_i \in T$ there is a von Neumann-Morgenstern expected **utility function** $u_{t_i} : X \to [0, 1]$, where $X = \Delta X_0$ denotes the set of lotteries over outcomes. Preferences are private values in the sense that an agent's utility depends exclusively on her type and the outcome she receives.

[5] To clarify, we do not think our results speak to which is preferable between the SP-L implementation of the Boston mechanism and the SP Gale-Shapley deferred acceptance mechanism. Recent empirical evidence (Agarwal and Somaini, 2014; Hwang, 2014) suggests that the welfare differences at stake may be small, i.e., the costs of exact SP may be small in this context. Our point is simply that *if* one believes that the analyses of Abdulkadiroğlu et al. (2011) and others imply that the Boston mechanism should be used in practice over Gale-Shapley, then one should consider whether our SP-L approach is better still.

We study mechanisms that are well defined for all possible market sizes, holding fixed X_0 and T. For each market size $n \in \mathbb{N}$, where n denotes the number of agents, an allocation is a vector of n outcomes, one for each agent, and there is a set $Y_n \subseteq (X_0)^n$ of **feasible allocations**. For instance, in an auction, the assumption that X_0 is fixed imposes that the number of potential types of objects is finite, and the sequence $(Y_n)_{\mathbb{N}}$ describes how the supply of each type of object changes as the market grows.

Definition 1. *Fix a set of outcomes X_0, a set of types T, and a sequence of feasibility constraints $(Y_n)_{\mathbb{N}}$. A* ***mechanism*** *$\{(\Phi^n)_{\mathbb{N}}, A\}$ consists of a finite set of actions A and a sequence of allocation functions*

$$\Phi^n : A^n \to \Delta((X_0)^n), \tag{2.1}$$

each of which satisfies feasibility: for any $n \in \mathbb{N}$ and $a \in A^n$, the support of $\Phi^n(a)$ is contained in the feasible set Y_n. A mechanism is ***direct*** *if $A = T$.*

We assume that mechanisms are **anonymous**, which requires that each agent's outcome depends only on her own action and the distribution of all actions. Formally, a mechanism is anonymous if the allocation function $\Phi^n(\cdot)$ is invariant to permutations for all $n \in \mathbb{N}$. Anonymity is a natural feature of many large-market settings. In Supplementary Appendix C we relax anonymity to **semi-anonymity** (Kalai, 2004). A mechanism is semi-anonymous if agents are divided into a finite set of groups, and an agent's outcome depends only on her own action, her group, and the distribution of actions within each group. Semi-anonymity accommodates applications in which there are asymmetries among classes of participants, such as double auctions in which there are distinct buyers and sellers, school choice problems in which students are grouped into different priority classes, and matching markets that are divided into two sides. Semi-anonymity also provides a convenient way to generalize results stated for i.i.d. distributions to more general distributions.

We adopt the following notation. Given a finite set S, the set of probability distributions over S is denoted ΔS, and the set of distributions with full support $\bar{\Delta} S$. Distributions over the set of types will typically be denoted as $\mu \in \Delta T$, and distributions over actions by $m \in \Delta A$. Throughout the analysis we will use the supremum norm on the sets ΔT, ΔA and X. Since the number of types, actions and outcomes is finite, all of these probability spaces are subsets of Euclidean space. Using this representation, we denote the distance between two outcomes $x, x' \in X$ as $\|x - x'\|$, and likewise for distributions over T and A. In particular, we use this topology in the definition of limit mechanisms below.

Given a vector of types $t \in T^n$, we use the notation $\text{emp}[t]$ to denote the empirical distribution of t on T. That is, for each type $\tau \in T$, $\text{emp}[t](\tau)$ is the fraction of coordinates of t that equal τ, and the vector $\text{emp}[t] = (\text{emp}[t](\tau))_{\tau \in T}$. Analogously, given a vector of actions $a \in A^n$, $\text{emp}[a]$ denotes the empirical distribution of a on A.

3 Strategy-proof in the Large

In this section we formally define strategy-proofness in the large (SP-L) and discuss its interpretation and its relationship to previous concepts.

3.1 Large-Market Limit

We begin by defining our notion of the large-market limit.

Given a mechanism $\{(\Phi^n)_{\mathbb{N}}, A\}$, define, for each n, the function $\phi^n : A \times \Delta A \to X$ according to

$$\phi^n(a_i, m) = \sum_{a_{-i} \in A^{n-1}} \Phi_i^n(a_i, a_{-i}) \cdot \Pr(a_{-i} | a_{-i} \sim iid(m)), \tag{3.1}$$

where $\Phi_i^n(a_i, a_{-i})$ denotes the marginal distribution of the i^{th} coordinate of $\Phi^n(a)$, i.e., the lottery over outcomes received by agent i when she plays a_i and the other $n-1$ agents play a_{-i}, and $\Pr(a_{-i}|a_{-i} \sim iid(m))$ denotes the probability that the action vector a_{-i} is realized given $n-1$ independent identically distributed (i.i.d.) draws from the action distribution $m \in \Delta A$. In words, $\phi^n(a_i, m)$ describes what an agent who plays a_i expects to receive, ex interim, if the other $n-1$ agents play i.i.d. according to action distribution m.

We use the interim allocation function ϕ^n to define the large-market limit.

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

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

In words, $\phi^\infty(a_i, m)$ equals the lottery that an agent who plays a_i receives, in the limit as the number of agents grows large, when the other agents play i.i.d. according to the probability distribution m.[6]

[6] The randomness in how we take the large-market limit is in contrast with early approaches to large-market analysis, such as Debreu and Scarf's (1963) replicator economy and Aumann's (1964) continuum economy. It is more closely related to the random economy method used in Immorlica and Mahdian's (2005) and Kojima and Pathak's (2009) studies of large matching markets.

It is easy to construct examples of mechanisms that do not have limits. For instance, if a mechanism is a uniform-price auction when n is even and is a pay-as-bid auction when n is odd, then the mechanism does not have a limit. However, we are not aware of a mechanism used in practice, or proposed for practical use, that does not have a limit. For the remainder of the paper we restrict attention to mechanisms that have limits.

Interpretation of the Limit and Relationship with Price Taking The randomness in how we take the large-market limit is economically important for the following reason: in our limit, the distribution of the empirical distribution of play is exogenous to any particular agent's own play. We state this claim formally in the Appendix as Lemma A.1. Intuitively, if a fair coin is tossed n times the distribution of the number of heads is stochastic, and the influence of the i^{th} coin toss on this distribution vanishes to zero as $n \to \infty$; whereas if the market grew large in a deterministic fashion one player's decision between heads or tails could be pivotal as to whether the number of heads is greater than or less than 50%.

We interpret treating the societal distribution of play as exogenous to one's own report as a generalized version of price taking. Suppose that a mechanism has prices that are a function of the empirical distribution of play. For example, in a uniform-price auction, price is determined based on where reported demand equals reported supply. In the limit, because the distribution of the empirical distribution of play is exogenous to each agent, the distribution of prices is exogenous to each agent. Now suppose that a mechanism does not have prices, but has price-like statistics that are functions of the empirical distribution of play and sufficient statistics for the outcomes received by agents who played each action. For example, in Bogomolnaia and Moulin's (2001) assignment mechanism, the empirical distribution of reports determines statistics called "run-out times", which describe at what time in their algorithm each object exhausts its capacity. As a second example, in Azevedo and Leshno (2011)'s matching model the empirical distribution of reports determines a set of statistics called "cutoffs" which describe the level of desirability necessary to achieve each possible match partner. In our large-market limit, each agent regards the distribution of these price-like statistics as exogenous to their own report.

