A Pseudo-Market Approach to Allocation with Priorities*

Yinghua He   Antonio Miralles   Marek Pycia   Jianye Yan

July 2017

Abstract

We propose a pseudo-market mechanism for no-monetary-transfer allocation of indivisible objects based on priorities such as those in school choice. Agents are given token money, face priority-specific prices, and buy utility-maximizing random assignments. The mechanism is asymptotically incentive compatible, and the resulting assignments are fair and constrained Pareto efficient. Aanund Hylland & Richard Zeckhauser (1979)'s position-allocation problem is a special case of our framework, and our results on incentives and fairness are also new in their classical setting.

Keywords: *Priority-based allocation, Efficiency, Stability, Incentive Compatibility, Pseudo-Market Approach*

JEL Codes: C78, D82, I29

*First posted draft: May 2010. This papers subsumes earlier drafts by Miralles, by He and Yan, by Pycia, by He, Miralles, and Yan, and the section on allocation with priorities in He's PhD dissertation. We thank Omer Ali, Eduardo Azevedo, Yeon-Koo Che, Renato Gomes, Christian Hellwig, Fuhito Kojima, Sanxi Li, Jordi Massó, Marcin Peski, Bernard Salanié, Takuro Yamashita, William Zame, Simpson Zhang, and audiences at numerous seminars and conferences for their valuable comments. Financial support from the NSF (1730636) and the French National Research Agency (l'Agence Nationale de la Recherche) through Projects DesignEdu (ANR-14-CE30-0004) and FDA (ANR-14-FRAL-0005) to He, that from the Ramón y Cajal contract (RYC2011-09665) of the Spanish Government, the Spanish Plan Nacional I+D+I (ECO2014-53051-P), the Generalitat de Catalunya (SGR2014-505), and the Severo Ochoa program (SEV-2015-0563) to Miralles, that from the PER Visitor Program at Columbia University and the William S. Dietrich II Economic Theory Center at Princeton University to Pycia, and that from the National Natural Science Foundation of China through grant 71401034 and the Fundamental Research Funds for the Central Universities in UIBE to Yan are gratefully acknowledged. He: Rice University and Toulouse School of Economics; yinghua.he@gmail.com. Miralles: Universitat Autonoma de Barcelona and Barcelona Graduate School of Economics; antonio.miralles@uab.cat. Pycia: UCLA and U Zurich. Yan: Center for New Structural Economics, Peking University; School of Banking and Finance, University of International Business & Economics; email: yanjianye@gmail.com.

We study the allocation of indivisible objects where monetary transfers are precluded and agents demand at most one object. Examples include student placement in public schools (where an object corresponds to a school seat and each object has multiple copies) and allocation of work or living space (where each object has exactly one copy). A common feature of these settings is that agents are prioritized. For instance, students who live in a school's neighborhood or have siblings in the school may enjoy admission priority at this school over those who do not, and the current resident may have priority over others in the allocation of the dormitory room he or she lives in.

Due to the lack of monetary transfers, objects in these environments are very often allocated by a centralized mechanism which maps agents' reported preferences to an allocation outcome. The outcome, known as assignment, can be either deterministic or random. The former dictates who gets what object, and the latter prescribes the probability shares of objects that each agent obtains and thus is a lottery over a set of deterministic assignments.

The standard allocation mechanisms used in practice and studied in the literature are ordinal: students are asked to rank schools or rooms, and the profile of submitted rankings determines the assignment. However, Antonio Miralles (2008) and Atila Abdulkadiroglu, Yeon-Koo Che & Yosuke Yasuda (2011) pointed out that we may implement Pareto dominant assignments by eliciting agents' cardinal utilities, which are their relative intensities of preferences over objects and their rates of substitution between probability shares in objects. Furthermore, Qingmin Liu & Marek Pycia (2012) and Marek Pycia (2014) showed that sensible ordinal mechanisms are asymptotically equivalent in large markets, while mechanisms eliciting cardinal utilities maintain their efficiency advantage.[1] Naturally, with more inputs, we expect a mechanism to deliver a better outcome, as cardinal preferences are more informative than ordinal ones. However, what has not been answered in the literature is how to use cardinal information efficiently.

This paper aims to fill this gap by providing a novel cardinal mechanism to improve upon

[1]The data on Boston and NYC school choice corroborates both the equivalence of ordinal mechanisms (see e.g., Parag A. Pathak & Tayfun Sonmez (2008) and Atila Abdulkadiroglu, Parag A. Pathak & Alvin E. Roth (2009)) and the inefficiency of ordinal mechanisms (Atila Abdulkadiroglu, Nikhil Agarwal & Parag A. Pathak 2015). For analysis of ordinal mechanisms see the seminal work of Atila Abdulkadiroglu & Tayfun Sonmez (2003) and Anna Bogomolnaia & Herve Moulin (2001). The literature discussion below also includes other papers emphasizing the need to elicit cardinal information.

the ordinal mechanisms. The mechanism is asymptotically incentive compatible, fair, and constrained efficient among ex-ante stable and fair mechanisms. A mechanism is ex-ante stable if, in any of its resulting assignment, no probability share of an object is given to an agent with lower priority at this object whenever a higher-priority agent is obtaining some probability shares in any of his/her less preferred objects (Onur Kesten & M. Utku Ünver 2015). Furthermore, every deterministic assignment that is compatible with an ex-ante stable random assignment eliminates all justified envy and thus satisfies stability (Abdulkadiroglu & Sonmez 2003). We use the strong fairness concept, equal claim, proposed by Yinghua He, Sanxi Li & Jianye Yan (2015); a mechanism satisfies equal claim if agents with the same priority at an object are given the same opportunity to obtain it.

We refer to our construction as the pseudo-market (PM) mechanism which elicits cardinal preferences from agents and delivers an assignment. If it is a random assignment, one can then conduct a lottery to implement one of the compatible deterministic assignments. To map reported preferences into assignments, PM internally solves a Walrasian equilibrium, where prices are priority-specific and the mechanism chooses probability shares to maximize each agent's expected utility given his/her reported preferences and an exogenous budget in token money. Budgets need not be equal across agents.

This Walrasian equilibrium used in the internal computation of the PM mechanism has a unique feature in its priority-specific prices: for each object, there exists a cut-off priority group such that agents in priority groups strictly below the cut-off face an infinite price for the object (hence, they can never be matched with the object), while agents in priority groups strictly higher than the cut-off face zero price for the object. By incorporating priorities in this manner, the PM mechanism extends the canonical Hylland & Zeckhauser (1979) mechanism which requires every agent to face the same prices and thus does not allow priorities. It is also a generalization of the Gale-Shapley Deferred-Acceptance (DA) mechanism, the most celebrated ordinal mechanism. Essentially, when both agents and objects have strict rankings over those on the other side, the DA mechanism eliminates all justified envy; whenever there are multiple agents in one priority group of an object, the tie has to be broken, usually in an exogenous way. The PM mechanism, instead, has ties broken endogenously and efficiently by using information on cardinal preferences. Agents with relatively higher cardinal preferences for an object obtain shares of that object before others who are in the same priority group.

We show that the PM mechanism is well-defined in the sense that it can always internally find a Walrasian equilibrium and deliver an assignment given any reported preference profile. Moreover, the mechanism is shown to be asymptotically incentive compatible in regular economies, where regularity guarantees that Walrasian prices are well defined as in the classical analysis of Walrasian equilibria (see e.g., Egbert Dierker (1974), Werner Hildenbrand (1974), and Matthew O. Jackson (1992)). The latter result is also new in the original Hylland & Zeckhauser (1979) problem and proves the long-standing conjecture they formulated.[2] As in the setting without priorities (see e.g., Abdulkadiroglu, Che & Yasuda (2011) and Pycia (2014)), the PM mechanism allows one to achieve higher social welfare than mechanisms eliciting only ordinal preferences such as the DA and the Probabilistic Serial mechanisms.

The PM mechanism is ex-ante stable because of our design of the priority-specific prices. Given an object s and its cut-off priority group, whenever a lower-priority agent obtains a positive share of s, a higher-priority agent must face a zero price for s, and, therefore, is never assigned to an object they prefer less than s.

We study fairness of the PM mechanism in the sense of *equal claim*, which requires that, for any given object, agents with the same priority are given the same opportunity to obtain this object.[3] Since prices for agents in the same priority group are by construction the same in the PM mechanism, we can conclude that equal claim is satisfied when agents are given equal budgets. Furthermore, we show that the PM mechanism in which agents have equal budgets is the only non-wasteful mechanism that satisfies ex-ante stability and equal claim.

Focusing on assignments that are ex-ante stable and equal-claim, we analyze efficiency: an assignment is constrained Pareto efficient if no other assignment that satisfies ex-ante

[2]Stating the true preferences in the PM mechanism is not always a dominant strategy for every agent. Hylland & Zeckhauser (1979) give an example where there are incentives for agents to misreport their preferences when objects do not rank agents. More generally, one calls a mechanism strategy-proof if reporting true preferences is a dominant strategy; Alvin E. Roth (1982) and Lin Zhou (1990) show strategy-proofness is in conflict with other desirable properties. In addition to proving the asymptotic incentive compatibility of the PM mechanism in regular economies, we also prove that it is limiting incentive compatible in the sense of Donald John Roberts & Andrew Postlewaite (1976).

[3]See He, Li & Yan (2015) for an analysis of this concept in the setting without priorities. Note that equal claim does not imply that same-priority agents at an object receive the same probability share of that object in the final assignment.

stability and equal claim dominates it in terms of agents' welfare.[4] An important corollary of our results is that a constrained Pareto-efficient assignment is always an outcome of the PM mechanism with equal budgets.

One may be interested in two-sided efficiency if the priority structure is closely related to object suppliers' preferences, e.g., when schools' priority ranking over students reflect a school district's preferences. An assignment is ex-ante two-sided Pareto efficient if it is not Pareto dominated by any other assignment with respect to both agents' expected utilities and objects' priorities treated as their ordinal preferences. When the welfare of objects is evaluated in terms of first-order stochastic dominance with respect to priorities,[5] PM always delivers assignments that satisfy ex-ante two-sided efficiency.

The PM mechanism is therefore a promising candidate that can be used in school choice, dormitory room allocation, and other allocation problems based on priorities. Moreover, it is flexible enough to accommodate additional constraints such as group-specific quotas.

Literature Review The early literature on school choice, the focal topic of priority-based allocation, e.g., Abdulkadiroglu & Sonmez (2003) and Haluk Ergin & Tayfun Sonmez (2006), followed the two-sided matching literature where it is common to assume that both sides strictly rank the other side. Implicitly, weak priorities are augmented with random lotteries to create strict priorities. It has been noted that when priorities are coarse, some issues arise. For example, stability no longer implies Pareto efficiency (Aytek Erdil & Haluk Ergin 2006); and, more importantly, how ties are broken affects the welfare of agents since it introduces artificial constraints. Extending the DA mechanism, Aytek Erdil & Haluk Ergin (2008) propose an algorithm for breaking priority ties and the computation of agent-efficient stable matchings when priority rankings are weak and only ordinal information is elicited. The two algorithms proposed by Kesten & Ünver (2015) offer further ways to break priority-ties in ordinal settings. However, Abdulkadiroglu, Pathak & Roth (2009) and Onur Kesten (2010) show that the inefficiency associated with a realized tie breaking in ordinal setting

[4]The literature on ordinal mechanisms that follows Bogomolnaia & Moulin (2001) defines efficiency in terms of first-order stochastic dominance; since we study expected-utility-maximizing agents, we can use the standard Pareto efficiency concept. It should be noted that there are priority structures and stable assignments that are Pareto dominated by assignments that are not stable; Abdulkadiroglu & Sonmez (2003) construct relevant examples in the ordinal setting, and the same examples remain valid in our setting.

[5]That is, an object or object supplier is better off if agents matched with this object in the new assignment first-order stochastically dominate those of the old one.

cannot be removed without harming student incentives.

Noting that agents may differ in their cardinal preferences, a strand of literature (e.g., Clayton Featherstone & Muriel Niederle (2008), Miralles (2008), Abdulkadiroglu, Che & Yasuda (2011), Peter Troyan (2012), Pycia (2014), Atila Abdulkadiroglu, Yeon-Koo Che & Yosuke Yasuda (2015), and Itai Ashlagi & Peng Shi (2016)) emphasizes the importance of eliciting signals of cardinal preferences from agents in matching mechanisms.[6] Ties in priorities can be broken with such signals, although the space of preference profiles or signals considered in these papers is restricted. Our PM mechanism elicits the entire relevant utility information in a general setting. Moreover, compared to discrete signals of cardinal preferences such as those in the popular Boston mechanism (defined in Appendix A), the PM mechanism has the advantage of being (asymptotically) incentive compatible. It has been shown theoretically (e.g., Pathak & Sonmez (2008)), experimentally (e.g., Yan Chen & Tayfun Sonmez (2006)), and empirically (e.g., Atila Abdulkadiroglu, Parag A. Pathak, Alvin E. Roth & Tayfun Sonmez (2006), Yinghua He (2012)), that strategic considerations may put less sophisticated agents at a disadvantage. More importantly, these effects do not disappear in large markets (Eduardo Azevedo & Eric Budish 2012). PM thus “levels the playing field” by eliminating this strategic concern while keeping the benefits of using cardinal preferences.

Our paper offers the first pseudo-market construction with priority constraints.[7] In addition, we also contribute to the growing literature on pseudo-market mechanisms in settings without priorities. The idea was first formulated by Hylland & Zeckhauser (1979). Miralles (2008) establishes a connection between the mechanism and the Boston mechanism in settings without priorities. Eric Budish (2011) and Eric Budish, Yeon-Koo Che, Fuhito Kojima & Paul Milgrom (2013) extend the pseudo-market mechanism to multi-unit demand settings such as course scheduling. Antonio Miralles & Marek Pycia (2014) show that every efficient assignment can be decentralized through prices, establishing the Second Welfare Theorem

[6] In recent work, SangMok Lee & Leeat Yariv (2014) and Yeon-Koo Che & Olivier Tercieux (2014) show that when agent’s utilities come from independent distributions, some ordinal mechanisms can be efficient.

[7] Notice that our paper subsumes Yinghua He (2011), Antonio Miralles (2011), Yinghua He & Jianye Yan (2012), and Yinghua He, Antonio Miralles & Jianye Yan (2012) who proposed this construction and proved it is well-defined. Subsequent work on personalized prices in pseudo-markets (e.g. Ashlagi & Shi (2016) and He, Li & Yan (2015)) did not address the question of when personalized-price mechanisms respect priority constraints.

for the no-transfer setting without priorities. He, Li & Yan (2015) make the point that any assignment, not necessarily efficient, can be decentralized by personalized prices.

Our analysis of fairness is related to Ashlagi & Shi (2016) who study a model with a continuum of agents without priorities and show that the equal-budget PM mechanism can implement any envy-free and Pareto efficient assignment. Envy-freeness is a weaker fairness property than equal claim, and the characterization of the equal-budget pseudo market in terms of envy-freeness and efficiency does not extend to large finite economies (see Antonio Miralles & Marek Pycia (2015)).

Our analysis of the PM mechanism's asymptotic incentive compatibility addresses a long-standing open problem posed by Hylland & Zeckhauser (1979). We build on the classic literature on the price-taking behavior of agents in exchange economies, e.g., Roberts & Postlewaite (1976) and Jackson (1992). The only earlier analysis of incentive compatibility of PM without priorities is Azevedo & Budish (2012) who show that it satisfies the strategy-proofness-in-the-large criterion that they introduce provided that budgets are equal and the number of utility types is finite and stays bounded as the market grows. Our result does not hinge on either of these assumptions.[8]

Organization of the Paper Section I sets up the model for the priority-based allocation problem. Section II defines the PM mechanism and establishes that it is well-defined. Section III investigate its incentive compatibility. We present fairness properties of the mechanism and its characterization in Section IV. Section V discusses results on its efficiency advantage relative to some well known mechanisms. The paper concludes in Section VI.

[8]The equal-budget PM mechanism satisfies a restriction of Azevedo and Budish's, EF-TB (envy-free but for tie breaking), among agents of the same priorities at all objects. Building on this observation we show that the equal-budget PM mechanism is strategy-proof-in-the-large provided that their environment assumptions hold true. We provide the details in Appendix C. Work on other related mechanisms includes Antonio Miralles (2012), Pycia (2014), and Isa Emin Hafalir & Antonio Miralles (2014), studying incentive compatible, efficient mechanisms in specific parametric settings. Tadashi Hashimoto (2013) constructs an *ex post* incentive-compatible mechanism that becomes efficient in large markets. Thanh Nguyen, Ahmad Peivandi & Rakesh Vohra (2015) introduce an optimization-based efficient mechanism that is strategy-proof-in-the-large. The asymptotic incentive properties of ordinal matching mechanisms have also been studied, e.g., Fuhito Kojima & Parag A. Pathak (2009), Fuhito Kojima & Mihai Manea (2010), Fuhito Kojima, Parag A. Pathak & Alvin E. Roth (2013), Liu & Pycia (2012), Marek Pycia (2011), SangMok Lee (2014), and Itai Ashlagi, Yash Kanoria & Jacob D. Leshno (2014).

I Model

We consider a priority-based allocation problem, or an economy, $\Gamma = \{\mathcal{S}, \mathcal{I}, Q, V, K\}$, where:

(i) $\mathcal{S} = \{s\}_{s=1}^{S}$ is a set of objects;

(ii) $\mathcal{I} = \{i\}_{i=1}^{I}$ is a set of agents, each of whom is to be matched with exactly one copy of an object;

(iii) $Q = [q_s]_{s=1}^{S}$ is a capacity vector, and $q_s \in \mathbb{N}$ is the supply of object s, $\forall s$. For simplicity, we assume that $\sum_{s=1}^{S} q_s = I$, i.e., there are just enough copies of objects to be allocated to agents; the extension to $\sum_{s=1}^{S} q_s \neq I$ is straightforward.

(iv) $V = [v_i]_{i\in\mathcal{I}}$, where $v_i = [v_{i,s}]_{s\in\mathcal{S}}$ and $v_{i,s} \in [0,1]$ is agent i's von Neumann-Morgenstern (vN-M) utility associated with object s.

(v) $K = [k_{s,i}]_{i\in\mathcal{I},s\in\mathcal{S}}$, where $k_{s,i} \in \mathcal{K} \equiv \left\{1, 2, ..., \overline{k}\right\}$ is the priority group of agent i at object s, and $\overline{k}\,(\leq I)$ is the maximum number of priority groups.[9] Roughly speaking, $[k_{s,i}]_{i\in\mathcal{I}}$ can be interpreted as s's weak ranking over all agents, and a **lower value** of $k_{s,i}$ means **higher priority**. That is, $k_{s,i} < k_{s,j}$ if and only if i has a higher priority at object s than j's. We allow both strict and coarse priority structures, in particular, the special case of interest when all agents have the same priority (the no-priority case).[10]

All objects and agents are acceptable to the other side, i.e., every agent considers every object better than being unassigned and is qualified to be assigned to any object. The analysis can be extended to the setting with unacceptable objects/agents. Agents are assigned to objects under the unit-demand constraint such that each agent must be matched with exactly one copy of an object. In the following, unless otherwise stated, we require non-wastefulness such that all copies of every object are to be assigned to some agents. Given the acceptability of everyone on both sides, wastefulness clearly leads to Pareto inefficiency.

An *assignment* is a matrix $\Pi = [\pi_i]_{i\in\mathcal{I}}$; $\pi_i = [\pi_{i,s}]_{s\in\mathcal{S}}$ and $\pi_{i,s} \in [0,1]$ is agent i's probability share of object s, or the probability that agent i is matched with object s. Given the

[9] It is innocuous to assume that every object has the same number of priority groups, as there might be no agent in a particular group of an object.

[10] Our results on incentive compatibility and fairness are also new in this classical case.

supply of the objects, an assignment is feasible if and only $\sum_{i\in\mathcal{I}}\pi_{i,s}\leq q_s$ for all s. The set of all feasible assignments is denoted by $\mathcal{A}$. Moreover, the unit-demand constraint implies that $\sum_{s\in\mathcal{S}}\pi_{i,s}=1$ for all i, and non-wastefulness leads to $\sum_{i\in\mathcal{I}}\pi_{i,s}=q_s$ for all s.

Because an assignment $\Pi\in\mathcal{A}$ is defined in terms of probability shares, Π is commonly known as random assignment; if, however, Π is degenerate, i.e., $\pi_{i,s}\in\{0,1\}$ for all i and s, it is also a deterministic assignment. Every feasible random assignment can be decomposed into a convex combination of deterministic assignments and can therefore be resolved into deterministic assignments (Kojima & Manea 2010), which generalizes the Birkhoff-von Newmann theorem (Garrett Birkhoff 1946, John von Neumann 1953). Notice that the convex combination may not be unique in general.

Given objects' priorities and supply, a matching mechanism is a mapping from agents' reported preferences, either cardinal or ordinal, to the space of feasible assignments, $\mathcal{A}$.

II The Pseudo-Market Mechanism

Given the structure of priorities K and the capacities Q, the pseudo-market (PM) mechanism maps a reported utility profile $V=[v_i]_{i\in\mathcal{I}}$ to a feasible assignment $[\pi_i]_{i\in\mathcal{I}}\in\mathcal{A}$ by internally finding a Walrasian equilibrium: it takes the exogenous budgets in token money $[b_i]_{i\in\mathcal{I}}$, $b_i\in(0,1]$, and finds a price matrix $P=[p_{s,k}]_{s\in\mathcal{S},k\in\mathcal{K}}\in\mathcal{P}\equiv[0,+\infty]^{S\times\bar{k}}$, where $p_{s,k}$ is the price of object s for agents in s's priority group k, by solving the utility maximization problem for every i,[11]

$$\pi_i\left(v_i,P\right)\quad\in\quad\arg\max_{\pi_{i,s}}\sum_{s\in\mathcal{S}}\pi_{i,s}v_{i,s}$$

subject to:

(i) unit-demand constraint: $\sum_{s\in\mathcal{S}}\pi_{i,s}=1$ for all i;[12]

[11] If $p_{s,k}=+\infty$, we define $+\infty\cdot 0=0$ and $+\infty\cdot\pi_{i,s}=+\infty$ if $\pi_{i,s}>0$.

[12] We model the unit-demand constraint as equalities. In other words, an agent's preferences over probability-share bundles that do not satisfy the equality are not defined in the mechanism (similar to preferences of having two spouses in a one-to-one marriage-matching model). This definition allows practitioner to announce in advance that every participant in the mechanism is guaranteed a copy of some object.

Alternatively, one can define the unit-demand constraint as weak inequalities, $\sum_{s\in\mathcal{S}}\pi_{i,s}\leq 1$. This amounts

(ii) feasibility constraint: $\sum_{i\in\mathcal{I}} \pi_{i,s}(v_i, P^*) \leq q_s$ for all objects s;

(iii) budget constraint: $\sum_{s\in\mathcal{S}} p_{s,k_{s,i}} \pi_{i,s} \leq b_i$ and the stipulation that if there are multiple bundles maximizing her expected utility then a cheapest one is chosen.

(iv) priority constraint:[13] $k^*(s)$ is the cut-off priority of object s if $\sum_{i\in\mathcal{I},\ k_{s,i}<k^*(s)} \pi_{i,s}(v_i, P^*) < q_s$ and $\sum_{i\in\mathcal{I},\ k_{s,i}\leq k^*(s)} \pi_{i,s}(v_i, P^*) = q_s$; moreover,

$$\begin{aligned} p^*_{s,k_{s,i}} &= 0, \text{ if } k_{s,i} < k^*(s), \\ p^*_{s,k_{s,i}} &\in [0, +\infty), \text{ if } k_{s,i} = k^*(s), \\ p^*_{s,k_{s,i}} &= +\infty, \text{ if } k_{s,i} > k^*(s). \end{aligned}$$

The PM mechanism accommodates personalized exogenous budgets, but to economize on notations, we focus on the mechanism with equal budgets such that $b_i = 1$ for all i and refer to the **equal-budget PM mechanism** simply as the **PM mechanism**. It should be noted that all results except fairness in Section IV extend to the mechanism with unequal budgets as long as budgets do not depend on reported utilities.

Given a reported utility profile, an assignment that can be resulted from the PM mechanism is a *PM assignment*. A price matrix in the internal Walrasian equilibrium of the mechanism is called a *PM price matrix* or simply *PM prices*.[14] To the extent that PM assignments crucially depends on PM prices, we study the properties of prices to determine assignment characteristics.

A unique feature of the PM mechanism is that the prices are designed to be priority-specific and increase when we move down on the priority list. If s is consumed completely by agents in priority groups higher than $k^*(s)$ (including $k^*(s)$), agents in s's priority groups strictly below $k^*(s)$ face an infinite price, while those in priority groups strictly higher than $k^*(s)$ face a zero price. Section IV discusses the implications of such a price structure.

to assuming that agents receive zero utility if unassigned. A PM mechanism can be similarly defined, and the weak inequalities can be satisfied as equalities if $v_{i,s} \geq 0$ for all i and s, given that there is enough supply.

[13] In the setting in which all agents have the same priority at all objects (i.e., no priorities), this priority constraint reduces to prices being non-negative and finite.

[14] By construction, a PM price matrix is also Walrasian equilibrium prices for the economy Γ augmented with the given budgets.

In this manner, the PM mechanism treats objects' priorities as agents' rights to obtain an object at a lower, sometimes zero, price. Whenever some agents with lower priorities can get a positive share of an object, an agent with a higher priority at that object can always get it for free. More importantly, agents can choose not to exercise the right if they wish, but they cannot trade priorities. This interpretation is similar to the consent in Kesten (2010) that allows agents to waive a certain priority at an object, but is in contrast to the treatment in the top-trading-cycles mechanism which implicitly allows agents to trade their priorities (Abdulkadiroglu & Sonmez 2003).

Our first main result is the existence of the PM prices and assignment.

Theorem 1 *Given any reported utility profile, there always exist a PM price matrix, and thus the PM mechanism can always deliver a PM assignment.*

Sketch of the Proof. The proof uses the traditional Kakutani's fixed point theorem, applied to a price matrix instead of a price vector. Our price space contains two features worth mentioning. First, some prices can be infinite but the price space remains compact.[15] Second, at high prices, e.g. at prices that are all above the agents' budget, no agent can afford to buy one unit of any object; the unit demand constraint is violated at such prices but, as we show, not at equilibrium prices. In the proof we rely on an "artificial outside option" that is infinitely supplied and always zero priced, and that for all agents is strictly worse than any other object in the original economy. We show that there is an equilibrium of this extended economy. Furthermore, because of our assumptions on total supply of objects and because excess supply implies zero price at all priority groups, no equilibrium would contain positive demand of the "artificial outside option". Thus, the equilibrium of the extended economy gives us an equilibrium of the original economy. ■

This key result shows that the PM mechanism is well-defined. The analogous result was proven by Hylland & Zeckhauser (1979) for the classical economy without priorities. The result is new for the case with priorities; the challenge in obtaining the result is the need to incorporate the priority condition (iv).[16] This condition is crucial in ensuring the fairness of

[15] Although not needed for the proof, we transform price matrix using the continuous function *arctangent*. Transformed prices are always between 0 and $\pi/2$, which may be convenient in practical computation.

[16] Our paper subsumes He and Yan's and Miralles's work, who independently proposed the construction of the PM mechanism with priority constraints.

the mechanism under priority constraints (Section IV).

An economy can have more than one PM price matrix and multiple PM assignments, and thus a complete specification of the mechanism must prescribe a price selection rule.[17] Our main results are robust to arbitrary selection rules, except those on incentive compatibility in the next section, which address the selection issue directly.

III Asymptotic Incentive Compatibility

Our next analysis focuses on asymptotic incentive compatibility in sequences of replica economies and considers the PM mechanism's incentive properties in large markets.[18] For any base economy $\Gamma = \{\mathcal{S}, \mathcal{I}, Q, V, K\}$, we use $\Gamma^{(n)} = \left\{\mathcal{S}, \mathcal{I}^{(n)}, Q^{(n)}, V^{(n)}, K^{(n)}\right\}$ to denote an n-fold replica of Γ which is an economy such that: (i) for each $i \in \mathcal{I}$, there are n copies of i in $\mathcal{I}^{(n)}$ whose preferences and priorities are exactly the same as i; (ii) $\mathcal{S}$ is constant in all economies; and (iii) $Q^{(n)} = nQ$, or equivalently $q_s^{(n)} = nq_s$ for all s and n. In the sequence of replica economies $\left\{\Gamma^{(n)}\right\}_{n\in\mathbb{N}}$, each $\Gamma^{(n)}$ has n copies of the base economy Γ. Notice that the set of PM prices is constant along any sequence of replica economies, provided that all agents report truthfully.[19]

We consider a natural analogue of regular economies from the general equilibrium literature (e.g., Dierker (1974), Hildenbrand (1974), Jackson (1992)).[20] To define this regularity concept, we use the Prohorov metric ρ to measure the distance between two distributions, μ and ν:

$$\rho(\mu,\nu) = \inf\left\{\varepsilon > 0 | \nu(E) \leq \mu(B_\varepsilon(E)) + \varepsilon \ \&\ \mu(E) \leq \nu(B_\varepsilon(E)) + \varepsilon,\ E \subseteq [0,1]^{S\times I} \text{ Borel}\right\}.$$

[17] In the market design literature, Alexander Kovalenkov (2002) is an exception to explicitly consider selection rules in an approximate Walrasian mechanism.

[18] Given the impossibility result in Zhou (1990) and the example in Hylland & Zeckhauser (1979), it is known that agents may have incentives to misreport their preferences in any finite market. See Appendix C for an analysis of the limit incentive compatibility concepts defined by Azevedo & Budish (2012) and Roberts & Postlewaite (1976).

[19] We make the replica assumption for simplicity but our analysis does not depend on it. See the end of this section for an extension beyond replica economies.

[20] For simplicity, we follow Jackson (1992) in defining regularity directly in terms of price behavior; alternatively we could express the definition of regularity in terms of properties of excess demand functions as in Dierker (1974) and Hildenbrand (1974).

A distribution of utilities μ^* is **regular** if there exists a neighborhood B of μ^* and a finite number $m > 0$ of continuous functions $\psi_1, ..., \psi_m$ from B to $[0, +\infty]^{S \times \bar{k}}$ such that for every distribution $\mu \in B$ the set of PM prices is $\{\psi_1(\mu), ..., \psi_m(\mu)\}$ and $\psi_i(\mu) \neq \psi_j(\mu)$ for every $i \neq j$. An economy Γ is regular if the corresponding distribution of utilities is regular. The proofs for this section show that if the base economy is regular then so is any replica economy.

Our second main result is the asymptotic incentive compatibility of the PM mechanism. A mechanism is **asymptotically incentive compatible** on a sequence of replica economies $\Gamma^{(n)}$ if for every agent the utility gain from submitting a utility profile different from the truth vanishes along the sequence. That is, for every $\varepsilon > 0$, there exists n^* such that $n > n^*$ implies that the utility gain from unilateral misreporting for every agent in $\Gamma^{(n)}$ is bounded by ε when everyone else is truth-telling.

Theorem 2 *There always exists a selection of PM prices in the definition of the PM mechanism such that the resulting PM mechanism is asymptotic incentive compatible on any sequence of replica economies whose base economy has a regular distribution of utilities.*

The above theorem shows that the utility gain from unilateral misreporting is bounded for *all agents* in a large enough economy. An analogue of this result remains true beyond replica economies: our proof of Theorem 2 also implies that the gain from manipulation for any agent who is present in all economies in a sequence vanishes as the economy grows, provided that the limit distribution of utilities is regular.[21]

Theorem 2 is new not only in the setting with priorities, but also in the canonical setting without priorities first studied by Hylland & Zeckhauser (1979).[22] While Hylland and Zeckhauser conjectured that their mechanism is asymptotically incentive compatible, their conjecture has so far remained open. The closest prior result was obtained by Azevedo & Budish (2012) who introduced the concept of strategy-proofness-in-the-large and in a discrete setting proved that every envy-free mechanism is incentive compatible in their sense. In particular, their result implies that Hylland and Zeckhauser's mechanism with equal budgets is strategy-proof-in-the-large in large economies with a bounded number of utility types.

[21] In addition, Appendix C shows that the PM mechanism is limiting incentive compatible in the sense of Roberts & Postlewaite (1976).

[22] See also e.g., Budish et al. (2013)

Their approach hinges both on the equality of budgets and on there being a bounded number of possible utility types; in contrast our result is valid in the standard model that allows a continuum of utility types and it is valid for any profile of budgets.