3.2 Definition of SP-L

A mechanism is strategy-proof if it is optimal for each agent to report truthfully, in any size market, given any realization of opponent reports.

Definition 3. *The mechanism* $\{(\Phi^n)_{\mathbb{N}}, T\}$ *is* ***strategy-proof (SP)*** *if, for all* n*, all* $t_i, t'_i \in T$*, and all* $t_{-i} \in T^{n-1}$

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

We say that a mechanism is strategy-proof in the large if it is optimal for each agent to report truthfully, in the large-market limit, given any full support i.i.d. distribution of opponent reports.

Definition 4. *The mechanism* $\{(\Phi^n)_{\mathbb{N}}, T\}$ *is* ***strategy-proof in the large (SP-L)*** *if, for any* $m \in \bar{\Delta}T$ *and all* $t_i, t'_i \in T$

$$u_{t_i}[\phi^\infty(t_i, m)] \geq u_{t_i}[\phi^\infty(t'_i, m)]. \tag{3.2}$$

Equivalently, the mechanism is SP-L if, for any $m \in \bar{\Delta}T$ *and* $\epsilon > 0$ *there exists* n_0 *such that, for all* $n \geq n_0$ *and all* $t_i, t'_i \in T$

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

Otherwise, the mechanism is ***manipulable in the large****.*

SP-L weakens SP in two ways. First, while SP requires that truthful reporting is optimal in any size market, SP-L requires that truthful reporting is optimal only in the limit as the market grows large. In large finite markets truthful reporting is only optimal in an approximate sense. Second, SP evaluates what report is best based on the (ex-post) realization of reports, whereas SP-L evaluates based on the (ex-interim) probability distribution of reports. A mechanism can be SP-L even if it has the property that, given $\epsilon > 0$, in any size market n one can find a type t_i and realization of opponent reports t_{-i} for which t_i has a misreport worth more than ϵ. What SP-L rules out is that there is a probability distribution of opponent reports with this property. Implicitly, SP-L takes a view on what information participants have in a large market when they decide how to play – they may have a (possibly incorrect) sense of the distribution of opponent preferences, but they do not know the exact realization of opponent preferences.

These two weakenings place SP-L between two commonly used notions of incentive compatibility. SP-L is weaker than the standard notion of asymptotic strategy-proofness, which requires that reporting truthfully is approximately optimal, in a large enough market, for any realization of opponent reports.[7] This distinction is important for the classification below;

[7]For example, Liu and Pycia (2011) define a mechanism as asymptotically strategy-proof if, given $\epsilon > 0$,

nearly all of the mechanisms that are classified as SP-L would fail this stronger criterion (e.g., uniform-price auctions, deferred acceptance), with the lone exception being the probabilistic serial mechanism. At the same time, SP-L is stronger than approximate Bayes-Nash incentive compatibility, which requires that truthful reporting is approximately optimal against the true probability distribution of opponent reports, which itself is assumed to be common knowledge. In contrast, SP-L requires truthful reporting to be approximately optimal for any probability distribution of opponent reports. This distinction is what allows SP-L mechanisms to maintain, at least approximately, some of the attractive features of SP design such as robustness, strategic simplicity, and fairness to unsophisticated agents.

Finally, the definition of the limit gives a useful way to think about SP-L as a generalization of price-taking. In the large-market limit the aggregate distribution of actions depends only on the distribution of one's opponents' actions, and not on one's own action. Thus, in the limit, agents take as given any statistic of the distribution of actions. In particular, in a mechanism that uses prices that are a function of the distribution of actions, agents take the distribution of prices as given. Thus, a mechanism is SP-L if reporting truthfully is optimal taking prices as given – or, more generally, taking the aggregate distribution of play as given. A mechanism is not SP-L if even an agent who takes prices as given – or, more generally, takes the aggregate distribution of play as given – wishes to misreport.

4 Classification of Non-SP Mechanisms

This section classifies a number of non-SP mechanisms into SP-L and manipulable in the large (Table 1 in the Introduction), and discusses how the classification organizes the evidence on manipulability in large markets. Specifically, all of the known mechanisms for which there is a detailed theoretical case that the mechanism has approximate incentives for truthtelling in large markets are SP-L (Section 4.2), and all of the known mechanisms for which there is empirical evidence that non-strategy-proofness causes serious problems even in large markets are manipulable in the large (Section 4.3).

Before proceeding, we make three brief observations regarding the classification. First, both the SP-L and the manipulable in the large columns of Table 1 include mechanisms that explicitly use prices (e.g., multi-unit auctions), as well as mechanisms that do not use prices

---

there exists n_0 such that for all $n \geq n_0$, types t_i, t_i', and a vector of $n-1$ types t_{-i},

$$u_{t_i}[\Phi_i^n(t_i, t_{-i})] \geq u_{t_i}[\Phi_i^n(t_i', t_{-i})] - \epsilon.$$

A similar definition is in Hatfield et al. (2015).

(e.g., matching mechanisms). For the mechanisms that do use prices, the SP-L ones are exactly those where an agent who takes prices as given wishes to report truthfully, such as the uniform-price auction. Second, the table is consistent with both Milton Friedman's (1960; 1991) argument in favor of uniform-price auctions over pay-as-bid auctions, and Alvin Roth's (1990; 1991; 2002) argument in favor of deferred acceptance over priority-match algorithms. Notably, while both Friedman's criticism of pay-as-bid auctions and Roth's criticism of priority-match algorithms were made on incentives grounds, the mechanisms they suggested in their place are not SP but are SP-L. Third, with the exception of probabilistic serial, none of the SP-L mechanisms satisfy a stronger, ex-post, notion of approximate strategy-proofness. That is, the classification would not conform to the existing evidence, nor to Friedman's and Roth's arguments, without the ex-interim perspective in the definition of SP-L.

4.1 Obtaining the Classification

To show that a mechanism is not SP-L it suffices to identify an example of a distribution of play under which agents may gain by misreporting, even in the limit. For SP-L mechanisms, this section gives two easy-to-check sufficient conditions for a mechanism to be SP-L, which directly yield the classification for all of the SP-L mechanisms in Table 1. Formal definitions of each mechanism and detailed derivations are in Supplementary Appendix D.[8]

The first sufficient condition is envy-freeness, a fairness criterion which requires that no player i prefers the assignment of another player j, for any realization of the reported types t.