IV Fairness under the Priority Constraint

We now discuss two complementary concepts of fairness: ex-ante stability of Kesten & Ünver (2015), which captures the lack-of-justified-envy aspects of fairness, and the equal-claim property of He, Li & Yan (2015), extended herein to the settings with priorities, which captures the equality aspects of fairness. Our main result shows that, given a priority structure, the set of equal-budget PM assignments is precisely the set of assignments that are ex-ante stable and satisfy equal claim.

A Ex-Ante Stability

A key property of PM assignments is the ex-ante stability introduced by Kesten & Ünver (2015). An assignment is **ex-ante stable** if it does not cause ex-ante justified envy. An assignment Π causes **ex-ante justified envy** of $i \in \mathcal{I}$ toward $j \in \mathcal{I} \setminus \{i\}$ if $\exists s, s' \in \mathcal{S}$ such that $v_{i,s} > v_{i,s'}$, $k_{s,i} < k_{s,j}$, $\pi_{j,s} > 0$, and $\pi_{i,s'} > 0$. In other words, agent i who has higher-priority at s than another agent j has ex-ante justified envy towards j if j has positive probability of obtaining object s, while with positive probability i obtains an object less preferable than s. If an assignment causes ex-ante justified envy, then its every implementation with positive probability generates deterministic assignments that are not justified-envy-free, or not stable, in the sense of (Abdulkadiroglu & Sonmez 2003). This is an important consideration as many school districts insist on avoiding justified envy, for example, NYC (Atila Abdulkadiroglu, Parag A. Pathak & Alvin E. Roth 2005) and Boston (Atila Abdulkadiroglu, Parag A. Pathak, Alvin E. Roth & Tayfun Sonmez 2005).

In defining the PM mechanism, we require that prices are zero above the cut-off priority group and infinity below the cut-off. This restriction is both sufficient and necessary for the ex-ante stability of PM. To see this necessity and sufficiency, we relax the PM construction by allowing the prices to be agent-specific so that the matrix of all prices is $[p_{i,s}]_{i \in \mathcal{I}, s \in \mathcal{S}} \in$

$[0,+\infty]^S$.[23] With personalized prices, without loss of generality, we can normalize each agent's possibly-unequal budget to be one.

We first restrict ourselves to the set of non-wasteful assignments. Given an economy Γ, and a given personalized price vector $P_i = [p_{i,s}]_{s\in\mathcal{S}} \in [0,+\infty]^S$, where $p_{i,s}$ is the price of object s for agent i, we construct the demand correspondence of agent i, $\pi_i^*(v_i, P_i)$, that maximizes $\sum_{s\in\mathcal{S}} \pi_{i,s} v_{i,s}$ subject to $\sum_{s\in\mathcal{S}} p_{i,s}\pi_{i,s} \le 1$ among feasible π_i such that $\sum_{s\in\mathcal{S}} \pi_{i,s} = 1$ and $\pi_{i,s} \ge 0$ for all objects s. The set of possible personalized prices $\mathcal{P}_\Gamma$ is the set of all possible personalized prices that can rationalize some assignment as a result of agents' utility maximization (given budgets). That is:

$$\mathcal{P}_\Gamma \equiv \left\{ P^* = [P_i^*]_{i\in\mathcal{I}} \in [0,+\infty]^{I\times S} \,|\exists \pi_i \in \pi^*(v_i, P_i^*),\ \sum_{i\in\mathcal{I}} \pi_{i,s}^* \le q_s,\ \forall i\in\mathcal{I},\ \forall s\in\mathcal{S} \right\}.$$

The set of associated assignments is $\Pi_\Gamma(P^*) \equiv \left\{[\pi_i]_{i\in\mathcal{I}} \in \mathcal{A} | \pi_i \in \pi^*(v_i, P_i^*),\ \forall i\in\mathcal{I}\right\}$ for $P^* \in \mathcal{P}_\Gamma$. Finally, $\Pi_\Gamma \equiv \cup_{P^*\in\mathcal{P}_\Gamma}\Pi_\Gamma(P^*)$ is the set of all possible assignments that can be supported as a result of agents' utility maximization. Every feasible assignment can be represented in this way, i.e., $\Pi_\Gamma = \mathcal{A}$ (see He, Li & Yan (2015) for details).

Because $\sum_{s\in\mathcal{S}} q_s = I$, the definition of the PM mechanism restricts its prices to be in the following set:

$$\mathcal{P}_\Gamma^{Stable} \equiv \left\{ P^* \in \mathcal{P}_\Gamma | \forall s, \forall \pi \in \Pi_\Gamma(P^*), \exists k', p_{i,s}^* = \begin{cases} 0 & \text{if} \sum\limits_{j\in\mathcal{I} \text{ s.t. } k_{s,j}\le k'} \pi_{j,s} < q_s \ \&\ k_{s,i} < k' \\ +\infty & \text{if} \sum\limits_{j\in\mathcal{I}, \text{ s.t. } k_{s,j}\le k'} \pi_{j,s} = q_s \ \&\ k_{s,i} > k' \end{cases} \right\}.$$

By Theorem 1, $\mathcal{P}_\Gamma^{Stable} \neq \varnothing$, and thus the set of assignments $\Pi_\Gamma^{Stable} \equiv \cup_{P^*\in\mathcal{P}_\Gamma^{Stable}}\Pi_\Gamma(P^*)$ is also non-empty. Furthermore, Π_Γ^{Stable} corresponds exactly to the set of ex-ante stable assignments.

Proposition 1 Π_Γ^{Stable} *is the set of all non-wasteful ex-ante stable assignments.*

The proof is provided in the appendix. Given the construction of the PM mechanism, we then obtain the following.

Corollary 1 *Every PM assignment is ex-ante stable.*

[23] In addition to earlier drafts of our paper, personalized prices were studied for instance in He, Li & Yan (2015).

While we normalize the budgets to be equal, a simple re-scaling of personalized prices shows that the above result is also true for the PM mechanism with unequal budgets.

The above construction can be naturally extended to possibly wasteful assignments in which $\sum_{i\in\mathcal{I}} \pi_{i,s} = q'_s \leq q_s$ for all schools. Indeed, our analysis goes through if we modify the "market-clearing" conditions in the PM construction and substitute q'_s in lieu of the actual capacities q_s.

B Equal claim

The PM mechanism satisfies the strong fairness criterion of equal claim, introduced by He, Li & Yan (2015) in a setting without priorities. This fairness criterion captures the idea that the mechanism treats agents in the same priority class in the same way.[24] An ex-ante stable assignment Π satisfies equal claim if Π is an expected-utility-maximization outcome and in this maximization all agents in any given priority group of s face the same price of s if budgets are equal, or the same (equal) ratio of price to their budgets when budgets are unequal.

Definition 1 *An ex-ante stable assignment* Π *satisfies equal claim if and only if, given equal budgets, there exists* $P^* \in \mathcal{P}_\Gamma^{Stable}$ *such that* $\Pi \in \Pi_\Gamma(P^*)$ *and that for any* s*,* $p^*_{i,s} = p^*_{j,s}$ *whenever* $k_{s,i} = k_{s,j}$*.*

This definition allows Π to be wasteful. Nonetheless, the definition of the PM mechanism, Theorem 1, and Corollary 1 together imply that in any economy Γ there exists a non-wasteful ex-ante stable assignment satisfying equal claim.

C Characterization

Our main result on fairness says that our PM mechanism is characterized by the above two fairness criteria.

Theorem 3 *Given an economy* Γ *and equal budgets, the set of PM assignments is equivalent to the set of non-wasteful assignments satisfying both ex-ante stability and equal claim.*

Using this result, we can furthermore prove that in the special case in which agents' preferences and objects' priorities are both strict, PM assignments are deterministic and stable.

[24]See He, Li & Yan (2015) for further discussion of this property.

Theorem 4 *In an economy Γ, if both agents and objects rank the other side strictly, every PM assignment is deterministic, and the set of PM assignments is equivalent to the set of stable deterministic assignments.*

V Efficiency

Our welfare analysis starts with investigating the welfare of both agents and suppliers of objects, where the latter's preferences are assumed to be represented in the priority structure. In section B, we turn to the agent welfare, which is the focus of the literature.

A "Welfare" of Both Sides

As objects' priority ranking over agents do not necessarily reflect any underlying preferences of their suppliers, it is natural to care only about the welfare of agents. However, there are exceptions, and needless to say, priorities are usually not randomly chosen. For example, in school choice, priority rules may reflect preferences of the local constituency such as minimizing transportation costs (distance-based priorities) and/or encouraging investment in studying (test-score-based priorities).

When one is interested in taking the welfare of both sides into account, PM assignment is two-sided Pareto efficient in terms of both agent preferences and object suppliers' preferences defined by priorities. In this context, we say that an assignment $\Pi' \in \mathcal{A}$ is **ex-ante two-side dominated** by another assignment $\Pi \in \mathcal{A}$ if:

$$\begin{aligned} \sum_{s \in \mathcal{S}} \pi_{i,s} v_{i,s} &\geq \sum_{s \in \mathcal{S}} \pi'_{i,s} v_{i,s}, \forall i \in \mathcal{I}, \\ \sum_{i \in \{k_{s,i} \leq k\}} \pi_{i,s} &\geq \sum_{i \in \{k_{s,i} \leq k\}} \pi'_{i,s}, \ \forall s \in \mathcal{S}, \forall k \in \mathcal{K}, \end{aligned}$$

and at least one inequality is strict. That is, every agent has a weakly higher expected utility in Π, and, for each object s, the assignment Π first-order stochastically dominates Π' with respect to the priority structure. An assignment is **ex-ante two-sided efficient** if it is not ex-ante two-side dominated by any other assignment. We then obtain:

Theorem 5 *Every PM assignment is ex-ante two-sided efficient.*

If the problem is indeed two-sided, i.e., objects' priorities represent some underlying possibly-weak preferences, our results then make the PM mechanism a promising candidate for two-sided matching with weak preferences.

B Welfare of Agents

Our characterization result (Theorem 3) implies that no mechanism that is ex-ante stable and fair in the sense of equal claim can dominate the PM mechanism in efficiency terms. We now illustrate via examples how the PM mechanism can dominate other mechanisms.[25] The first example (subsection B.1) shows it outperforms the best possible ordinal mechanism even if the latter ignores priority constraints. Subsection B.2 compares the PM with the DA and again shows the efficiency advantage of the PM mechanism. In subsection B.3, we focus on the Boston mechanism, which is known to elicit signals of cardinal preferences from agents. Indeed, we show that the Boston mechanism can achieve the same PM assignment under some conditions (Proposition 2), but they otherwise differ.

An assignment $\Pi' \in \mathcal{A}$ is ex-ante **Pareto dominated for agents** by another assignment $\Pi \in \mathcal{A}$ if:

$$\sum_{s \in \mathcal{S}} \pi_{i,s} v_{i,s} \geq \sum_{s \in \mathcal{S}} \pi'_{i,s} v_{i,s}, \forall i \in \mathcal{I},$$

and at least one inequality is strict. An assignment is *ex ante* **agent-efficient** if it is not Pareto dominated for agents by any other feasible assignment. The definition applies to both random and deterministic assignments, and every deterministic assignment in any decomposition of an ex-ante agent-efficient assignment is Pareto optimal for agents.

In general, the PM mechanism cannot achieve ex-ante agent-efficiency due to the priority structure, and Theorem 3 implies that any agent-efficient assignment satisfying ex-ante stability and equal claim is a PM assignment.

The unique feature of PM is that it elicits and uses cardinal preferences to make the assignment. This implies that the outcome has the potential to be more efficient than ordinal mechanisms. In a one-sided setting (i.e., no priorities), Abdulkadiroglu, Che & Yasuda (2011) show cardinal mechanisms can dominate ordinal ones, and, building on subsequent analysis by Pycia (2014), we extend this result to the setting with priorities.

The following definition is useful for the comparison.

[25] The fact that some of the mechanisms we study can be dominated is known, see Ergin & Sonmez (2006), Abdulkadiroglu, Che & Yasuda (2011), Troyan (2012), and Pycia (2014).

Definition 2 *An assignment Π^* is ordinally efficient if there does not exist $\Pi \neq \Pi^*$ such that:*

$$\sum\nolimits_{s' \ s.t. \ v_{i,s'} \geq v_{i,s}} \pi^*_{i,s'} \leq \sum\nolimits_{s' \ s.t. \ v_{i,s'} \geq v_{i,s}} \pi_{i,s'}, \forall s \in \mathcal{S}, i \in \mathcal{I},$$

where at least one inequality is strict. Π^* is symmetric ordinal efficient if furthermore $\pi^*_{i,s} = \pi^*_{j,s}$, $\forall s$, whenever i and j have the same ordinal preferences.

B.1 The Cost of Ordinality

The following example, based on Pycia (2014), illustrates the extent to which restricting ourselves to ordinal mechanisms may result in an efficiency loss.

Example 1 *Let us consider the following economy with four agents $(i_1, ..., i_4)$ and four objects $(s_1, ..., s_4)$ with one copy of each available:*

Cardinal Preferences

	Objects			
Agent	s_1	s_2	s_3	s_4
i_1	1	ϵ	ϵ^2	0
i_2	1	$1-\epsilon$	ϵ^2	0
i_3	1	ϵ^2	$1-\epsilon$	0
i_4	1	ϵ^2	ϵ	0

$0 < \epsilon < 0.5$

Priority Structure

	Objects			
Agent	s_1	s_2	s_3	s_4
i_1	1	2	2	1
i_2	1	2	1	1
i_3	1	1	2	1
i_4	1	2	2	2

Smaller number means higher priority.

Note that no pair of agents has the same priorities at all objects. The following prices and assignment is an equilibrium outcome of the PM mechanism:

PM Priority-Specific Prices

	Objects			
Agent	s_1	s_2	s_3	s_4
i_1	2	1	1	0
i_2	2	1	0	0
i_3	2	0	1	0
i_4	2	1	1	0

PM Assignment

	Objects				*Expected*
Agent	s_1	s_2	s_3	s_4	*Utility*
i_1	1/2	0	0	1/2	1/2
i_2	0	1	0	0	$1-\epsilon$
i_3	0	0	1	0	$1-\epsilon$
i_4	1/2	0	0	1/2	1/2

Like in Pycia (2014), we can replicate this example and compare the PM mechanism with the "best" *ordinal mechanisms ignoring the priority constraint because Liu & Pycia (2012) showed that in large economies, all regular, asymptotically strategy-proof, asymptotically symmetric, and asymptotically efficient ordinal mechanisms deliver outcomes asymptotically equivalent to the symmetric ordinal efficient assignments.*

PS: Symmetric Ordinally Efficient Assignment

	Objects				Expected
Agent	s_1	s_2	s_3	s_4	*Utility*
i_1	1/4	1/2	0	1/4	$(1+2\epsilon)/4$
i_2	1/4	1/2	0	1/4	$(3-2\epsilon)/4$
i_3	1/4	0	1/2	1/4	$(3-2\epsilon)/4$
i_4	1/4	0	1/2	1/4	$(1+2\epsilon)/4$

Can be achieved by the Probabilistic Serial.

The above assignment can be implemented by the Probabilistic Serial (PS) whose definition is in Appendix A.

***Conclusion**: Given $\epsilon \in (0, 0.5)$, the PS assignment is Pareto dominated by the above PM assignment in terms of agent welfare, despite the fact that the PS assignment ignores priorities.*[26] *The PM assignment delivers an total welfare 0 to 50 percent higher.*

B.2 Comparison with the Gale-Shapley Deferred-Acceptance Mechanism

The Gale-Shapley Deferred-Acceptance (DA) mechanism, whose definition is also available in Appendix A, is a mechanism that has attracted the most attention both in the literature as well as in practice. When it is implemented in settings where priorities are coarse/weak, some tie-breaking rule is needed. For example, following reforms in NYC and Boston, the school choice program uniformly randomly chooses a single tie-breaking order for equal-priority students at each school and then employs the student-proposing DA using the modified priority structure.

From the perspective of tie-breaking, one may view the PM mechanism as a version of the DA mechanism with coarse priorities on one side. The unique feature of the PM mechanism is that the ties are broken endogenously according to cardinal preferences. The following example shows a case where the PM dominates the DA.

[26] For PS that takes priorities into account, see Mustafa Oğuz Afacan (2015). Such extensions of PS yield assignments that are dominated by those from standard PS.

Example 2 *In the same setting as in Example 1, the assignment from the DA with single tie-breaking (DA-STB) is as follows:*

The DA-STB Assignment

	Objects				*Expected*
Agent	s_1	s_2	s_3	s_4	*Utility*
i_1	1/4	1/6	1/12	1/2	$(3 + 2\epsilon + \epsilon^2)/12$
i_2	1/4	7/24	11/24	0	$(13 - 7\epsilon + 11\epsilon^2)/24$
i_3	1/4	11/24	7/24	0	$(13 - 7\epsilon + 11\epsilon^2)/24$
i_4	1/4	1/12	1/6	1/2	$(3 + 2\epsilon + \epsilon^2)/12$

DA-STB: The DA mechanism with single tie-breaking.

Conclusion*: The DA-STB assignment is Pareto dominated by the PM assignment in terms of agents' expected utility for $\epsilon \in (0, 0.5)$; the latter has a total welfare that is 16 to 89 percent higher.*

It should be noted that Kesten & Ünver (2015) extend the DA mechanism and propose two variants to deal with the tie-breaking on the object side. Since their mechanisms still rely on ordinal preferences of agents, the cost of ordinality (Section B.1) still applies. Empirically, Abdulkadiroglu, Agarwal & Pathak (2015) use data from the NYC high school match to show that possible improvements upon the DA mechanism from various ordinal mechanisms are rather limited. In fact, the best outcomes that the mechanisms in Kesten & Ünver (2015) can achieve are constrained ordinal efficiency, which are necessarily dominated by ordinal efficient outcomes.[27]

We also note a special case in which agents' preferences and objects' priorities are both strict. In this case, it must be that $\overline{k} = I$ and that there is exactly one agent in each priority group of any object. Noting that any DA assignment when agents report true ordinal preferences is stable, we have the following result as a corollary of Theorem 4:

Corollary 2 *If both agents and objects rank those on the other side strictly, any DA assignment when agents report true ordinal preferences is a PM assignment.*

B.3 Comparison with the Boston Mechanism

The PM mechanism is closely related to another commonly used mechanism, the Boston mechanism, whose definition is available in Appendix A. It has been noted in the literature that the

[27] In recent work, Che & Tercieux (2014) provide modifications of DA to improve asymptotic efficiency.

Boston mechanism elicits signals of agents' cardinal preferences, and indeed sometimes the Boston mechanism can yield PM assignments in Nash equilibrium.

Proposition 2 *A PM assignment is also a Bayesian Nash equilibrium assignment of the Boston mechanism, if every agent has strict preferences and consumes a bundle that either includes only free objects (according to her own prices), or includes one object with a positive price in* $(1, +\infty)$ *(according to her own price) and all others free to all agents.*

Note that the above result is a sufficient condition, and the following example shows there are other cases where the PM and Boston coincide.

Example 3 ***(The Boston Coincides with the PM Mechanism)*** *In the same setting as in Example 1, one can verify that the following strategies constitute a Nash equilibrium under the Boston mechanism (with single tie-breaking), and the equilibrium outcome is exactly the PM assignment.*

BM Equilibrium Strategies

Agent	*Rank-Order List*			
i_1	s_1	s_4	$\dots$	$\dots$
i_2	s_2	$\dots$	$\dots$	$\dots$
i_3	s_3	$\dots$	$\dots$	$\dots$
i_4	s_1	s_4	$\dots$	$\dots$

"..." indicates an arbitrary school.

BM Equilibrium Assignment

	Objects				*Expected*
Agent	s_1	s_2	s_3	s_4	*Utility*
i_1	1/2	0	0	1/2	1/2
i_2	0	1	0	0	$1-\epsilon$
i_3	0	0	1	0	$1-\epsilon$
i_4	1/2	0	0	1/2	1/2

When agents do not have strict preferences, or at least one of them spends her budget on more than one object with positive and finite prices, in general, a PM assignment is not an equilibrium outcome of the Boston mechanism. More importantly, in addition to the Boston mechainism's disadvantages discussed in the literature review, not every equilibrium assignment of the Boston is a PM assignment. The following example show this clearly.

Example 4 ***(The Boston Differs from the PM Mechanism)*** *Let us consider the following economy with three agents* $(i_1, ..., i_3)$ *and three objects* $(s_1, ..., s_3)$ *with one copy of each available.*

Moreover, there are no priorities. The unique PM price matrix and assignment are as follows:

Cardinal Preferences

	Objects		
Agent	s_1	s_2	s_3
i_1	1	0.9	0
i_2	1	0.6	0
i_3	1	0.1	0

PM Prices

	Objects		
Agent	s_1	s_2	s_3
i_1	15/8	9/8	0
i_2	15/8	9/8	0
i_3	15/8	9/8	0

PM Assignment

	Objects		
Agent	s_1	s_2	s_3
i_1	0	8/9	1/9
i_2	7/15	1/9	19/45
i_3	8/15	0	7/15

Note that in the PM assignment, i_2 purchases a positive probability share of both s_1 and s_2. Moreover, the Nash equilibrium of the Boston mechanism, which is unique in terms of outcomes, is that i_1 top ranks s_2, while i_2 and i_3 top ranking s_1, leading to an assignment different from the PM assignment:

BM Equilibrium Strategies

Agent	*Rank-Order List*		
i_1	s_2	...	...
i_2	s_1	s_3	...
i_3	s_1	s_3	...

"..." indicates an arbitrary school.

BM Assignment

	Objects		
Agent	s_1	s_2	s_3
i_1	0	1	0
i_2	1/2	0	1/2
i_3	1/2	0	1/2

VI Concluding Remarks

This paper studies the allocation of indivisible goods based on priorities when monetary transfers are not possible and agents have unit demand. We propose a pseudo-market (PM) mechanism, which elicits agents' cardinal preferences and delivers an assignment as bundles of probability shares in objects. When doing so, the PM mechanism internally finds a Walrasian equilibrium in which agents are endowed with budgets of token money and purchase bundles to maximize their expected utility. The prices in the Walrasian equilibrium depend on agents' cardinal preferences and are priority-specific. More specifically, everyone in any given priority group of an object faces the same price, while those who are in higher priority groups of an object face a lower, sometimes zero, price of that object.

The mechanism has desirable properties. After showing the mechanism is well-defined, we prove that it is asymptotically incentive compatible for agents to report cardinal preferences in a sequence of replica economies. Moreover, the mechanism delivers an assignment, which can be random or deterministic, that satisfies ex-ante stability or eliminates ex-ante justified envy. The structure of PM prices also guarantees that everyone in the same priority group of an object has an equal claim to that object, whenever budgets are equal. The mechanism can deliver all assignments that are not dominated by any assignment satisfying the above criteria. Because of the explicit use of cardinal preferences, the PM mechanism has an efficiency advantage over other popular mechanisms.

These properties of the mechanism make it a promising candidate for real-life applications to settings such as school choice. Schools often prioritize student applications, and the priority structure is usually determined by the school district or local laws. In most cases, a school's priority ranking over students is not strict, which makes the PM a natural candidate to run seat allocation. The mechanism guarantees that the resulting assignment is ex-ante stable and thus that it can be implemented as a lottery over deterministic assignments that are stable. Furthermore, as Abdulkadiroglu, Che & Yasuda (2011) point out, in settings such as school choice, students may have similar ordinal preferences. Therefore, without information on cardinal preferences, the efficiency that a mechanism can achieve may be limited.[28] Indeed, using data from the high school match in NYC, Abdulkadiroglu, Agarwal & Pathak (2015) show the potentials of eliciting cardinal utilities in improving student welfare. By explicitly using students' cardinal preferences, the PM mechanism allows school districts to achieve such efficiency gains.[29]

The major concern with implementing the PM mechanism is the difficulty of eliciting cardinal preferences from agents. For instance, Bogomolnaia & Moulin (2001) argue that agents participating in the allocation problem may have limited rationality/information and thus do not know exactly their cardinal preferences. However, the evidence in Eric Budish & Judd Kessler (2014) from an experiment of allocating course schedules to students via a pseudo-market mechanism shows that the difficulty in reporting cardinal preferences does not prevent the mechanism from outperforming its alternatives on multiple dimensions. Besides, more training and more time to acquire information on cardinal preferences for agents may also lower this difficulty.

From a different perspective, one may consider the requirement of reporting cardinal preferences as an incentive for agents to investigate whether an object is a good fit for her. Such information

[28]Whereas ordinal inefficiency may vanish in large markets (Yeon-Koo Che & Fuhito Kojima 2010), the cardinal inefficiency of ordinal mechanisms persists (Pycia 2014).

[29]Note that one can accommodate group-specific quotas within the PM similarly to how they might be accommodated within the DA mechanism (see e.g., Abdulkadiroglu & Sonmez (2003)): to accommodate such quotas, one can divide each school into multiple sub-schools each of which has a quota equal to the one for the corresponding group and gives that group the highest priority.

acquisition can even be welfare-improving (Yan Chen & Yinghua He 2015).

Another concern is that agents, especially children and their parents in educational settings, may not be able to play the preference-revelation game optimally, even when telling the truth is always optimal. For example, Alex Rees-Jones (2016) uses survey data and reports that 5.38 percent of the agents do not report true ordinal preferences under the DA mechanism, which is a strategy-proof ordinal mechanism; Avinash Hassidim, Assaf Romm & Ran Shorrer (2016) report an even higher rate, 19 percent, using data on agents' behavior. However, misreporting behavior may not affect the outcome or the assignment, because agents tend to omit objects that are unlikely to be obtained by them (Gabrielle Fack, Julien Grenet & Yinghua He 2015, Georgy Artemov, Yeon-Koo Che & Yinghua He 2017). Indeed, as documented in Hassidim, Romm & Shorrer (2016), at most 1.4 percent of the agents misreport and end up with sub-optimal outcomes; even lower rates are reported in Artemov, Che & He (2017). Nonetheless, as an incentive compatible mechanism, the PM mechanism allows market designer to find ways to convince agents to report their true preferences.

References

Abdulkadiroglu, Atila, and Tayfun Sonmez. 2003. “School Choice: A Mechanism Design Approach.” *American Economic Review*, 93(3): 729–747.

Abdulkadiroglu, Atila, Nikhil Agarwal, and Parag A. Pathak. 2015. “The Welfare Effects of Coordinated Assignment: Evidence from the NYC HS Match.” *National Bureau of Economic Research Working Paper Series*, No. 21046.

Abdulkadiroglu, Atila, Parag A. Pathak, Alvin E. Roth, and Tayfun Sonmez. 2005. “The Boston Public School Match.” *The American Economic Review, Papers and Procedings*, 95(2): 368–371.

Abdulkadiroglu, Atila, Parag A. Pathak, Alvin E. Roth, and Tayfun Sonmez. 2006. “Changing the Boston School Choice Mechanism: Strategy-proofness as Equal Access.” *Mimeo.*

Abdulkadiroglu, Atila, Parag A. Pathak, and Alvin E. Roth. 2005. “The New York City High School Match.” *American Economic Review, Papers and Proceedings*, 95(2): 364–367.

Abdulkadiroglu, Atila, Parag A. Pathak, and Alvin E. Roth. 2009. “Strategy-proofness versus Efficiency in Matching with Indifferences: Redesigning the New York City High School Match.” *American Economic Review*, 99(5): 1954–78.

Abdulkadiroglu, Atila, Yeon-Koo Che, and Yosuke Yasuda. 2011. “Resolving Conflicting Preferences in School Choice: the Boston Mechanism Reconsidered.” *American Economic Review*, 101(1): 399–410.

Abdulkadiroglu, Atila, Yeon-Koo Che, and Yosuke Yasuda. 2015. “Expanding “Choice” in School Choice.” *American Economic Journal: Microeconomics*, 7(1): 1–42.

Afacan, Mustafa Oğuz. 2015. “The Object Allocation Problem with Random Priorities.” *mimeo.*

Artemov, Georgy, Yeon-Koo Che, and Yinghua He. 2017. “Strategic ‘Mistakes’: Implications for Market Design Research.” *mimeo.*

Ashlagi, Itai, and Peng Shi. 2016. “Optimal allocation without money: an engineering approach.” *Management Science*, 62(4): 1078–1097.

Ashlagi, Itai, Yash Kanoria, and Jacob D. Leshno. 2014. “Unbalanced random matching markets: The stark effect of competition.” *Manuscript, MIT.*

Azevedo, Eduardo, and Eric Budish. 2012. “Strategyproofness in the Large as a Desideratum for Market Design.” *mimeo.*

Birkhoff, Garrett. 1946. “Three Observations on Linear Algebra.” *Universidada Nacional de Tucuman, Revista*, Serie A(vol. 5): 147–151.

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

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

Budish, Eric, and Judd Kessler. 2014. "Changing the Course Allocation Mechanism at Wharton." *Chicago Booth Research Paper*, , (15-08).

Budish, Eric, Yeon-Koo Che, Fuhito Kojima, and Paul Milgrom. 2013. "Designing Random Allocation Mechanisms: Theory and Applications." *American Economic Review*, 103(2): 585–623.

Chen, Yan, and Tayfun Sonmez. 2006. "School Choice: An Experimental Study." *Journal of Economic Theory*, 127(1): 202–231.

Chen, Yan, and Yinghua He. 2015. "Information Acquisition and Provision in School Choice." *mimeo.*

Che, Yeon-Koo, and Fuhito Kojima. 2010. "Asymptotic Equivalence of Probabilistic Serial and Random Priority Mechanisms." *Econometrica*, 78(5): 1625–1672.

Che, Yeon-Koo, and Olivier Tercieux. 2014. "Efficiency and stability in large matching markets." *Manuscript.*

Dierker, Egbert. 1974. *Topological methods in Walrasian economics. Lecture Notes in Economics and Mathematical Systems*, Springer Berlin Heidelberg.

Erdil, Aytek, and Haluk Ergin. 2006. "Two-Sided Matching with Indifferences." *Mimeo.*

Erdil, Aytek, and Haluk Ergin. 2008. "What's the Matter with Tie-Breaking? Improving Efficiency in School Choice." *American Economic Review*, 98(3): 669–89.

Ergin, Haluk, and Tayfun Sonmez. 2006. "Games of School Choice under the Boston Mechanism." *Journal of Public Economics*, 90(1-2): 215–237.

Fack, Gabrielle, Julien Grenet, and Yinghua He. 2015. "Beyond Truth-Telling: Preference Estimation with Centralized School Choice." C.E.P.R. Discussion Papers CEPR Discussion Papers 10907.

Featherstone, Clayton, and Muriel Niederle. 2008. "Ex Ante Efficiency in School Choice Mechanisms: An Experimental Investigation." *NBER Working Paper Series*, No. 14618.

Hafalir, Isa Emin, and Antonio Miralles. 2014. "Welfare-Maximizing Assignment of Agents to Hierarchical Positions." *Manuscript.*

Hashimoto, Tadashi. 2013. "The generalized random priority mechanism with budgets." *Manuscript.*

Hassidim, Avinash, Assaf Romm, and Ran Shorrer. 2016. ""Strategic" Behavior in a Strategy-proof Environment." mimeo.

He, Yinghua. 2011. *Essays on public policies. PhD Dissertation*, Columbia University.

He, Yinghua. 2012. "Gaming the Boston School Choice Mechanism in Beijing." *Toulouse School of Economics Working Papers*, 12-345. Toulouse School of Economics.

He, Yinghua, and Jianye Yan. 2012. "Competitive Equilibrium from Equal Incomes for Two-Sided Matching." Toulouse School of Economics (TSE) TSE Working Papers 12-344.

He, Yinghua, Antonio Miralles, and Jianye Yan. 2012. "Competitive Equilibrium from Equal Incomes for Two-Sided Matching." *mimeo.*

He, Yinghua, Sanxi Li, and Jianye Yan. 2015. "Evaluating assignment without transfers: A market perspective." *Economics Letters*, 133(C): 40–44.

Hildenbrand, Werner. 1974. *Core and equilibria of a large economy.* Princeton, NJ: Princeton University Press.

Hylland, Aanund, and Richard Zeckhauser. 1979. "The Efficient Allocation of Individuals to Positions." *Journal of Political Economy*, 87(2): 293–314.

Jackson, Matthew O. 1992. "Incentive Compatibility and Competitive Allocations." *Economics Letters*, 40(3): 299–302.

Kesten, Onur. 2010. "School Choice with Consent." *Quarterly Journal of Economics*, 125(3): 1297–1348.

Kesten, Onur, and M. Utku Ünver. 2015. "A Theory of School-Choice Lotteries." *Theoretical Economics*, 10(2): 543–595.

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

Kojima, Fuhito, and Parag A. Pathak. 2009. "Incentives and Stability in Large Two-Sided Matching Markets." *American Economic Review*, 99(3): 608–27.