Definition 5. *A direct mechanism* $\{(\Phi^n)_{\mathbb{N}}, T\}$ *is* ***envy-free (EF)*** *if, for all* i, j, n, t*:*

$$u_{t_i}[\Phi^n_i(t)] \geq u_{t_i}[\Phi^n_j(t)].$$

Theorem 1 below shows that EF implies SP-L. The intuition for the proof is as follows. In anonymous mechanisms, the gain to player i from misreporting as player j can be decomposed as the sum of the gain from receiving j's bundle, holding fixed the aggregate distribution of types, plus the gain from affecting the aggregate distribution of types (expression (A.2) in Appendix A). Envy-freeness directly implies that the first component in this decomposition

[8]Two of these mechanisms do not fit the framework used in the body of the paper. Deferred acceptance is a semi-anonymous mechanism, and the Walrasian mechanism has an infinite set of bundles. For details of how we accommodate these generalizations, see Supplementary Appendix D.

is non-positive. Lemma A.1 then implies that the second component becomes negligible in large markets. More precisely, the effect of misreporting on the distribution of the empirical distribution of reports vanishes at a rate of essentially $\sqrt{n}$, which yields both that EF implies SP-L and the convergence rate for EF mechanisms as stated in Theorem 1.

Most of the mechanisms in the SP-L column of Table 1 are EF, with the only exceptions being approximate CEEI and deferred acceptance.[9] Fortunately, these mechanisms satisfy a weakening of EF that we show is also sufficient. Specifically, each of these mechanisms involves a certain form of tie-breaking lottery, and after this lottery is realized no agent envies another agent with a lower lottery number. Formally,[10]

Definition 6. *A direct mechanism $\{(\Phi^n)_{\mathbb{N}}, T\}$ is* ***envy-free but for tie breaking (EF-TB)*** *if for each n there exists a function $x^n : (T \times [0,1])^N \to \Delta(X_0^n)$, symmetric over its coordinates, such that*

$$\Phi^n(t) = \int_{l \in [0,1]^n} x^n(t,l)dl$$

and, for all i, j, n, t, and l, if $l_i \geq l_j$ then

$$u_{t_i}[x_i^n(t,l)] \geq u_{t_i}[x_j^n(t,l)].$$

The following theorem shows that either condition guarantees that a mechanism is SP-L.

Theorem 1. *If a mechanism is EF-TB (and in particular if it is EF), then it is SP-L. The maximum possible gain from misreporting converges to 0 at a rate of $n^{-\frac{1}{2}+\epsilon}$ for EF mechanisms, and $n^{-\frac{1}{4}+\epsilon}$ for EF-TB mechanisms. Formally, if a mechanism is EF (EF-TB), then given $\mu \in \bar{\Delta}T$ and $\epsilon > 0$ there exists $C > 0$ such that, for all t_i, t_i' and n, the gain from deviating,*

$$u_{t_i}[\phi_i^n(t_i', \mu)] - u_{t_i}[\phi_i^n(t_i, \mu)],$$

is bounded above by

$$C \cdot n^{-\frac{1}{2}+\epsilon} \quad (C \cdot n^{-\frac{1}{4}+\epsilon}).$$

The theorem shows that either condition can be used to classify new or existing mechanisms as SP-L. It also gives reasonable rates of convergence for the maximum possible gain from manipulating a mechanism.

[9] Both approximate CEEI and deferred acceptance include as a special case the random serial dictatorship mechanism, which Bogomolnaia and Moulin (2001) show is not envy-free.

[10] This definition is for anonymous mechanisms. The definition for semi-anonymous mechanisms, which is needed for deferred acceptance, is contained in Supplementary Appendix C.

The proof of the theorem for the EF-TB case builds upon the argument for the EF case, by showing that EF-TB mechanisms have small amounts of envy before lotteries are drawn (Lemma A.2). This is accomplished with three basic ideas. First, how much player i envies player j prior to the lottery draw equals the average envy by all type t_i players towards type t_j players, as a consequence of anonymity. Second, it is possible to bound this average envy, after a given lottery draw l, by how evenly distributed the lottery numbers in the vector l are. Intuitively, if players of types t_i and t_j receive evenly distributed lottery numbers, average envy has to be small. The final step is an application of a probabilistic bound known as the Dvoretzky–Kiefer–Wolfowitz inequality, which guarantees that lottery numbers are typically very evenly distributed.

4.2 Relationship to the Theoretical Literature on Large Markets

The SP-L column of Table 1 organizes a large literature demonstrating the approximate incentive compatibility of specific mechanisms in large markets. We obtain results for Walrasian mechanisms as in Roberts and Postlewaite (1976) and Jackson and Manelli (1997), double auctions as in Rustichini et al. (1994) and Cripps and Swinkels (2006), uniform-price auctions as in Swinkels (2001), deferred acceptance mechanisms as in Immorlica and Mahdian (2005) and Kojima and Pathak (2009), the probabilistic serial mechanism as in Kojima and Manea (2010), approximate CEEI as in Budish (2011), and the Hylland and Zeckhauser (1979) and generalized Hylland-Zeckhauser (Budish et al., 2013) pseudomarket mechanisms, whose large-market incentive properties had not previously been studied.

The single concept of SP-L and Theorem 1 classifies all of these mechanisms. In contrast, the prior literature has employed different notions of approximate incentive compatibility and different analysis techniques, tailored for each mechanism.[11] Of course, analyses that are tailored to specific mechanisms can yield a more nuanced understanding of the exact forces pushing players away from truthful behavior in finite markets, as in the first-order

[11] This note elaborates on the different concepts used in the literature. Roberts and Postlewaite (1976) ask that truthful reporting is ex-post approximately optimal for all opponent reports where equilibrium prices vary continuously with reports. Rustichini et al. (1994) study the exact Bayes-Nash equilibria of double auctions in large markets, and bound the rate at which strategic misreporting vanishes with market size. Swinkels (2001) studies both exact Bayes-Nash equilibria and ϵ-Bayes-Nash equilibria of the uniform-price and pay-as-bid auctions. Kojima and Pathak (2009) study ϵ-Nash equilibria of the doctor-proposing deferred acceptance algorithm assuming complete information about preferences on the hospital side of the market and incomplete information about preferences on the doctor side of the market. In an appendix they also consider ϵ-Bayes-Nash equilibria, in which there is incomplete information about preferences on both sides of the market. Kojima and Manea (2010) show that probabilistic serial satisfies exact SP, without any modification, in a large enough finite market. Budish (2011) shows that approximate CEEI satisfies exact SP in a continuum economy.

condition analysis of Rustichini et al. (1994) or the rejection chain analysis of Kojima and Pathak (2009).

4.3 Relationship to Empirical Literature on Manipulability

For each of the manipulable in the large mechanisms in Table 1, there is explicit empirical evidence that participants strategically misreport their preferences in practice. Furthermore, misreporting harms design objectives such as efficiency or fairness. In this section we briefly review this evidence.

Consider first multi-unit auctions for government securities. Empirical analyses have found considerable bid shading in discriminatory auctions (Hortaçsu and McAdams, 2010), but negligible bid shading in uniform-price auctions, even with as few as 13 bidders (Kastl, 2011; Hortaçsu et al. (2015)). Friedman (1991) argued that the need to play strategically in pay-as-bid auctions reduces entry of less sophisticated bidders, giving dealers a sheltered market that facilitates collusion. In uniform-price auctions, by contrast, "You do not have to be a specialist" to participate, since all bidders pay the market-clearing price. Consistent with Friedman's view, Jegadeesh (1993) shows that pay-as-bid auctions depressed revenues to the US Treasury during the Salomon Squeeze in 1991, and Malvey and Archibald (1998) find that the US Treasury's adoption of uniform-price auctions in the mid-1990s broadened participation. Cross-country evidence is also consistent with Friedman's argument, as Brenner et al. (2009) find a positive relationship between a country's using uniform-price auctions and indices of ease of doing business and economic freedom, whereas pay-as-bid auctions are positively related with indices of corruption and of bank-sector concentration.

Next, consider the Boston mechanism for school choice. Abdulkadiroğlu et al. (2006) find evidence of a mix of both sophisticated strategic misreporting and unsophisticated naive truthtelling; see also recent empirical work by Agarwal and Somaini (2014) and Hwang (2014). Sophisticated parents strategically misreport their preferences by ranking a relatively unpopular school high on their submitted preference list. Unsophisticated parents, on the other hand, frequently play a dominated strategy in which they waste the highest positions on their rank-ordered list on popular schools that are unattainable for them. In extreme cases, participants who play a dominated strategy end up not receiving any of the schools they ask for.

Next, consider the mechanisms used in practice for the multi-unit assignment problem of course allocation. In the bidding points auction, Krishna and Ünver (2008) use both field and laboratory evidence to show that students strategically misreport their preferences, and that this harms welfare. Budish (2011) provides additional evidence that some students get very poor outcomes under this mechanism; in particular students sometimes get zero of the courses they bid for. In the Harvard Business School draft mechanism, Budish and Cantillon (2012) use data consisting of students' stated preferences and their underlying true preferences to show that students strategically misreport their preferences. They show that misreporting harms welfare relative both to a counterfactual in which students report truthfully, and relative to a counterfactual in which students misreport, but optimally. They also provide direct evidence that some students fail to play best responses, which supports the view that Bayes-Nash equilibria are less robust in practice than dominant-strategy equilibria.

For labor market clearinghouses, Roth (1990, 1991, 2002) surveys a wide variety of evidence that shows that variations on priority matching mechanisms perform poorly in practice, while variations on Gale and Shapley's deferred acceptance algorithm perform well. Roth emphasizes that the former are unstable under truthful play whereas the latter are stable under truthful play. By contrast, we emphasize that the former are not SP-L whereas the latter are SP-L.

5 SP-L is Approximately Costless in Large Markets Relative to Bayes-Nash

In this section we will show that, in large markets, SP-L is in a well-defined sense approximately costless to impose relative to Bayes-Nash incentive compatibility. The exception is that there can be a large cost if the Bayes-Nash mechanism is very sensitive to agents' beliefs, but this itself is likely to be undesirable in practical market design settings.

5.1 Preliminaries

It will be useful to extend the function Φ^n linearly to distributions over vectors of actions. Given a distribution $\bar{m} \in \Delta(A^n)$ over vectors of actions, let

$$\Phi^n(\bar{m}) = \sum_{a \in A^n} \bar{m}(a) \cdot \Phi^n(a). \tag{5.1}$$

Likewise, given an action a_i and a distribution $\bar{m} \in \Delta(A^{n-1})$ over $n-1$ actions, let

$$\Phi_i^n(a_i, \bar{m}) = \sum_{a_{-i} \in A^{n-1}} \bar{m}(a_{-i}) \cdot \Phi_i^n(a_i, a_{-i})$$

and given distributions $\hat{m}, m \in \Delta A$ let

$$\phi^{\infty}(\hat{m}, m) = \sum_{a_i \in A} \hat{m}(a_i) \cdot \phi^{\infty}(a_i, m)$$

Given a mechanism $\{(\Phi^n)_{\mathbb{N}}, A\}$, a **strategy** σ is defined as a map from T to ΔA. Given a strategy σ and a vector of types t, let $\sigma(t) \in \Delta(A^n)$ denote the associated distribution over vectors of actions. Given a strategy σ and a probability distribution over types $\mu \in \Delta T$, let $\sigma(\mu) \in \Delta A$ denote the distribution over actions induced by strategy σ when player types are drawn according to μ.

Definition 7. *Given a mechanism $\{(\Phi^n)_{\mathbb{N}}, A\}$ with limit $\phi^{\infty}(\cdot,\cdot)$, and a probability distribution over types $\mu \in \Delta T$, the strategy $\sigma^*_{\mu} : T \to \Delta A$ is a* ***limit Bayes-Nash equilibrium at prior μ*** *if, for all $t_i \in T$ and $a'_i \in A$:*

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

Typically, a mechanism's Bayes-Nash equilibria vary with the prior. For instance, in a pay-as-bid auction how much bidders shade their bid in equilibrium varies with the distribution of bidders' values, and in the Boston mechanism how students misreport their preferences in equilibrium depends on the distribution of students' preferences. We define a family of limit equilibria as a set containing a limit equilibrium for each possible prior.[12]

Definition 8. *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 $\mu \in \Delta T$, the strategy $\sigma^*_{\mu}(\cdot)$ is a limit BNE at prior μ.*

Our continuity condition is defined on a family of limit equilibria.

Definition 9. *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 continuous at prior $\mu_0 \in \bar{\Delta} T$ if, given $\epsilon > 0$, there exists n_0 and a neighborhood $\mathcal{N}$ of μ_0 such that, for any $n \geq n_0$, $t_i \in T$, $t_{-i} \in T^{n-1}$ where $\text{emp}[t_i, t_{-i}] \in \mathcal{N}$, and $\mu, \mu' \in \mathcal{N}$, we have:*

$$\left\| \Phi^n_i(\sigma^*_{\mu}(t_i), \sigma^*_{\mu}(t_{-i})) - \Phi^n_i(\sigma^*_{\mu'}(t_i), \sigma^*_{\mu'}(t_{-i})) \right\| < \epsilon.$$

The family of equilibria is ***continuous*** *if it is continuous at every full support prior.*

[12] In an earlier version of this paper we showed that the analysis goes through essentially unchanged if we use a family of exact Bayes-Nash equilibria in large finite markets rather than a family of limit Bayes-Nash equilibria. Please see Appendix C.2 of Azevedo and Budish (2013).

In words, a family of equilibria is continuous if a small change in the prior μ has only a small effect on agent t_i's outcome in a large enough market. We show below in Section 6 that the Boston mechanism has a continuous family of equilibria (Proposition 1). In Section 7.1 and Supplementary Appendix B we also consider a weaker notion of continuity that allows for points of discontinuity so long as they are in a certain sense knife-edge. The theorem goes through under this condition as well but in a slightly weaker form.

5.2 Construction Theorem

We now establish that, given a mechanism with a continuous family of Bayes-Nash equilibria, there exists an SP-L mechanism that implements approximately the same outcomes as the original mechanism.

Theorem 2. *Given any mechanism $\{(\Phi^n)_{\mathbb{N}}, A\}$ with a continuous family of limit Bayes-Nash equilibria $(\sigma^*_\mu)_{\mu\in\Delta T}$, there exists a direct, SP-L mechanism $\{(F^n)_{\mathbb{N}}, T\}$ such that, in the large market limit, for any prior, truthful play of the direct mechanism produces the same outcomes as equilibrium play of the original mechanism. Formally, letting f^∞ denote the limit of the direct mechanism, for any full-support prior μ and any type t_i we have*

$$f^\infty(t_i, \mu) = \phi^\infty(\sigma^*_\mu(t_i), \sigma^*_\mu(\mu)).$$

Proof Sketch. The proof of Theorem 2 is by construction. We provide a detailed sketch as follows, with full details contained in Appendix A.