Kojima, Fuhito, Parag A. Pathak, and Alvin E. Roth. 2013. "Matching with Couples: Stability and Incentives in Large Markets." *The Quarterly Journal of Economics*, 128(4): 1585–1632.

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

Lee, SangMok. 2014. "Incentive Compatibility of Large Centralized Matching Markets." *Manuscript, University of Pennsylvania.*

Lee, SangMok, and Leeat Yariv. 2014. "On the efficiency of stable matchings in large markets." *Manuscript.*

Liu, Qingmin, and Marek Pycia. 2012. "Ordinal Efficiency, Fairness, and Incentives in Large Markets." *mimeo.*

Miralles, Antonio. 2008. "School Choice: The Case for the Boston Mechanism." *Mimeo.*

Miralles, Antonio. 2011. "Pseudomarkets with Priorities in Large Random Assignment Economies." *Mimeo.* Working Papers.

Miralles, Antonio. 2012. "Cardinal Bayesian allocation mechanisms without transfers." *Journal of Economic Theory*, 147(1): 179–206.

Miralles, Antonio, and Marek Pycia. 2014. "Prices and Efficient Assignments Without Transfers." *Manuscript.*

Miralles, Antonio, and Marek Pycia. 2015. "Large vs. Continuum Assignment Economies." *Manuscript.*

Nguyen, Thanh, Ahmad Peivandi, and Rakesh Vohra. 2015. "Assignment Problems with Complementarities." *Manuscript.*

Pathak, Parag A., and Tayfun Sonmez. 2008. "Leveling the Playing Field: Sincere and Sophisticated Players in the Boston Mechanism." *American Economic Review*, 98(4): 1636–52.

Pycia, Marek. 2011. "Ordinal Efficiency, Fairness, and Incentives in Large Multi-Unit-Demand Assignments." *mimeo.*

Pycia, Marek. 2014. "The Cost of Ordinality." *mimeo.*

Rees-Jones, Alex. 2016. "Suboptimal Behavior in Strategy-Proof Mechanisms: Evidence from the Residency Match." mimeo.

Roberts, Donald John, and Andrew Postlewaite. 1976. "The Incentives for Price-Taking Behavior in Large Exchange Economies." *Econometrica*, 44(1): 115–27.

Roth, Alvin E. 1982. "The Economics of Matching: Stability and Incentives." *Mathematics of Operations Research*, 7(4): 617–628.

Troyan, Peter. 2012. "Comparing school choice mechanisms by interim and ex-ante welfare." *Games and Economic Behavior*, 75(2): 936–947.

von Neumann, John. 1953. "A Certain Zero-Sum Two-Person Game Equivalent to the Optimal Assignment Problem." In *Contributions to the Theory of Games.* Vol. II, , ed. Harold William Kuhn and Albert William Tucker, 5–12. Princeton, N. J.:Princeton University Press.

Zhou, Lin. 1990. "On A Conjecture by Gale about One-Sided Matching Problems." *Journal of Economic Theory*, 52(1): 123–135.

Appendices

A Alternative Mechanisms

This appendix gives the definitions of three mechanisms: the Probabilistic Serial, the Boston mechanism (also known as the immediate-acceptance mechanism), and the Gale-Shapley deferred-acceptance mechanism.

The **Probabilistic Serial** is defined by the following symmetric simultaneous eating algorithm. It is proposed for one-sided matching where objects do not rank agents. Each object s is considered as an infinitely divisible object with supply q_s that agents eat in the time interval $[0, 1]$.

Round 1. Each agent eats away from her favorite object at the same unit speed, and the algorithm proceeds to the next step when an object is completely exhausted.

Generally, in:

Round k (k>1). Each agent eats away from her most-preferred object among the remaining ones at the same unit speed, and the algorithm proceeds to the next step when an object is completely exhausted.

The process terminates after any round k when every agent has eaten exactly one total unit of objects (i.e., at time 1). The assignment of an agent i is then given by the amount of each object she has eaten during the run of the algorithm.

The **Boston mechanism** solicits rank-ordered lists of objects from agents, uses pre-defined rules, including tie-breaking rules, to determine objects' strict ranking over agents, and has multiple rounds:

Round 1. Each object considers all the agents who rank it first and assigns its copies in order of their priority at that object until either there are no copies of the object left or no such agents left.

Generally, in:

Round (k>1). The k-th choice of the agents who have not yet been assigned is considered. Each object that still has available copies assigns the remaining copies to agents who rank it as kth choice in order of their priority at that object until either there are no copies of that object left or no such agent left.

The process terminates after any round k when every agent is assigned a copy of some object, or if the only agents who remain unassigned listed no more than k choices.

The **Gale-Shapley Deferred-Acceptance (DA) mechanism** can be agent-proposing or object-proposing. In the former, the mechanism collects objects' supplies and their priority structure over agents, as well as agents' submitted rank-ordered lists of objects. When necessary, tie-breaking rules are applied to form strict rankings of objects over agents. The process then has several rounds:

Round 1. Every agent applies to her first choice. Each object rejects the least preferred agents in excess of its supply and *temporarily holds* the others.

Generally, in:

Round (k>1). Every agent who is rejected in Round $(k-1)$ applies to the next choice on her list. Each object pools new applicants and those who are held from Round $(k-1)$ together and rejects the least preferred agents in excess of its supply. Those who are not rejected are *temporarily held* by the objects.

The process terminates after any Round k when no rejections are issued. Each object is then matched with agents it is currently holding. The object-proposing DA mechanism is similarly defined.

B Proofs

A Proof of Theorem 1

First, we make the price space compact by transforming $[0,+\infty]^{S\times\overline{k}}$ to $\mathcal{Z}\equiv[0,\pi/2]^{S\times\overline{k}}$ such that $\forall P\in[0,+\infty]^{S\times\overline{k}}$, there is a $Z\in\mathcal{Z}$ and $Z=[z_{s,k}]_{s\in\mathcal{S},k\in\mathcal{K}}=[\arctan(p_{s,k})]_{s\in\mathcal{S},k\in\mathcal{K}}$, with $\arctan(+\infty)\equiv\pi/2$ and $\tan(\pi/2)\equiv+\infty$.[30] Since the arctangent function, arctan, is a positive monotonic transformation, the reverse statement is also true such that $\forall Z\in\mathcal{Z}$, there is a $P\in[0,+\infty]^{S\times\overline{k}}$ and $P=TAN(Z)\equiv[\tan(z_{s,k})]_{s\in\mathcal{S},k\in\mathcal{K}}$.

For the purposes of this proof, we also augment our economy with an "artificial outside option" with infinite supply, s_0, at which every agent has the same priority and necessarily face a zero price. For every agent, her valuation of that this object is $v_{s_0}<\min_{i\in\mathcal{I}}\min_{s\in\mathcal{S}}v_{i,s}$, which guarantees that the demand for s_0 is zero in equilibrium.[31] A price-adjustment process for our extended economy $\overline{\Gamma}$ is defined as,

$$
\begin{aligned}
&H\left[Z,G\left(\mathcal{TAN}(Z),u\right)\right]\\
\equiv\;&\left\{Y=[y_{s,k}]_{s\in\mathcal{S},k\in\mathcal{K}}\;\middle|\;\begin{array}{l}y_{s,k}\left([d_{s,k}]_{k\in\mathcal{K}}\right)=\min\left\{\frac{\pi}{2},\max\left[0,z_{s,k}+\left(\sum_{\kappa=1}^{k}d_{s,\kappa}-\frac{q_s}{I}\right)\right]\right\}\\ \forall\,[d_{s,k}]_{s\in\mathcal{S},k\in\mathcal{K}}\in G\left(\mathcal{TAN}(Z),u\right)\end{array}\right\},
\end{aligned}
$$

where $u=(u_1,\cdots,u_I)$ are agents' reported utility profile, and $G\left(\mathcal{TAN}(Z),u\right)$ is the per capita demand correspondence for each priority group of each object in the extended economy. Demands are well defined in the extended simplex Δ^{S+1} since object s_0 always has a zero price.

Since G is the average of individual demand correspondences, it is then upper hemicontinuous

[30] We could alternatively work in the original space $[0,+\infty]^{S\times\overline{k}}$ which is also compact. One may find the transformation to $[0,\pi]^{S\times\overline{k}}$ useful when solving PM equilibrium computationally. Here and in the following, with some abuse of notation, π, without subscript, is the mathematical constant, i.e., the ratio of a circle's circumference to its diameter.

[31] v_{s_0} could even be negative. As we impose the unit-demand constraint, everyone's total probability shares of all objects must be exactly one. Therefore, agents may demand an object of negative utility.

and convex-valued, and thus $H[Z,G]$ has the same properties because it is a continuous function. $H[Z,G]$ therefore satisfies all the conditions of Kakutani's fixed-point theorem, and there must exist a fixed point Z^* such that $Z^* \in H[Z^*, G(\mathcal{TAN}(Z^*), u)]$.

Note that not all fixed points, Z^*, lead to prices that satisfy the conditions of the PM mechanism. For example, the prices implied by Z^* may be higher for higher priority groups. More precisely, given Z^*, there may exist $[d_{s,k}]_{s\in\mathcal{S},k\in\mathcal{K}} \in G$ such that $\forall s$ and $\forall k$,

$$z^*_{s,k} = \min\left\{\frac{\pi}{2}, \max\left[0, z^*_{s,k} + \left(\sum_{\kappa=1}^{k} d_{s,\kappa} - \frac{q_s}{I}\right)\right]\right\}.$$

In other words, $\sum_{\kappa=1}^{\bar{k}} d_{s,\kappa} = q_s/I$, $\forall s$; there exists $k^*(s)$ for each s such that (a) $d_{s,k^*(s)} > 0$, $\sum_{\kappa=1}^{k^*(s)} d_{s,\kappa} = q_s/I$, and $z^*_{s,k^*(s)} \in \left[0, \frac{\pi}{2}\right)$; (b) if $k < k^*(s)$, $\sum_{\kappa=1}^{k} d_{s,\kappa} < q_s/I$ and $z^*_{s,k} = 0$; and (c) if $k > k^*(s)$, $d_{s,k} = 0$, and $z^*_{s,k}$ can be some value in $[0, \pi/2]$. The "indeterminacy" in (c) happens because a finite price can sometimes be high enough to deter consumption by some agents.

We therefore need to impose some selection rule. If $d_{s,k} = 0$, and $z^*_{s,k} \in [0, \pi/2)$ for some $k > k^*(s)$, there must exist another Z^{**} such that $Z^{**} \in H[Z^{**}, G(\mathcal{TAN}(Z^{**}), u)]$ and that if $k \le k^*(s)$, $z^{**}_{s,k} = z^*_{s,k}$; and that if $k > k^*(s)$, $z^{**}_{s,k} = \pi/2$. That is, the highest price is selected.

Now we show that an equilibrium contains no positive demand for the "artificial outside option" s_0. If that were the case, then there would be at least one school $s \in S$ which is in excess supply: $\sum_{\kappa=1}^{\bar{k}} d_{s,\kappa} < q_s/I$. This however implies that $z^*_{s,k} = 0$ for all k, that is, the price for s is zero for everyone. Consider an individual i who has purchased a positive probability share of s_0. Since $v_{i,s_0} < v_{i,s}$, agent i's bundle is not optimal, since substituting all probability shares in s_0 for those in s (at no cost) would strictly increase her expected utility. Therefore, P^{**} induces demands in the simplex of the original economy, Δ^S, and it is an equilibrium price matrix for the original economy.

In summary, $P^{**} = TAN(Z^{**})$ satisfies the conditions of prices defined in the PM mechanism and indeed clears the market. Therefore, a PM price matrix exists, which implies the existence of PM assignment.

B Proof of Theorem 2

Let us represent each economy by a probability measure. Let $T = [0,1]^S \times \mathcal{K}^S$ be the compact space of utility-priority profiles endowed with the standard Euclidean distance. For any profile $(v,k) \in T$ and scalar $\varepsilon > 0$, let $B_\varepsilon(v,k)$ be the ball of profiles within distance ε of (v,k). Let $\mathcal{M}$ be the space of compact-support Borel probability measures on T. An economy can be conveniently represented by a probability measure μ on T, where $\mu(v,k)$ is the proportion of agents with utility-priority profile (v,k) in the economy. Therefore, each of the sequence of replica economies can be represented by the same measure. We extend our use of the Prohorov metric ρ to measure the distance between measures on T,

$$\rho(\mu,\nu) = \inf\left\{\varepsilon > 0 | \nu(E) \le \mu(B_\varepsilon(E)) + \varepsilon \text{ and } \mu(E) \le \nu(B_\varepsilon(E)) + \varepsilon,\ \forall E \subset T\right\}.$$

Notice that the entire set of regular economies can be partitioned into open and disjoint subsets such that for every subset B there is a finite number $m > 0$ of continuous functions $\psi_1, ..., \psi_m$ from B to $[0, +\infty]^{S\times\bar{k}}$ such that the set of transformed PM price matrices $\Psi(\mu) = \{\psi_1(\mu), ..., \psi_m(\mu)\}$ for every $\mu \in B$. Indeed, consider an open ball of regular economies around each regular economy. Non-disjoint balls must have the same set of price functions. Taking a union of open sets with the same set of price functions gives us an open set with these price functions that is disjoint from regular economies with other price sets.

Let us set $\psi^{(n)}(\mu) = \psi(\mu) = \psi_1(\mu)$ for regular economies, and set both $\psi^{(n)}(\mu)$ and $\psi(\mu)$ to be an arbitrary price vector otherwise. By construction, this price function is continuous at every regular economy.

Take the n-replica regular economy $\Gamma^{(n)}$ that is represented by $\mu^{(n)}(=\mu)$. Suppose agent i submits a report u instead of v_i, and the resulting measure on utility profiles is $\mu_u^{(n)}$. By definition of the Prohorov metric, $\mu_u^{(n)}$ is close to $\mu^{(n)}$ in which everyone is truth-telling. For large enough n we have that $\mu_u^{(n)}$ is in the same price-function-ball as $\mu^{(n)} = \mu$. Since $\psi^{(n)}$ is continuous on each price-function ball, agent i can affect prices by only a small amount: given every $\varepsilon > 0$, for every n sufficiently large and for all u_i,

$$\left|\arctan\left(\psi^{(n)}\left(u_i, V_{-i}^{(n)}\right)\right) - \arctan\left(\psi^{(n)}\left(v_i, V_{-i}^{(n)}\right)\right)\right| < \varepsilon.$$

We therefore specify a price selection rule for PM; since agents' utilities are continuous in prices, Theorem 2 follows.

For completeness, we provide below a detailed analysis of the latter statement, including a useful technical lemma.

Let $\mathcal{P}_{u_i}^{(n)}$ denote the set of PM prices when one copy of i reports u_i while all others reporting truthfully ($V_{-i}^{(n)}$) in $\Gamma^{(n)}$. Therefore, $\mathcal{P}_{v_i}^{(n)}$ is the set of PM prices when everyone in $\Gamma^{(n)}$ is truth-telling, and $\cup_{u_i\in[0,1]^S}\mathcal{P}_{u_i}^{(n)}$ is the set of prices that i can obtain through unilateral manipulation of her reports. Furthermore, $\cup_{u_i\in[0,1]^S}\mathcal{P}_{u_i}^{(1)}$ is the set of obtainable PM prices associated with the base economy Γ. Similar to the lemma in Roberts & Postlewaite (1976), we have the following:

Lemma B1 *Given the sequence of replica economies, $\left\{\Gamma^{(n)}\right\}_{n\in\mathbb{N}}$, and a agent i,*

(i) $\cup_{u_i\in[0,1]^S}\mathcal{P}_{u_i}^{(n)}$ is closed for all n.

(ii) The sets of PM prices that i can obtain by unilateral manipulation in $\left\{\Gamma^{(n)}\right\}_{n\in\mathbb{N}}$ have a nesting structure: $\cup_{u_i\in[0,1]^S}\mathcal{P}_{u_i}^{(n)} \subseteq \cup_{u_i\in[0,1]^S}\mathcal{P}_{u_i}^{(n')}$ for all $n > n'$.

(iii) If $P \notin \mathcal{P}_{v_i}^{(1)}$, there exists n^* such that $n > n^*$ implies $P \notin \mathcal{P}_{u_i}^{(n)}$, and thus $\mathcal{P}_{v_i}^{(1)} = \cap_{n\in\mathbb{N}}\left(\cup_{u_i\in[0,1]^S}\mathcal{P}_{u_i}^{(n)}\right)$.

Proof of Lemma B1. We prove the lemma step by step.