Suppose in a market of size n the agents report types $t = (t_1, \ldots, t_n)$. Our constructed mechanism calculates the empirical distribution of reports, $\text{emp}[t]$, and then plays the limit Bayes-Nash equilibrium of the original mechanism associated with this empirical distribution:

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

In words, F^n plays action $\sigma^*_{\text{emp}[t]}(t_i)$ for agent i who reports t_i, where $\text{emp}[t]$ is not the true distribution of agents' types μ_0 (which is not known to the mechanism) but rather the *empirical* distribution of agents' reported types. Intuitively, F^n acts as a proxy agent playing the original mechanism Φ^n on each agent's behalf, and uses a strategy that would be the limit Bayes-Nash equilibrium in a world in which the empirical distribution of agents' reports were in fact the true distribution of agents' preferences, and, additionally, this distribution was common knowledge.

To prove that this constructed mechanism produces outcomes under truthful play that coincide with equilibrium play of the original mechanism, suppose that the prior is μ_0 and that all agents report truthfully to the constructed mechanism. 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 thus receives $F_i^n(t_i, t_{-i}) = \Phi_i^n(\sigma^*_{\hat{\mu}}(t_i), \sigma^*_{\hat{\mu}}(t_{-i}))$. As the market grows large, the realized distribution of $\hat{\mu}$ converges almost surely to the true distribution μ_0, by the law of large numbers. Hence, by continuity, agent i's allocation is converging to

$$f^\infty(t_i, \mu_0) = \phi^\infty(\sigma_{\mu_0}(t_i), \sigma_{\mu_0}(\mu_0)).$$

As required, this is exactly what agent i would receive under the original mechanism, in the large-market limit, in the Bayes-Nash equilibrium corresponding to the true prior μ_0.

To prove that the constructed mechanism is SP-L, suppose that the agents other than i misreport their preferences, according to some distribution $m \neq \mu_0$. 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 $F_i^n(t_i, t'_{-i}) = \Phi_i^n(\sigma^*_{\hat{m}}(t_i), \sigma^*_{\hat{m}}(t'_{-i}))$. As the market grows large, the distribution of $\hat{m}$ will converge in probability to m, so, by continuity, agent i's allocation is converging to

$$f^\infty(t_i, m) = \phi^\infty(\sigma_m(t_i), \sigma_m(m)).$$

This is what agent i would receive under the original mechanism, in the large-market limit, *in the Bayes-Nash equilibrium corresponding to prior m.* Even though the other agents are systematically misreporting their preferences, it is optimal for agent i 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 she reports her preferences truthfully to the constructed mechanism.[13] Hence, in the limit, it is optimal for i to report truthfully for any distribution of opponent reports, i.e., the constructed mechanism is SP-L. □

Relationship to the Revelation Principle The construction is related to the traditional Bayes-Nash direct revelation mechanism construction (Myerson, 1979). In a traditional Bayes-Nash direct revelation mechanism, the mechanism designer and participants have a common knowledge prior about payoff types, say μ_0. The mechanism announces a Bayes-

[13]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.

Nash equilibrium strategy $\sigma^*_{\mu_0}(\cdot)$, and plays $\sigma^*_{\mu_0}(t_i)$ on behalf of an agent who reports t_i. Truthful reporting is a Bayes-Nash equilibrium.

In contrast, our constructed mechanism does not depend on a prior. Instead, the mechanism *infers* a prior from the empirical distribution of agents' play (cf. Segal (2003); Baliga and Vohra (2003)). If agents indeed play truthfully, this inference is correct in the limit. But if the agents misreport, so that the empirical $\hat{m}$ is very different from the prior μ_0, our mechanism 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 her preferences truthfully remains happy to have done so even if the other agents misreport, unlike in a traditional Bayes-Nash direct revelation mechanism, and our mechanism is SP-L rather than Bayes-Nash. Moreover, our mechanism is prior free and consistent with the Wilson doctrine, unlike a traditional Bayes-Nash direct revelation mechanism.

6 Application: The Boston Mechanism

Theorem 2 can be applied to provide a new perspective on an ongoing market design debate concerning the Boston mechanism for student assignment.

The earliest papers on the Boston mechanism, Abdulkadiroğlu and Sönmez (2003) and Abdulkadiroğlu et al. (2006), criticized the mechanism on the grounds that it is not strategy-proof, and proposed that the strategy-proof Gale-Shapley deferred acceptance algorithm be used instead.[14] These papers had a major policy impact as they led to the Gale-Shapley algorithm's eventual adoption for use in practice (cf. Roth (2008)).

A second generation of papers on the Boston mechanism, Abdulkadiroğlu et al. (2011); Miralles (2009); Featherstone and Niederle (2011), made a more positive case for the mechanism. They argued that while the Boston mechanism is not strategy-proof, it has Bayes-Nash equilibria that are attractive. In particular, it has Bayes-Nash equilibria that yield greater student welfare than do the dominant strategy equilibria of the Gale-Shapley procedure. Perhaps, these papers argue, the earlier papers were too quick to dismiss the Boston mechanism in favor of strategy-proof deferred acceptance.

However, these second-generation papers rely on students being able to reach the attractive Bayes-Nash equilibria. This raises several potential questions: is common knowledge a reasonable assumption? will students be able to calculate the desired equilibrium? will

[14] In two-sided matching, the Gale-Shapley algorithm is strategy proof for the proposing side of the market and SP-L for the non-proposing side of the market. In school choice only the student side of the market is strategic, with schools being non-strategic players whose preferences are determined by public policy.

unsophisticated students be badly harmed?

We will show, using Theorem 2, that in a large market it is possible to obtain the attractive welfare properties of the Bayes-Nash equilibria identified by these second-generation papers on the Boston mechanism, but without the robustness problems associated with Bayes-Nash mechanisms.

6.1 Definition of the Boston Mechanism

Denote the set of schools by $X_0 = S \cup \{\emptyset\}$. In a market of size n, there are $\lfloor q_s \cdot n \rfloor$ seats available in school s, where $q_s \in (0,1)$ denotes the proportion of the market that s can accommodate and $\lfloor \cdot \rfloor$ is the floor function. It is assumed that X_0 includes a null school $\emptyset$ in excess supply. An agent of type $t_i \in T$ has a strict utility function u_{t_i} over X_0. The utility of the null school is normalized to 0. In particular, all agents strictly prefer any of the proper schools to the null school.

We consider a simplified version of the Boston mechanism with a single round. The action space is the set of proper schools $A = S$, so that each student points to a school. If the number of students pointing to school s is lower than the number of seats, then all of those students are allocated to school s. If there are more students who point to s than its capacity, then students are randomly rationed, and those who do not obtain a seat in s are allocated to the null school. Formally, given a vector of reports a, the allocation $\Phi_i^n(a)$ assigns i to school a_i with probability

$$\min\{\frac{\lfloor q_{a_i} \cdot n \rfloor}{\text{emp}_{a_i}[a] \cdot n}, 1\},$$

and to the null school with the remaining probability. Consequently, the limit mechanism is

$$\phi^\infty(s,m) = \min\{\frac{q_s}{m_s}, 1\} \cdot s,$$

which denotes receiving school s with the probability $\min\{\frac{q_s}{m_s}, 1\}$, which we term the probability of acceptance to school s, and school $\emptyset$ with the remaining probability.