(i) $\cup_{u_i\in[0,1]^S}\mathcal{P}_{u_i}^{(n)}$ is closed.

Consider a sequence of price matrices $P^{(m)} \to \bar{P}$, where $P^{(m)} \in \cup_{u_i\in[0,1]^S}\mathcal{P}_{u_i}^{(n)}$. That is, for each m, there is a sequence of $u_i^{(m)}$ such that $P^{(m)} \in \mathcal{P}_{u_i^{(m)}}^{(n)}$. Since $u_i^{(m)}$ is bounded, there must exist a convergent subsequence, which is also denoted as $u_i^{(m)} \to \bar{u}_i$. Besides, the corresponding subsequence of price matrices, still denoted as $P^{(m)}$, converges to $\bar{P}$. This implies:

$$\pi^{(m)}\left(u_i^{(m)}, P^{(m)}\right) + (n-1)\,\pi^{(m)}\left(v_i, P^{(m)}\right) + n\sum_{j\neq i}\pi^{(m)}\left(v_j, P^{(m)}\right) = nQ,$$

where $\pi^{(m)}\left(u_i, P^{(m)}\right)$ denotes an element in the set $\pi\left(u_i, P^{(m)}\right)$. Due to their boundedness, there is a subsequence of $\left\{\pi^{(m)}\left(u_i^{(m)}, P^{(m)}\right), \pi^{(m)}\left(v_i, P^{(m)}\right)\right\}$ that converges to $\{\bar{\pi}_{u_i}, \bar{\pi}_{v_i}\}$.

The maximum theorem implies that $\pi(u_i, P)$ is upper hemicontinuous in (u_i, P), and therefore $\bar{\pi}_{u_i} \in \pi\left(\bar{u}_i, \bar{P}\right)$, $\bar{\pi}_{v_i} \in \pi\left(v_i, \bar{P}\right)$, and $\bar{\pi}_{v_j} \in \pi\left(v_j, \bar{P}\right)$. The equality above leads to:

$$\bar{\pi}_{u_i} + (n-1)\,\bar{\pi}_{v_i} + n\sum_{j\neq i}\pi^{(m)}\bar{\pi}_{v_j} = nQ,$$

which proves that $\bar{P} \in \cup_{u_i\in[0,1]^S}\mathcal{P}_{u_i}^{(n)}$ and hence that $\cup_{u_i\in[0,1]^S}\mathcal{P}_{u_i}^{(n)}$ is closed.

(ii) The nesting structure of $\cup_{u_i\in[0,1]^S}\mathcal{P}_{u_i}^{(n)}$.

To simplify notations, in the following, let us assume that the demand correspondence $\pi(u_i, P)$ is single-valued for all i, all u_i, and all P. The proof can easily be extended to allow $\pi(u_i, P)$ to be set-valued.

$P \in \cup_{u_i\in[0,1]^S}\mathcal{P}_{u_i}^{(n)}$ means that there exists $u_i^{(n)}$ such that P clears the market given reports $\left(u_i^{(n)}, v_i\right)$:

$$\pi\left(u_i^{(n)}, P\right) + (n-1)\,\pi(v_i, P) + n\sum_{j\neq i}\pi(v_j, P) = nQ.$$

To have P as a PM price matrix in $\Gamma^{(n')}$, there has to exist some $u_i^{(n')} \in [0,1]^S$ such that:

$$\pi\left(u_i^{(n')}, P\right) + \left(n'-1\right)\pi(v_i, P) + n'\sum_{j\neq i}\pi(v_j, P) = n'Q.$$

Differencing the two equations and rearranging the terms lead to:

$$\pi\left(u_i^{(n')}, P\right) = \frac{n'}{n}\pi\left(u_i^{(n)}, P\right) + \frac{n-n'}{n}\pi(v_i, P).$$

Since $\pi\left(u_i^{(n)}, P\right)$ and $\pi(v_i, P)$ are affordable to i, the convex combination of the two must be affordable to i. Therefore, there must exist some $u_i^{(n')}$ such that the above equation is satisfied.

(iii) $\mathcal{P}_{v_i}^{(1)} = \cap_{n\in\mathbb{N}}\left(\cup_{u_i\in[0,1]^S}\mathcal{P}_{u_i}^{(n)}\right)$

It is straightforward to verify that $\mathcal{P}_{v_i}^{(1)} \subseteq \mathcal{P}_{v_i}^{(n)} \subseteq \cup_{u_i\in[0,1]^S}\mathcal{P}_{u_i}^{(n)}$ for all n. We then show that for any $P \notin \mathcal{P}_{v_i}^{(1)}$, there exists n^* such that $n > n^*$ implies $P \notin \mathcal{P}_{u_i}^{(n)}$.

Suppose that $P \notin \mathcal{P}_{v_i}^{(1)}$ but the statement in the lemma is false. The nesting structure implies that $P \in \cup_{u_i\in[0,1]^S}\mathcal{P}_{u_i}^{(n)}$, for all n. Therefore, there exists a sequence of reports by the given copy of agent i, $\left\{u_i^{(n)}\right\}_{n\in\mathbb{N}}$, such that the market clears at P:

$$\pi\left(u_i^{(n)}, P\right) + (n-1)\,\pi(v_i, P) + n\sum_{j\neq i}\pi(v_j, P) = nQ.$$

Rearranging the above equation yields:

$$\pi\left(u_i^{(n)}, P\right) - \pi(v_i, P) = n\left(Q - \sum_{j\in\mathcal{I}}\pi(v_j, P)\right),$$

where the left-hand-side term is bounded due to the unit demand constraint. Moreover, $P \notin \mathcal{P}_{v_i}^{(1)}$ implies $Q - \sum_{j\in\mathcal{I}}\pi(v_j, P) \neq 0$, which means the right-hand-side of the equation diverges when n increases. Therefore, there must exist $\bar{n}$ such that the above equation cannot be satisfied for $n > \bar{n}$. This contradiction proves the lemma. ■

We are now ready to finish the proof of our main incentive compatibility theorem.

Proof of Theorem 2. Suppose that for a copy of agent type i, also denoted as i, there exists a subsequence of replica economies $\left\{\Gamma^{(n_m)}\right\}_{n_m\in\mathbb{N}}$ where she gains at least ε by unilateral misreporting. Let $P^{(m)} = \psi^{(n_m)}\left(u_i^{(n_m)}, V_{-i}^{(n_m)}; K^{(n)}\right)$ where $P^{(m)}$ is the price matrix with which PM implements the assignment in economy $\Gamma^{(n_m)}$ after i's unilateral manipulation. Since $\left\{\arctan\left(P^{(m)}\right)\right\}_{n_m\in\mathbb{N}}$ is bounded, there is a subsequence (also denoted as $\left\{\arctan\left(P^{(m)}\right)\right\}_{n_m\in\mathbb{N}}$) converging to some $\arctan\left(\bar{P}\right)$. Because $\cup_{u_i\in[0,1]^S}\mathcal{P}_{u_i^{(n_m)}}^{(n_m)}$ and thus $\arctan\left(\cup_{u_i\in[0,1]^S}\mathcal{P}_{u_i^{(n_m)}}^{(n_m)}\right)$ are closed (Lemma B1), we have:

$$\arctan\left(\bar{P}\right) \in \arctan\left(\cap_{n\in\mathbb{N}}\left(\cup_{u_i\in[0,1]^S}\mathcal{P}_{u_i^{(n_m)}}^{(n_m)}\right)\right) = \arctan\left(\mathcal{P}_{v_i}^{(1)}\right),$$

which, together with the continuity of $\psi^{(n_m)}\,(=\psi)$ as shown at the beginning of this subsection, further implies $\bar{P} = \psi(V, K)$. In other words, $\bar{P}$ is the PM price matrix in Γ selected by ψ when everyone is truth-telling.

We define the indirect utility function $W_{u_i}(P)$ as the expected utility (with respect to true preferences v_i) that i can obtain when reporting u_i given price P. By the maximum theorem, i's utility maximization problem implies that $W_{u_i}(P)$ is continuous in P. Moreover, the utility from manipulation, $W_{u_i}\left(P^{(m)}\right)$, is always bounded above by $W_{v_i}\left(P^{(m)}\right)$, and $W_{v_i}\left(P^{(m)}\right)$ goes to $W_{v_i}\left(\bar{P}\right)$ when m goes to infinite. Therefore, the (sub)sequence of $W_{u_i}\left(P^{(m)}\right)$ is bounded above by the utility from truth-telling:

$$\limsup_{m\to\infty} W_{u_i}\left(P^{(m)}\right) \leq \limsup_{m\to\infty} W_{v_i}\left(P^{(m)}\right) = W_{v_i}\left(\bar{P}\right).$$

This contradiction proves that the statement in the theorem is true for a given copy of i.

To prove the statement holds true for each copy of each agent type, we note that there is a finite number of agent types in $\Gamma^{(n)}$. There thus must exist n^* such that $n > n^*$ implies that the utility gain from unilateral misreporting for any agent is uniformly bounded by ε given that everyone else is truth-telling. ■

C Other Proofs

Proof of Proposition 1. Given an ex-ante stable assignment Π, $\forall s \in \mathcal{S}$, all the priority groups belong to one of the three categories:

(a) cut-off group, i.e., $k^*(s)$ such that $\sum_{j\in\mathcal{I},\ k_{s,j}<k^*(s)} \pi_{j,s} < q_s$, $\sum_{j\in\mathcal{I},\ k_{s,j}\leq k^*(s)} \pi_{j,s} = q_s$, and $\sum_{j\in\mathcal{I},\ k_{s,j}>k^*(s)} \pi_{j,s} = 0$;

(b) groups that have higher priority than $k^*(s)$ at s in Π, i.e., a set $\overline{\mathcal{K}}_s \subset \mathcal{K}$ such that $k \in \overline{\mathcal{K}}_s$ iff $k < k^*(s)$;

(c) groups that have lower priority than $k^*(s)$ at s in Π, i.e., a set $\underline{\mathcal{K}}_s \subset \mathcal{K}$ such that $k \in \underline{\mathcal{K}}_s$ iff $k > k^*(s)$.

Note that $k^*(s)$ always exists and is unique for all s and for any given Π, while $\overline{\mathcal{K}}_s$ or $\underline{\mathcal{K}}_s$ may be empty. As long as there are at least two priority groups, $\overline{\mathcal{K}}_s = \varnothing$ implies $\underline{\mathcal{K}}_s \neq \varnothing$, and vice versa.

(i) We first show that every Π in Π_Γ^{Stable} is ex-ante stable.

If $P \in \mathcal{P}_\Gamma^{Stable}$, then

$$p_{i,s} = \begin{cases} 0 & \text{if } k_{s,i} \in \overline{\mathcal{K}}_s \\ \in [0,+\infty] & \text{if } k_{s,i} = k^*(s) \\ +\infty & \text{if } k_{s,i} \in \underline{\mathcal{K}}_s \end{cases}$$

Fix $\Pi \in \Pi_\Gamma(P)$ for some $P \in \mathcal{P}_\Gamma^{Stable}$. $\forall i,j \in \mathcal{I}$, $\forall s, s' \in \mathcal{S}$ such that $v_{i,s} > v_{i,s'}$ and $k_{s,i} < k_{s,j}$, if $\pi_{j,s} > 0$, we must have $\pi_{i,s'} = 0$ since $p_{i,s} = 0$ according to the definition of $\mathcal{P}_\Gamma^{Stable}$. Equivalently, $\pi_{i,s'} > 0$ is not optimal for i facing $(p_{i,1}, \cdots, p_{i,S})$, which proves every Π in Π_Γ^{Stable} is *ex ante* stable.

Proof of Theorem 3. By Proposition 1 as well as Corollary 1, PM assignments are ex-ante stable. Moreover, by the definition of equal claim among ex-ante stable assignment, PM assignments also satisfy equal claim.

For any given assignment that satisfies ex-ante stability and equal claim, Proposition 1 and the definition of equal claim imply that the assignment can be rationalized by prices that satisfy the PM construction. Therefore, the assignment is a PM assignment. ■

Proof of Theorem 4. Given a stable matching, for each object s, we may find $k^*(s) = \max_{i\in\{j\in\mathcal{S}|j \text{ is matched with } s\}} \{k_{s,i}\}$, which is the lowest priority group of s among those who are matched with s. We may then define the following price system:

$$p_{s,k_{s,i}} = \begin{cases} 0, \text{ if } k_{s,i} \leq k^*(s) \\ +\infty, \text{ if } k_{s,i} > k^*(s) \end{cases}, \forall s.$$

This price system satisfies the requirement of the PM mechanism. We need to show that agents maximize their expected utility given the prices.

The only possible deviation for an agent i is to choose some object s which is free to her. That is, she is in a higher priority group of s than someone who is already accepted by s. If this deviation is profitable to i, (i, s) forms a blocking pair (or i has justified envy at s). By the definition of stability, there is no such pair. This proves that every stable matching is an PM assignment.

Similarly, for any PM assignment, there exists a corresponding price matrix that guarantees that prices are either zero or infinite, which implies that the assignment is deterministic. For deterministic assignments, ex-ante stability is equivalent to stability, and Theorem 3 implies that PM assignments in this case are stable. ■

Proof of Proposition 2. Let P^* be a PM price matrix.

Suppose that $s_{i,1}$ is the non-free object (according to her own price) on which agent i spends her budget, and that $s_{i,2}$ is her most preferred object among all free ones. By assumption, $1 < p^*_{s_{i,1},k_{s_{i,1},i}} < +\infty$. Since each agent has strict preferences over objects, $s_{i,2}$ is unique and $v_{i,s_{i,1}} > v_{i,s_{i,2}}$. By assumption, if i's consumption includes a positive probability share of $s_{i,2}$, $s_{i,2}$ must be free to everyone. i's assignment $\left\{\pi^*_{i,s}\right\}_{s\in\mathcal{S}}$ must be such that:

$$\pi^*_{i,s_{i,1}} = 1/p^*_{s_{i,1},k_{s_{i,1},i}},\ \pi^*_{i,s_{i,2}} = 1 - \pi^*_{i,s_{i,1}},\text{ and } \pi^*_{i,s} = 0,\ \forall s \neq s_{i,1}, \neq s_{i,2};$$

Alternatively, if i does not spend any budget on any non-free objects,

$$\pi^*_{i,s'_{i,2}} = 1,\text{ and } \pi^*_{i,s} = 0\ \forall s \neq s'_{i,2}.$$

Note that such $s'_{i,2}$ may or may not be free to every agent.

Consider that agent i's submitted rank-order list in BM is $L^*_i = (s_{i,1}, s_{i,2})$ if i spends some of her budget or $L^*_i = \left(s'_{i,2}\right)$ if she does not spend any budget at all. It can be verified that given these rank-order lists, BM clears the market in two rounds and delivers the same assignment as the PM mechanism. The only thing left to check is that this is a Nash equilibrium.

(i) If $L^*_i = \left(s'_{i,2}\right)$, suppose there exists s' s.t. $v_{i,s'} > v_{i,s'_{i,2}}$. If not, there is no profitable deviation for i, as she is matched with her most preferred object already. If i ranks s' above $s'_{i,2}$, she cannot be matched with s', because all those top ranking s' must be in a higher priority group of s'. Otherwise, s' would cost i a finite amount, which would allow her to purchase some shares in s' under the PM mechanism. Certainly, ranking s' below $s'_{i,2}$ does not change the assignment. Similarly, i cannot benefit by ranking objects less preferable than $s'_{i,2}$ in her list.

(ii) Now suppose $L^*_i = (s_{i,1}, s_{i,2})$ and L'_i is a profitable deviation for i. Given the assumptions, we have the following results:

(a) Object $s_{i,1}$ is not available after the first round of BM;

(b) i may obtain a positive share of s' by ranking it first if in PM $p_{s',k_{s',i}} < +\infty$ (i.e., i's priority at object s' is at least as high as the cut-off group).

(c) Only objects available in the second round and rounds later are those ranked as second choice by some agents. In other words, they are those who have zero prices for everyone in the PM mechanism.

Therefore, if L_i' still has $s_{i,1}$ as her first choice, it cannot be profitable, because she can get $1/p^*_{s_{i,1},k_{s_{i,1},i}}$ of $s_{i,1}$ and at best $1-1/p^*_{s_{i,1},k_{s_{i,1},i}}$ shares of $s_{i,2}$.