6.2 Limit Bayes-Nash Equilibria of the Boston Mechanism

Let $\Sigma^*(\mu)$ denote the set of limit Bayes-Nash equilibria of the Boston mechanism given prior μ. Let $P^*(\mu)$ denote the set of vectors of probabilities of acceptance over all equilibria in $\Sigma^*(\mu)$. The next Proposition establishes existence and some regularity properties of the equilibria of the Boston mechanism.

Proposition 1 (Structure of the set of limit equilibria)**.** *The correspondence $\Sigma^*(\mu)$ is non-empty, convex-valued and continuous in $\bar{\Delta}T$. The correspondence $P^*(\mu)$ is non-empty, single-valued, and continuous in $\bar{\Delta}T$.*

Given a prior μ, the Boston mechanism may have multiple equilibria.[15] Nevertheless, the probability of acceptance to each school is the same in any equilibrium because $P^*(\cdot)$ is single-valued. The intuition is that lowering the probability of acceptance to a school weakly reduces the set of students who would optimally point to it, and weakly increases the set of students who would point to other schools. Therefore, an argument similar to uniqueness arguments in competitive markets with gross substitutes shows that equilibrium probabilities of acceptance are unique. Moreover, equilibrium delivers well-behaved outcomes because probabilities of acceptance vary continuously.

Together, Proposition 1 and Theorem 2 yield the following corollary:

Corollary 1 (SP-L implementation of the Boston mechanism)**.** *The Boston mechanism has a continuous family of limit Bayes-Nash equilibria. For any such family $(\sigma^\infty_\mu)_{\mu \in \Delta T}$, the direct mechanism constructed according to equation (5.2) is SP-L, and, in the large market limit, for any prior, truthful play of the direct mechanism produces the same outcomes as Bayes-Nash equilibrium play of the Boston mechanism.*

Interestingly, the SP-L mechanism that we construct according to (5.2) closely resembles the Hylland and Zeckhauser (1979) pseudo-market mechanism for single-unit assignment.[16] In our constructed mechanism, agents report their types, the mechanism computes the equilibrium market-clearing probabilities P^*_s associated with the distribution of reports, and each student points to their most-preferred school given their reported types and the computed probabilities. In Hylland and Zeckhauser (1979)'s mechanism, agents report their types, the mechanism computes equilibrium market-clearing prices p^*_s given the distribution of reports, and each student purchases the lottery they like best given their reported types and the computed prices.

[15] To see why multiple equilibria are possible, consider for example an equilibrium where types t_1 and t_2 both point with probability $1/2$ to each school s_1 and s_2. In such case, there are other equilibria where types t_1 and t_2 change the proportion in which they point to each school in opposite directions, keeping the probabilities of acceptance the same.

[16] See also Miralles (2009), which contains a very nice description of the connection between the Boston mechanism's Bayes-Nash equilibria and Hylland and Zeckhauser (1979).

7 Extensions and Discussion

7.1 Relaxing Continuity

Theorem 2 assumes continuity of the given Bayes-Nash mechanism's family of equilibria. While this assumption has an intuitive appeal in that it asks that a mechanism's outcomes not be too sensitive to tiny changes in the prior, it is a strong assumption. Many well-known mechanisms violate it. For example, in pay-as-bid and uniform-price auctions, even though a small change in the prior typically has only a small effect on agents' bids, this small change in bids can have a large (i.e., discontinuous) effect on the number of units some bidder wins or the market-clearing price.

In Supplementary Appendix B we show that a weaker version of Theorem 2 obtains under a condition that we call quasi-continuity. Quasi-continuity allows for a family of equilibria to have discontinuities, with respect to both the prior on which agents' strategies are based and the empirical distribution of reports, but requires that the discontinuities are in a certain sense knife-edge. Roughly, any discontinuity is surrounded by regions in which outcomes are continuous. Under this condition, the conclusion of the theorem (Theorem B.1) is as follows. If the mechanism is continuous at a given prior μ_0, then, as before, there exists an SP-L mechanism that gives agents the same outcomes as the given Bayes-Nash mechanism, in the large-market limit. If the mechanism is not continuous at μ_0, then there exists an SP-L mechanism that gives agents a convex combination of the outcomes they would obtain under the original mechanism, for a set of priors in an arbitrarily small neighborhood of μ_0.

A question that remains open for future research is to fully characterize the conditions under which there is no gap between Bayes-Nash and SP-L in large markets. We have counterexamples that fail quasi-continuity in which our constructed mechanism does not approximate the original Bayes-Nash mechanism, even for the weaker form of approximation described above (cf. Supplementary Appendix B.2). However, the counterexamples that we have found are far from market design applications, and also the fact that our construction leaves a gap between Bayes-Nash and SP-L only proves that our method of proof does not work, it does not prove that there is a gap. It would also be desirable to obtain results analogous to Theorems 2 and B.1 in which the continuity conditions are defined not on family of equilibria, but on mechanisms themselves.

Given these open questions, we do not see Theorems 2 and B.1 as providing definitive proof that there is never an advantage to using Bayes-Nash over SP-L in large markets. Rather, we see the results as suggesting that, for the purposes of practical market design, a researcher is justified searching in the space of SP-L mechanisms rather than broadening her search to include Bayes-Nash. For there to be a meaningful gain to using Bayes-Nash over SP-L in large markets, the Bayes-Nash mechanism must fail quasi-continuity, which means that its outcomes are extremely sensitive to agents' beliefs and reports. In addition, the researcher must believe the usual conditions required for Bayes-Nash equilibrium, such as common knowledge and strategic sophistication, which seems unrealistic in the context of a highly discontinuous mechanism.

7.2 Semi-Anonymity

Our analysis focuses on mechanisms that are anonymous, meaning that each agent's outcome is a symmetric function of her own action and the distribution of all actions. In Supplementary Appendix C we generalize key definitions and results to the case of semi-anonymous mechanisms, as defined in Kalai (2004). A mechanism is semi-anonymous if each agent belongs to one of a finite number of groups, and her outcome is a symmetric function of her own action, her group, and the distribution of actions within each group. This generalization is useful for two reasons. First, it allows our analysis to cover more mechanisms. For instance, double auctions and matching markets are naturally modeled as semi-anonymous mechanisms. Second, it allows results and concepts stated for i.i.d. distributions to be extended to more general distributions.