If L_i' has s' ($s' \neq s_{i,1}$) as her first choice, to be profitable, $v_{i,s'} > v_{i,s_{i,2}}$ and s' cannot be of zero price or infinite price to i in PM. If s' is of zero price to i, i could have obtained s' instead of $s_{i,2}$ in PM; if s' is of infinite price to i, i could never obtain any shares of s'. i thus must be in the cut-off priority group of s'. Given L^*_{-i} and the rules of BM, by ranking s' as first choice, i can obtain:

$$\pi'_{i,s'} = \frac{q_{s'} - \sum_{j\in\left\{j\in\mathcal{I}:k_{s',j}<k_{s',i}\right\}}\pi^*_{j,s'}}{p^*_{s',k_{s',i}}\left(q_{s'} - \sum_{j\in\left\{j\in\mathcal{I}:k_{s',j}<k_{s',i}\right\}}\pi^*_{j,s'}\right)+1},$$

where $q_{s'} - \sum_{j\in\left\{j\in\mathcal{I}:k_{s',j}<k_{s',i}\right\}}\pi^*_{j,s'}$ is the remaining quota at s' after those who are in higher priority groups claim their shares; and $p^*_{s',k_{s',i}}\left(q_{s'} - \sum_{j\in\left\{j\in\mathcal{I}:k_{s',j}<k_{s',i}\right\}}\pi^*_{j,s'}\right)$ is the total expenditure on s' by agents who are in the same priority group of s' as i; and more importantly it is the total number of such agents other than i who have ranked s' as first choice given L^*_{-i}. This is because everyone spends her budget on at most one object and $p^*_{s',k_{s',i}} > 1$ by assumption. This lead to $p^*_{s',k_{s',i}}\pi'_{i,s'} < 1$, which implies that $\pi'_{i,s'}$ is affordable to i in PM.

Moreover, given L^*_{-i} and any L_i', besides the first-choice object (s'), i can only obtain some shares in objects that are free to everyone in the PM. Therefore, the assignment resulting from a potentially deviation is still affordable to i in PM, which implies that it cannot be profitable.

This complete the proof that $\left(L^*_i, L^*_{-i}\right)$ is a Bayesian Nash equilibrium in BM. ■

(ii) We show that if $\Pi \in \mathcal{A}$ is ex-ante stable, then $\exists P^* \in \mathcal{P}^{Stable}_{\Gamma}$ such that $\Pi \in \Pi_{\Gamma}(P^*)$. It suffices to show that $\forall i \in \mathcal{I}$, $[\pi_{i,s}]_{s\in\mathcal{S}}$ is the optimal choice facing $\left[p^*_{i,s}\right]_{s\in\mathcal{S}}$ and $\left[p^*_{i,s}\right]_{s\in\mathcal{S}}$ is in $\mathcal{P}^{Stable}_{\Gamma}$.

Given Π, we can still define three sets of priorities, $\overline{\mathcal{K}}_s$, $\{k^*(s)\}$, and $\underline{\mathcal{K}}_s$. Across agents, the only restriction on prices in $\mathcal{P}^{Stable}_{\Gamma}$ is that prices for agents with priorities in $\overline{\mathcal{K}}_s \cup \underline{\mathcal{K}}_s$ and not in cut-off groups must be the same (either zero or infinite). An immediate finding is that $\forall k_{s,i} \in \underline{\mathcal{K}}_s$, we can set $p^*_{i,s} = +\infty$ since $\pi_{i,s} = 0$ for all such i and s, which satisfies the property of $\mathcal{P}^{Stable}_{\Gamma}$.

Given Π, we can further group the objects into three distinct sets for agent i, $\mathcal{S} = \underline{\mathcal{S}}_i \cup \mathcal{S}^c_i \cup \overline{\mathcal{S}}_i$:

$$\underline{\mathcal{S}}_i = \{s \in \mathcal{S} | k_{s,i} \in \underline{\mathcal{K}}_s\}; \mathcal{S}^c_i = \{s \in \mathcal{S} | k_{s,i} = k^*(s)\}; \overline{\mathcal{S}}_i = \left\{s \in \mathcal{S} | k_{s,i} \in \overline{\mathcal{K}}_s\right\}.$$

Also note that $\forall i \in \mathcal{I}$, $\mathcal{S} \backslash \underline{\mathcal{S}}_i \neq \varnothing$, and we consider the following possibilities:

(a) $\mathcal{S}^c_i = \varnothing$: The ex-ante stability implies that i is matched with her most-preferred object within $\mathcal{S} \backslash \underline{\mathcal{S}}_i = \overline{\mathcal{S}}_i$ with probability 1, thus $p^*_{i,s} = 0$ $\forall s \in \mathcal{S} \backslash \underline{\mathcal{S}}_i = \overline{\mathcal{S}}_i$ supports this assignment as a utility-maximization outcome and satisfies the properties of $\mathcal{P}^{Stable}_{\Gamma}$.

(b) $\overline{\mathcal{S}}_i = \varnothing$: This implies that $\mathcal{S} \backslash \underline{\mathcal{S}}_i = \mathcal{S}^c_i$. By adjusting the prices of objects in $\mathcal{S}^c_i$, one can make $[\pi_{i,s}]_{s\in\mathcal{S}}$ an optimal choice of i. This is feasible because there are no restrictions on prices of objects in $\mathcal{S}_i^c$.

(c) $\mathcal{S}_i^c \neq \varnothing$ and $\overline{\mathcal{S}}_i \neq \varnothing$: We denote the most-preferred object within $\overline{\mathcal{S}}_i$ for i as $\overline{s}_i$, then the ex-ante stability implies that $\forall s \in \mathcal{S}\backslash\underline{\mathcal{S}}_i$, $\pi_{i,s} = 0$ if $v_{i,\overline{s}_i} > v_{i,s}$. Let us set $p^*_{i,s} = 0$ for $\forall s \in \overline{\mathcal{S}}_i$, which satisfies the properties of $\mathcal{P}_\Gamma^{Stable}$.

Denote $\mathcal{S}_i^c(\overline{s}_i) \equiv \{s \in \mathcal{S}_i^c | v_{i,s} \geq v_{i,\overline{s}_i}\}$. If $\pi_{i,\overline{s}_i} = 0$, i must only consume objects in $\mathcal{S}_i^c(\overline{s}_i)$. Given zero prices for all objects in $\overline{\mathcal{S}}_i$ and infinite prices for objects in $\underline{\mathcal{S}}_i$, one can find a vector of personalized prices for all objects in $\mathcal{S}_i^c(\overline{s}_i)$ to make $[\pi_{i,s}]_{s\in\mathcal{S}}$ i's optimal choice. Note that this can be done independently for all agents. If instead $\pi_{i,\overline{s}_i} > 0$, it implies that i only consumes objects in $\{\overline{s}_i\} \cup \mathcal{S}_i^c$. Similarly, one can find a price vector for objects in $\mathcal{S}_i^c(\overline{s}_i)$ to make $[\pi_{i,s}]_{s\in\mathcal{S}}$ i's optimal choice.

This proves that there always exists a price matrix $P^* \in \mathcal{P}_\Gamma^{Stable}$ such that each $\left[p^*_{i,s}\right]_{s\in\mathcal{S}}$ supports $[\pi_{i,s}]_{s\in\mathcal{S}}$ if Π is *ex ante* stable. ■

Proof of Theorem 5. We define the following rules regarding infinity:

$$0 * +\infty = 0; +\infty \geq +\infty.$$

Suppose a PM assignment, $\left[\pi^*_{i,s}\right]_{i\in\mathcal{I},s\in\mathcal{S}}$, is ex-ante *Pareto dominated* by another assignment $[\pi_{i,s}]_{i\in\mathcal{I},s\in\mathcal{S}}$, i.e.,

$$\sum_{s\in\mathcal{S}} \pi_{i,s} v_{i,s} \geq \sum_{s\in\mathcal{S}} \pi^*_{i,s} v_{i,s}, \forall i \in \mathcal{I}, \tag{1}$$

$$\sum_{i\in\{k_{s,i}\leq k\}} \pi_{i,s} \geq \sum_{i\in\{k_{s,i}\leq k\}} \pi^*_{i,s}, \ \forall s \in \mathcal{S}, \forall k \in \mathcal{K}, \tag{2}$$

and at least one inequality is strict.

For any agent whose most preferred object is free or has the associated price no more than one, she obtains that object for sure, and there is no other assignment that makes her better off. If for agent i, $\sum_{s\in\mathcal{S}} \pi_{i,s} v_{i,s} > \sum_{s\in\mathcal{S}} \pi^*_{i,s} v_{i,s}$, it must be such that $\sum_{s\in\mathcal{S}} p_{s,k_{s,i}} \pi_{i,s} > 1$ and $\sum_{s\in\mathcal{S}} \pi^*_{i,s} p_{s,k_{s,i}} = 1$. Otherwise $\left[\pi^*_{i,s}\right]_{s\in\mathcal{S}}$ would not be optimal for i.

Moreover, for agents other than i who do not obtain their most preferred objects, it must be that $\sum_{s\in\mathcal{S}} p_{s,k_{s,j}} \pi_{j,s} \geq \sum_{s\in\mathcal{S}} p_{s,k_{j,s}} \pi^*_{j,s}$, since $\left[\pi^*_{j,s}\right]_{s\in\mathcal{S}}$ is the cheapest among bundles delivering the same expected utility. Therefore,

$$\sum_{s\in\mathcal{S}} p_{s,k_{s,i}} \pi_{i,s} + \sum_{j\neq i} \sum_{s\in\mathcal{S}} p_{s,k_{s,j}} \pi_{j,s} > \sum_{s\in\mathcal{S}} p_{s,k_{s,i}} \pi^*_{i,s} + \sum_{j\neq i} \sum_{s\in\mathcal{S}} p_{s,k_{s,j}} \pi^*_{j,s}.$$

However, because prices are higher for agents in lower priority groups, equation (2) implies that:

$$\sum_{j\in\mathcal{I}} \sum_{s\in\mathcal{S}} p_{s,k_{s,j}} \pi_{j,s} \leq \sum_{j\in\mathcal{I}} \sum_{s\in\mathcal{S}} p_{s,k_{s,j}} \pi^*_{j,s},$$

which leads to a contradiction.

Suppose instead that for object s, equation (2) is satisfied for all k, and $\exists \underline{k} \in \left\{1, ..., \overline{k}-1\right\}$, such that

$$\sum_{i \in \{k_{s,i} \leq \underline{k}\}} \pi_{i,s} > \sum_{i \in \{k_{s,i} \leq \underline{k}\}} \pi^*_{i,s}.$$

This implies,

$$\sum_{j \in \mathcal{I}} p_{s,k_{s,j}} \pi_{j,s} < \sum_{j \in \mathcal{I}} p_{s,k_{s,j}} \pi^*_{j,s},$$

again because prices are higher for agents in lower priority group. Aggregating over all objects,

$$\sum_{j \in \mathcal{I}} \sum_{s \in \mathcal{S}} p_{s,k_{s,j}} \pi_{j,s} < \sum_{j \in \mathcal{I}} \sum_{s \in \mathcal{S}} p_{s,k_{s,j}} \pi^*_{j,s}.$$

However, based on the same arguments as above, equation (1) implies that $\sum_{s \in \mathcal{S}} p_{s,k_{s,j}} \pi_{j,s} \geq \sum_{s \in \mathcal{S}} p_{s,k_{s,j}} \pi^*_{j,s}$, $\forall j \in \mathcal{I}$, and thus,

$$\sum_{j \in \mathcal{I}} \sum_{s \in \mathcal{S}} p_{s,k_{s,j}} \pi_{j,s} \geq \sum_{j \in \mathcal{I}} \sum_{s \in \mathcal{S}} p_{s,k_{s,j}} \pi^*_{j,s}.$$

This leads to another contradiction.

Therefore, $\left[\pi^*_{i,s}\right]_{i \in \mathcal{I}, s \in \mathcal{S}}$, must be two-sided ex-ante *efficient.* ■

C Other Incentive Compatibility Concepts for Large Markets

A SP-L of the Equal-Budget PM Mechanism

We now show that the PM mechanism giving equal budgets to agents of the same priority type satisfy Azevedo and Budish's SP-L (strategy-proof-in-the-large) property provided the domain of agents' types is finite. Let T be a finite domain of agents' vNM utility profiles, $\overline{\Delta} T$ be the set of full-support lotteries over T. Random mechanisms ϕ_n are defined on cartesian products over T.

Proposition C1 *Assume that the number of agents in every priority profile grows to infinity along a sequence of economies $(E_n)_{n \in \mathbb{N}}$. The sequence of PM mechanisms $(\phi_n)_{n \in \mathbb{N}}$ on $(E_n)_{n \in \mathbb{N}}$ that give equal budgets to agents of the same priority type is SP-L that is for any $\epsilon > 0$ and any $m \in \overline{\Delta} T$, there exists n_0 such that for all $n \geq n_0$ and all $t, t' \in T$ we have*

$$\sum_{t_{-i} \in T^{n-1}} u_t \left(\phi_n \left(t, t_{-i}\right)\right) m \left(t_{-i}\right) \geq \sum_{t_{-i} \in T^{n-1}} u_t \left(\phi_n \left(t', t_{-i}\right)\right) m \left(t_{-i}\right) - \epsilon.$$

Proof. Fix any $m \in \overline{\Delta} T$ and a priority profile π. Let $\mathcal{I}_\pi$ be the set of agents of priority profile π; let us define for these agents the random mechanism $\tilde{\phi}_n$ that takes as arguments the profile of preferences of agents in I_π and assigns them the lottery over $\phi\left(t_{\mathcal{I}_\pi}, t_{-\mathcal{I}_\pi}\right)$ where the preferences of agents of priority profiles different from π are drawn according to m. Since ϕ is envy free among $\mathcal{I}_\pi$, hence so is $\tilde{\phi}$. Thus, Proposition 1 of Azevedo and Budish (2013) implies that $\tilde{\phi}$ is SP-L, and hence we can conclude that ϕ is SP-L. ■

B Limiting Individual Incentive Compatibility

This appendix proves that the PM mechanism satisfies the concept of limiting individual incentive compatibility as in Roberts & Postlewaite (1976).

Definition C1 *Let $\left\{\Gamma^{(n)}\right\}_{n\in\mathbb{N}}$ be a sequence of economies and let i be an agent in each $\Gamma^{(n)}$. A mechanism is limiting individually incentive compatible for i in $\left\{\Gamma^{(n)}\right\}_{n\in\mathbb{N}}$ if for any ε there exists n^* such that $n > n^*$ implies that for each π_i attainable by i in $\Gamma^{(n)}$ there exists a competitive assignment π_i^* to i in $\Gamma^{(n)}$ (everyone is truth-telling) such that $\sum_{s\in\mathcal{S}} \pi_{i,s}^* v_{i,s} > \sum_{s\in\mathcal{S}} \pi_{i,s} v_{i,s} - \varepsilon$.*

Therefore, this concept focuses on the incentive for an individual agent to misreport while everyone else is truth-telling. In particular, it does not require a price selection rule, because only the existence of such a truth-telling equilibrium is required. The following shows that the PM mechanism satisfies this property in a sequence of economies.

B.1 Sequence of Economies

We first define per capita demand functions and take into account that agents in different priority groups face different prices, and thus the per capita demand is priority-specific. Let $F_i(P)$ be the augmented set of feasible consumption bundles for agent i,

$$F_i(P) \equiv \begin{cases} \left\{ \pi_i = [\pi_{i,s}]_{s\in\mathcal{S}} \;\middle|\; \begin{array}{c} \pi_{i,s} \geq 0, \forall s, \ \sum_{s\in\mathcal{S}} \pi_{i,s} = 1, \\ \text{and } \sum_{s\in\mathcal{S}} \pi_{i,s} p_{s,k_{s,i}} \leq 1 \end{array} \right\}, & \text{if } p_{s,k_{s,i}} \leq 1 \text{ for some } s; \\ \left\{ \pi_i = [\pi_{i,s}]_{s\in\mathcal{S}} \;\middle|\; \begin{array}{c} \pi_{i,s} \geq 0, \forall s, \ \sum_{s\in\mathcal{S}} \pi_{i,s} = \frac{1}{\min_{t=1,\dots,S}\left\{p_{t,k_{t,i}}\right\}}, \\ \text{and } \sum_{s\in\mathcal{S}} \pi_{i,s} p_{s,k_{s,i}} \leq 1 \end{array} \right\}, & \text{if } p_{s,k_{s,i}} > 1, \forall s. \end{cases}$$

When there are no affordable bundles such that $\sum_{s\in\mathcal{S}} \pi_{i,s} = 1$, the second part of the definition assumes that every agent is allowed to spend all their money on the cheapest objects. $F_i(P)$ is then non-empty, closed, and bounded.[32]

Let $U_i = \sum_{s\in\mathcal{S}} \pi_{i,s} v_{i,s}$ be i's expected utility function. Define $G_i(P, v_i)$ as the set of bundles that i would choose from $F_i(P)$ to maximize U_i. Formally,

$$G_i(P, v_i) = \left\{ \pi_i \in F_i(P) \;\middle|\; \begin{array}{c} \forall \pi_i' \in F_i(P), U_i(\pi_i) > U_i(\pi_i'), \\ \text{or } U_i(\pi_i) \geq U_i(\pi_i') \text{ and } \sum_{s\in\mathcal{S}} \pi_{i,s} p_s \leq \sum_{s\in\mathcal{S}} \pi_{i,s}' p_s \end{array} \right\}.$$

[32] It is important to note that P cannot be an equilibrium whenever the second part of $F_i(P)$'s definition is invoked.