7.3 Related Literature

Our paper is related to several lines of literature. First, there is a large theory literature that has studied how market size can ease incentive constraints for specific mechanisms. Important papers in this tradition include Roberts and Postlewaite (1976) on the Walrasian mechanism, Rustichini et al. (1994) and Cripps and Swinkels (2006) on double auctions, and Immorlica and Mahdian (2005) and Kojima and Pathak (2009) on the Gale-Shapley deferred acceptance algorithm. We discussed this literature in detail in Section 4.2. It is important to highlight that the aim of our paper is quite different from, and complementary to, this literature. Whereas papers such as Roberts and Postlewaite (1976) provide a defense of a *specific pre-existing mechanism* based on its approximate incentives properties in large markets, our paper aims to justify SP-L as a *general desideratum* for market design. In particular, our paper can be seen as providing justification for focusing on SP-L when designing *new* mechanisms. Another point of difference versus this literature is that our criterion itself is new; see fn. 11 for full details of the approximate incentives criteria used in this prior literature.

Second, there is an empirical literature that studies how participants behave in real-world non-SP market designs. One example is Abdulkadiroğlu et al. (2006), who show, in the context of the school choice system in Boston, that sophisticated students strategically misreport their preferences, while unsophisticated students frequently play dominated strategies; see Hwang (2014) and Agarwal and Somaini (2014) for related studies. Another example is Budish and Cantillon (2012), who show that students at Harvard Business School strategically misreport their preferences for courses, often sub-optimally, and that this misreporting harms welfare relative to both truthful play and optimal equilibrium behavior. We discuss this literature in more detail in Section 4.3. We view this literature as providing support for the concept of SP-L, since all of the examples we are aware of in which there is evidence of substantial harm from misreporting involves mechanisms that not only are not SP, but are not even SP-L.

Third, our paper is related to the literature on the role of strategy-proofness in practical market design. Wilson (1987) famously argued that practical market designs should aim to be robust to agents' beliefs, and Bergemann and Morris (2005) formalized the sense in which SP mechanisms are robust in the sense of Wilson. Several recent papers have argued that SP 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 our notion of SP-L approximates the appeal of SP, while at the same time being considerably less restrictive. Also, the distinction we draw between mechanisms that are SP-L and mechanisms that are manipulable even in large markets 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), Erdil and Klemperer (2010), Carroll (2013) and especially Pathak and Sönmez (2013), each of which proposes a method to compare non-SP mechanisms based on the degree to which they violate SP.[17] Parkes et al. (2001), Day and Milgrom (2008) and Erdil and Klemperer (2010) propose cardinal measures of a combinatorial auction's manipulability, based, respectively, on Euclidean distance from Vickrey prices, the worst-case incentive to misreport, and marginal incentives to misreport. Each of these papers then seeks to design a combinatorial auction that minimizes manipulability subject to other design objectives, such as efficiency.

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

Carroll (2013), like Day and Milgrom (2008), proposes a worst-case measure of manipulability, though, like us, he considers incentives to manipulate from an ex-interim rather than ex-post perspective. Carroll (2013) then compares voting rules based on the rate at which worst-case incentives to manipulate converge to zero. Pathak and Sönmez (2013) propose a partial order by which to compare non-SP 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 helps to explain several recent policy decisions in which school authorities switched from one manipulable mechanism to another, and offers another perspective on Milton Friedman's argument in favor of uniform-price auctions. We view our approach as complementary to these alternative approaches. A distinguishing feature of our approach is that it yields an explicit second-best design desideratum, namely that mechanisms be SP-L.

8 Conclusion

A potential interpretation of our results is that they suggest that SP-L be viewed as a necessary condition for good design in large anonymous settings. Our criterion provides a common language for criticism of mechanisms ranging from Friedman's (1960) criticism of pay-as-bid auctions, to Roth's (1990; 1991) criticism of priority-matching mechanisms, to Abdulkadiroğlu and Sönmez's (2003) criticism of the Boston mechanism for school choice. The issue is not simply that these mechanisms are manipulable, but that they are manipulable even in the large-market limit; even the kinds of agents we think of as "price takers" will want to misreport their preferences. The evidence we review in Section 4 suggests that manipulability in the large is a costly problem in practice, whereas the record for SP-L mechanisms, though incomplete, is quite positive. Our result in Section 5 then indicates that manipulability in the large can be avoided at approximately zero cost. Together, these results suggest that using a mechanism that is manipulable in the large is a preventable design mistake.

Whether SP-L can also be viewed as sufficient depends upon the extent to which the large-market abstraction is compelling in the problem of interest. Unfortunately, there is not a simple bright-line answer to the question of "how large is large".[18] But – just as

[18] Indeed, even in theoretical 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 double auctions. For instance, Rustichini et al. (1994) are able to show, in 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. Of course, in any specific context, the analyst's case that the market is large can be strengthened with empirical or computational evidence;

economists in other fields instinctively understand that there are some contexts where it is necessary to explicitly model strategic interactions, and other contexts where it may be reasonable to assume price-taking behavior – we hope that market designers will pause to consider whether it is necessary to restrict attention to SP mechanisms, or whether SP-L may be sufficient for the problem at hand.

see, for instance, Roth and Peranson (1999).

References

Abdulkadiroğlu, Atila and Tayfun Sönmez, "School Choice: A Mechanism Design Approach," *The American Economic Review*, 2003, *93* (3), 729–747.

_ , **Parag A. Pathak, Alvin E. Roth, and Tayfun Sönmez**, "Changing the Boston Mechanism: Strategyproofness as Equal Access," *Mimeo, Harvard University*, 2006.

_ , _ , **and** _ , "Strategy-Proofness Versus Efficiency in Matching with Indifferences: Redesigning the Nyc High School Match," *The American Economic Review*, 2009, *99* (5), 1954–1978.

_ , **Yeon-Koo Che, and Yosuke Yasuda**, "Resolving Conflicting Preferences in School Choice: The Boston Mechanism Reconsidered," *The American Economic Review*, 2011, *101* (1), 399–410.

Agarwal, Nikhil and Paulo Somaini, "Demand Analysis Using Strategic Reports: An Application to a School Choice Mechanism," *NBER Working Paper No 20775*, 2014.

Aumann, Robert J., "Existence of Competitive Equilibria in Markets with a Continuum of Traders," *Econometrica*, 1964, pp. 1–17.

Ausubel, Lawrence M. and Paul Milgrom, "The Lovely but Lonely Vickrey Auction," *Combinatorial Auctions (Cramton et al., eds)*, 2006.

Azevedo, Eduardo and Eric Budish, "Strategyproofness in the Large," *Working Paper. First version October 2011.*, 2013.

_ **and Jacob D. Leshno**, "A Supply and Demand Framework for Two-Sided Matching Markets," *Mimeo, Harvard University*, 2011.

Baliga, Sandeep and Rakesh Vohra, "Market Research and Market Design," *Advances in Theoretical Economics*, 2003, *3* (1).

Bergemann, Dirk and Stephen Morris, "Robust Mechanism Design," *Econometrica*, 2005, *73* (6), 1771–1813.

Bogomolnaia, Anna and Herve Moulin, "A New Solution to the Random Assignment Problem," *Journal of Economic Theory*, 2001, *100* (2), 295–328.

Brenner, Menachem, Dan Galai, and Orly Sade, "Sovereign Debt Auctions: Uniform or Discriminatory?," *Journal of Monetary Economics*, 2009, *56(2)*, 267–74.

Budish, Eric, "The Combinatorial Assignment Problem: Approximate Competitive Equilibrium from Equal Incomes," *Journal of Political Economy*, 2011, *119(6)*, 1061–1103.

_ **and Judd Kessler**, "Changing the Course Allocation Mechanism at Wharton," *Working Paper*, 2015.

Budish, Eric B. and Estelle Cantillon, "The Multi-Unit Assignment Problem: Theory and Evidence from Course Allocation at Harvard," *American Economic Review*, 2012, *102(5)*, 2237–71.

Budish, Eric, Gerard Cachon, Judd Kessler, and Abraham Othman, "Course Match: A Large-Scale Implementation of Approximate Competitive Equilibrium from Equal Incomes for Combinatorial Allocation," *Working Paper*, 2015.

_ , **Yeon-Koo Che, Fuhito Kojima, and Paul Milgrom**, "Designing Random Allocation Mechanisms: Theory and Applications," *American Economic Review*, 2013, *103(2).*

Carroll, Gabriel, "A Quantitative Approach to Incentives: Application to Voting Rules," *Mimeo, Stanford*, 2013.

Cramton, Peter, Yoav Shoham, and Richard Steinberg (eds.), *Combinatorial Auctions* 2006.

Cripps, Martin W. and Jeroen M. Swinkels, "Efficiency of Large Double Auctions," *Econometrica*, 2006, *74* (1), 47–92.

Day, Robert and Paul Milgrom, "Core-Selecting Package Auctions," *International Journal of Game Theory*, 2008, *36* (3), 393–407.

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

Ehlers, Lars and Bettina Klaus, "Coalitional Strategy-Proof and Resource-Monotonic Solutions for Multiple Assignment Problems," *Social Choice and Welfare*, 2003, *21(2)*, 265–80.

Erdil, Aytek and Paul Klemperer, "A New Payment Rule for Core-Selecting Package Auctions," *Journal of the European Economic Association*, 2010.

Featherstone, Clayton and Muriel Niederle, "School Choice Mechanisms Under Incomplete Information: An Experimental Investigation," 2011.

Friedman, Milton, *A Program for Monetary Stability*, Vol. 541 of *The Miller Lectures*, Fordham University Press, 1960.

_ , "How to Sell Government Securities," *Wall Street Journal*, 1991, p. A8.

Fudenberg, Drew and Jean Tirole, *Game Theory*, MIT Press, 1991.

Green, Jerry and Jean-Jacques Laffont, "Characterization of Satisfactory Mechanisms for the Revelation of Preferences for Public Goods," *Econometrica*, 1977, pp. 427–438.

Hatfield, John, "Strategy-Proof, Efficient, and Nonbossy Quota Allocations," *Social Choice and Welfare*, 2009, *33(3)*, 505–15.

Hatfield, John William, Fuhito Kojima, and Scott Duke Kominers, "Strategy-Proofness, Investment Efficiency, and Marginal Returns: an Equivalence," *Mimeo, Stanford University*, 2015.

Holmstrom, Bengt, "Groves Schemes on Restricted Domains," *Econometrica*, 1979, *47*, 1137–44.

Hortaçsu, Ali and David McAdams, "Mechanism Choice and Strategic Bidding in Divisible Good Auctions: An Empirical Analysis of the Turkish Treasury Auction Market," *Journal of Political Economy*, 2010.

_ **, Jakub Kastl, and Allen Zhang**, "Bid Shading and Bidder Surplus in U.S. Treasury Auction System," *Working Paper*, 2015.

Hurwicz, Leonid, "On Informationally Decentralized Systems," *Decision and Organization: A Volume in Honor of Jacob Marschak*, 1972, *12*, 297.

Hwang, Sam, "A Robust Redesign of High School Match," *Working Paper*, 2014.

Hylland, Aanund and Richard Zeckhauser, "The Efficient Allocation of Individuals to Positions," *The Journal of Political Economy*, 1979, pp. 293–314.

Immorlica, Nicole and Mohammad 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, Matthew O. and Alejandro M. Manelli, "Approximately Competitive Equilibria in Large Finite Economies," *Journal of Economic Theory*, 1997, *77* (2), 354–376.

Jegadeesh, Narasimhan, "Treasury Auction Bids and the Salomon Squeeze," *Journal of Finance*, 1993, pp. 1403–1419.

Kalai, Ehud, "Large Robust Games," *Econometrica*, 2004, *72* (6), 1631–1665.

Kastl, Jakub, "Discrete Bids and Empirical Inference in Divisible Good Auctions," *Review of Economic Studies*, 2011.

Kojima, Fuhito and Mihai Manea, "Incentives in the Probabilistic Serial Mechanism," *Journal of Economic Theory*, 2010, *145* (1), 106–123.

_ **and Parag A. Pathak**, "Incentives and Stability in Large Two-Sided Matching Markets," *The American Economic Review*, 2009, *99* (3), 608–627.

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

Liu, Qingmin and Marek Pycia, "Ordinal Efficiency, Fairness, and Incentives in Large Markets," 2011.

Malvey, Paul F. and Christine M. Archibald, "Uniform-Price Auctions: Update of the Treasury Experience," *Department of the Treasury, Office of Market Finance*, 1998.

Milgrom, Paul, *Putting Auction Theory to Work*, Cambridge University Press, 2004.

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

Miralles, Antonio, "School Choice: The Case for the Boston Mechanism," *Auctions, Market Mechanisms and Their Applications*, 2009, pp. 58–60.

Myerson, Roger B., "Incentive Compatibility and the Bargaining Problem," *Econometrica: Journal of the Econometric Society*, 1979, pp. 61–73.

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

Parkes, David C., Jayant R. Kalagnanam, and Marta Eso, "Achieving Budget-Balance with Vickrey-Based Payment Schemes in Exchanges," in "Proceedings of the Seventeenth International Joint Conference on Artificial Intelligence," Vol. 17 AAAI Press August 2001, pp. 1161–1168.

Pathak, Parag A. and Tayfun 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," *American Economic Review*, 2013, *103(1)*, 80–106.

Roberts, Donald J. and Andrew Postlewaite, "The Incentives for Price-Taking Behavior in Large Exchange Economies," *Econometrica*, 1976, pp. 115–127.

Roth, Alvin E., "The Economics of Matching: Stability and Incentives," *Mathematics of Operations Research*, 1982, pp. 617–628.

_ , "New Physicians: A Natural Experiment in Market Organization," *Science*, 1990, *250*, 1524–8.

_ , "A Natural Experiment in the Organization of Entry Level Labor Markets: Regional Markets for New Physicians and Surgeons in the U.K.," *American Economic Review*, 1991, *81*, 415–40.

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

Rustichini, Aldo, Mark A. Satterthwaite, and Steven R. Williams, "Convergence to Efficiency in a Simple Market with Incomplete Information," *Econometrica*, 1994, pp. 1041–1063.

Segal, Ilya, "Optimal Pricing Mechanisms with Unknown Demand," *The American Economic Review*, 2003, *93* (3), 509–529.

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

Swinkels, Jeroen M., "Efficiency of Large Private Value Auctions," *Econometrica*, 2001, *69* (1), 37–68.

Wilson, Robert, "Game-Theoretic Analyses of Trading Processes," in Truman F. Bewley, ed., *Advances in Economic Theory: Fifth World Congress*, Cambridge, UK: Cambridge University Press, 1987, chapter 2, pp. 33–70.