Since $G_i(P, v_i)$ is obtained from the closed, bounded, and non-empty set $F_i(P)$ by maximizing (and minimizing) continuous functions, $G_i(P, v_i)$ must be non-empty. $G_i(P, v_i)$ is a convex set, because $U_i(\pi_i)$ and $\sum_{s\in\mathcal{S}} \pi_{i,s} p_{s,k_{s,i}}$ are linear functions of π_i. Define $G(P, v)$ as the set of per capita demand for each priority group for each object that can emerge when prices equal P and each agent i chooses a vector in $G_i(P, v_i)$, that is, $\forall P \in \mathcal{P}$:

$$G(P,V) = \left\{ D = [d_{s,k}]_{s\in\mathcal{S},k\in\mathcal{K}} \;\middle|\; \begin{array}{c} d_{s,k} = \frac{1}{|\mathcal{I}|}\sum_{\{i\in\mathcal{I}|k_{s,i}=k\}} \pi_{i,s},\ \forall s, \forall k \\ [\pi_{i,s}]_{s\in\mathcal{S}} \in G_i(P, v_i) \end{array} \right\}.$$

It can be verified that $G(P, V)$ is also closed, bounded, and upper hemicontinuous.

The following definition is needed to define the sequence of economies.

Definition C2 *A sequence of correspondences $f^{(n)}(P)$ uniformly converge to $f(P)$ if and only if, for any $\varepsilon > 0$, there exists $N \in \mathbb{N}$, such that when $n \geq N$,*

$$\sup_P d_H\left(f^{(n)}(P), f(P)\right) \leq \varepsilon,$$

where d_H is Hausdorff distance, i.e.,

$$d_H\left(f^{(n)}(P), f(P)\right) = \max\left\{ \begin{array}{c} \sup_{Y\in f(P)} \inf_{Y^{(n)}\in f^{(n)}(P)} \left\|Y^{(n)} - Y\right\|, \\ \sup_{Y^{(n)}\in f^{(n)}(P)} \inf_{Y\in f(P)} \left\|Y^{(n)} - Y\right\| \end{array} \right\},$$

where $\|\cdot\|$ is the Euclidean distance.

Let $\left\{\Gamma^{(n)}\right\}_{n\in\mathbb{N}}$ be a sequence of economies where $\Gamma^{(n)} = \left\{\mathcal{S}, \mathcal{I}^{(n)}, Q^{(n)}, V^{(n)}, K^{(n)}\right\}$ and $\forall n \in \mathbb{N}$:

(i) $\mathcal{I}^{(n)} \subset \mathcal{I}^{(n')}$ and $q_s^{(n)} < q_s^{(n')}$ for all s if $n < n'$; $\left|\mathcal{I}^{(n)}\right| = \sum_{s\in\mathcal{S}} q_s^{(n)}$; and $q_s^{(n)}/\left|\mathcal{I}^{(n)}\right| = q_s/I$;

(ii) $K^{(n)}$ is such that the associated priority groups satisfy $\left|\left\{i \in \mathcal{I}^{(n)} | k_{s,i} = k\right\}\right| / \left|\mathcal{I}^{(n)}\right| = C_{s,k}$, for all k and s, where $C_{s,k}$ is a constant.

(iii) the number of objects, $S = |\mathcal{S}|$, is constant;

(iv) the corresponding per capita demand $G^{(n)}\left(P, V_{-i}^{(n)}\right) \to g(P)$ uniformly as $n \to \infty$.

Remark C1 *Analogous to the regularity imposed in the main text, the above restrictions on the sequence of economies can also be interpreted as regular conditions.*

Remark C2 *$g(P)$ is a convex-valued, closed, bounded, and upper hemicontinuous correspondence, since $G^{(n)}\left(P, V^{(n)}\right)$ has these properties. This definition includes two special cases: (i) a sequence of replica economies where $G^{(n)}\left(P, V^{(n)}\right) = g(P)$, for all $n \in \mathbb{N}$; and (ii) a sequence of economies in which agents' preference-priority profiles are i.i.d. drawn from a joint distribution of preferences and priorities, while holding constant the relative size of each priority group at each object.*

B.2 Results and Proofs

We first present a result on the set of PM prices and then another on the limiting incentive compatibility.

Lemma C2 *In the sequence of economies* $\left\{\Gamma^{(n)}\right\}_{n\in\mathbb{N}}$*, let* $\mathcal{P}_{u_i}^{(n)} \subset [0,+\infty]^{S\times\overline{k}}$ *be the set of PM prices given* $\left(u_i, V_{-i}^{(n)}\right)$*. Then* $\lim_{n\to\infty} d_H\left(\mathcal{P}_{v_i}^{(n)}, \mathcal{P}_{u_i}^{(n)}\right) = 0$*,* $\forall u_i \in [0,1]^S$*, for any* i *in all* $\mathcal{I}^{(n)}$*.*

Proof. This is proven by the following three steps.

(1) Misreporting cannot affect per capita demand by priority groups in the limit.

First, recall that per capita demand of each priority group at each object is $G(P,v)$ for $P \in [0,+\infty]^{S\times\overline{k}} \equiv \mathcal{P}$ and v is the tuple of all agents' preferences.

Since each agent can increase or decrease the total demand of a priority group of an object at most by one copy, $\forall\, [d_{s,k}]_{s\in\mathcal{S},k\in\mathcal{K}} \in G^{(n)}\left(P, \left(u_i, V_{-i}^{(n)}\right)\right)$, there must exist $\left[d'_{s,k}\right]_{s\in\mathcal{S},k\in\mathcal{K}} \in G^{(n)}\left(P, \left(v_i, V_{-i}^{(n)}\right)\right)$, such that, $\forall s$, $\forall k$,

$$d'_{s,k} - \frac{1}{\left|\mathcal{I}^{(n)}\right|} \leq d_{s,k} \leq d'_{s,k} + \frac{1}{\left|\mathcal{I}^{(n)}\right|}.$$

Similarly, $\forall \left[d'_{s,k}\right]_{s\in\mathcal{S},k\in\mathcal{K}} \in G^{(n)}\left(P, \left(v_i, V_{-i}^{(n)}\right)\right)$, there exists $[d_{s,k}]_{s\in\mathcal{S},k\in\mathcal{K}} \in G^{(n)}\left(P, \left(u_i, V_{-i}^{(n)}\right)\right)$, such that $\forall s$, $\forall k$,

$$d_{s,k} - \frac{1}{\left|\mathcal{I}^{(n)}\right|} \leq d'_{s,k} \leq d_{s,k} + \frac{1}{\left|\mathcal{I}^{(n)}\right|}.$$

Therefore, given any P,

$$\sup_{u_i\in[0,1]^S} d_H\left(G^{(n)}\left(P, \left(u_i, V_{-i}^{(n)}\right)\right), G^{(n)}\left(P, \left(v_i, V_{-i}^{(n)}\right)\right)\right) \leq \frac{\sqrt{S\overline{k}}}{\left|\mathcal{I}^{(n)}\right|},$$

which implies that, given any P,

$$\lim_{n\to\infty} \sup_{u_i\in[0,1]^S} d_H\left(G^{(n)}\left(P, \left(u_i, V_{-i}^{(n)}\right)\right), G^{(n)}\left(P, \left(v_i, V_{-i}^{(n)}\right)\right)\right) = 0. \tag{3}$$

By definition, $G^{(n)}\left(P, \left(v_i, V_{-i}^{(n)}\right)\right) \to g(P)$ uniformly. Therefore, Equation (3) implies that $G^{(n)}\left(P, \left(u_i, V_{-i}^{(n)}\right)\right)$ converges to $g(P)$ uniformly as $n \to \infty$.

(2) Price Adjustment Process

Similar to the proof for Theorem 1, define $Z \equiv [z_{s,k}]_{s\in\mathcal{S},k\in\mathcal{K}} \in [0,\pi/2]^{S\times\overline{k}} \equiv \mathcal{Z}$, where $z_{s,k} = \arctan(p_{s,k})$, $\forall s$, $\forall k$.

A price adjustment process for $\Gamma^{(n)}$ is defined as,

$$H\left[Z, G^{(n)}\left(\mathcal{TAN}(Z), \left(v_i, V_{-i}^{(n)}\right)\right)\right]$$
$$\equiv \left\{ Y = [y_{s,k}]_{s\in\mathcal{S},k\in\mathcal{K}} \;\middle|\; \begin{array}{c} y_{s,k}\left([d_{s,k}]_{k\in\mathcal{K}}\right) = \min\left\{\pi/2, \max\left[0, z_{s,k} + \left(\sum_{\kappa=1}^{k} d_{s,\kappa} - q_s/I\right)\right]\right\} \\ \forall [d_{s,k}]_{s\in\mathcal{S},k\in\mathcal{K}} \in G^{(n)}\left(\mathcal{TAN}(Z), \left(v_i, V_{-i}^{(n)}\right)\right) \end{array} \right\},$$

where, $\mathcal{TAN}(Z) \equiv [\tan(z_{s,k})]_{s\in\mathcal{S},k\in\mathcal{K}}$. It is straightforward to verify that the correspondence H is a mapping from $\mathcal{Z}$ to $\mathcal{Z}$, given $\left(v_i, V_{-i}^{(n)}\right)$. Similarly,

$$H[Z, g(\mathcal{TAN}(Z))]$$
$$\equiv \left\{ Y = [y_{s,k}]_{s\in\mathcal{S},k\in\mathcal{K}} \;\middle|\; \begin{array}{c} y_{s,k}\left([d_{s,k}]_{k\in\mathcal{K}}\right) = \min\left\{\pi/2, \max\left[0, z_{s,k} + \left(\sum_{\kappa=1}^{k} d_{s,\kappa} - q_s/I\right)\right]\right\} \\ \forall [d_{s,k}]_{s\in\mathcal{S},k\in\mathcal{K}} \in g(\mathcal{TAN}(Z)); \end{array} \right\}.$$

Claim: $H\left[Z, G^{(n)}\left(\mathcal{TAN}(Z), \left(v_i, V_{-i}^{(n)}\right)\right)\right] \to H[Z, g(\mathcal{TAN}(Z))]$ uniformly as $n \to \infty$.

The uniform convergence of $G^{(n)}\left(P, \left(v_i, V_{-i}^{(n)}\right)\right)$ to $g(P)$ means that $\forall\varepsilon > 0$, $\exists N \in \mathbb{N}$, such that when $n > N$, $\forall P \in \mathcal{P}$, i.e., $\forall Z \in \mathcal{Z}$,

$$\sup_{\left[d_{s,k}^{(n)}\right]_{s\in\mathcal{S},k\in\mathcal{K}} \in G^{(n)}\left(P,\left(v_i,V_{-i}^{(n)}\right)\right)} \; \inf_{[d_{s,k}]_{s\in\mathcal{S},k\in\mathcal{K}} \in g(P)} \left\| \left[d_{s,k}^{(n)} - d_{s,k}\right]_{s\in\mathcal{S},k\in\mathcal{K}} \right\| \leq \varepsilon, \text{ and}$$
$$\sup_{[d_{s,k}]_{s\in\mathcal{S},k\in\mathcal{K}} \in g(P)} \; \inf_{\left[d_{s,k}^{(n)}\right]_{s\in\mathcal{S},k\in\mathcal{K}} \in G^{(n)}\left(P,\left(v_i,V_{-i}^{(n)}\right)\right)} \left\| \left[d_{s,k}^{(n)} - d_{s,k}\right]_{s\in\mathcal{S},k\in\mathcal{K}} \right\| \leq \varepsilon.$$

By the definition of the Euclidean distance, the first inequality implies that,

$$\sup_{\left[d_{s,k}^{(n)}\right]_{s\in\mathcal{S},k\in\mathcal{K}} \in G^{(n)}} \; \inf_{[d_{s,k}]_{s\in\mathcal{S},k\in\mathcal{K}} \in g(P)} \left\| \left[\begin{array}{c} \min\left\{\frac{\pi}{2}, \max\left[0, \arctan(p_{s,k}) + \left(\sum_{\kappa=1}^{k} d_{s,\kappa}^{(n)} - \frac{q_s}{I}\right)\right]\right\} \\ - \min\left\{\frac{\pi}{2}, \max\left[0, \arctan(p_{s,k}) + \left(\sum_{\kappa=1}^{k} d_{s,\kappa} - \frac{q_s}{I}\right)\right]\right\} \end{array} \right]_{s\in\mathcal{S},k\in\mathcal{K}} \right\|$$
$$\leq \varepsilon.$$

Or, equivalently,

$$\sup_{Y^{(n)} \in H\left[Z, G^{(n)}\left(\mathcal{TAN}(Z), \left(v_i, V_{-i}^{(n)}\right)\right)\right]} \; \inf_{Y \in H[Z, g(\mathcal{TAN}(Z))]} \left\| Y^{(n)} - Y \right\| \leq \varepsilon. \tag{4}$$

Similarly, we have,

$$\sup_{Y \in H[Z, g(\mathcal{TAN}(Z))]} \; \inf_{Y^{(n)} \in H\left[Z, G^{(n)}\left(\mathcal{TAN}(Z), \left(v_i, V_{-i}^{(n)}\right)\right)\right]} \left\| Y^{(n)} - Y \right\| \leq \varepsilon. \tag{5}$$

Since (4) and (5) are satisfied for all $n > N$ and $\forall Z \in \mathcal{Z}$, $H\left[Z, G^{(n)}\left(\mathcal{TAN}(Z), \left(v_i, V_{-i}^{(n)}\right)\right)\right]$ converges to $H[Z, g(\mathcal{TAN}(Z))]$ uniformly.

From the proof for Theorem 1, $H\left[Z, G^{(n)}\right]$ is upper hemicontinuous and convex-valued and thus satisfies all the conditions of Kakutani's fixed-point theorem.

Claim: Given $\left(v_i, V_{-i}^{(n)}\right)$ and any PM price matrix $P \in \mathcal{P}$, its positive monotonic transformation $Z \in \mathcal{Z}$ is a fixed point of $H\left[Z, G^{(n)}\left(\mathcal{TAN}(Z), \left(v_i, V_{-i}^{(n)}\right)\right)\right]$.

If P^* is a PM price matrix, there must exist a unique $k^*(s) \in \mathcal{K}$ for each s such that, for some $[d_{s,k}]_{s \in \mathcal{S}, k \in \mathcal{K}} \in G^{(n)}\left(P^*, \left(v_i, V_{-i}^{(n)}\right)\right)$,
(i) $p^*_{s,k^*(s)} \in [0, +\infty)$ and $\sum_{\kappa=1}^{k^*(s)} d_{s,\kappa} = \frac{q_s}{I}$,
(ii) $\sum_{\kappa=1}^{k} d_{s,\kappa} < \frac{q_s}{I}$ and $p^*_{s,k} = 0$ if $k < k^*(s)$, and
(iii) $d_{s,k} = 0$ and $p^*_{s,k} = +\infty$ if $k > k^*(s)$.

Let $P^* = \mathcal{TAN}(Z^*)$, given the same $[d_{s,k}]_{s \in \mathcal{S}, k \in \mathcal{K}}$, we must have

$$\begin{aligned}
\min\left\{\frac{\pi}{2}, \max\left[0, z^*_{s,k} + \left(\sum_{\kappa=1}^{k} d_{s,\kappa} - \frac{q_s}{I}\right)\right]\right\} &= 0 = z^*_{s,k}, \text{ if } k < k^*(s); \\
\min\left\{\frac{\pi}{2}, \max\left[0, z^*_{s,k} + \left(\sum_{\kappa=1}^{k} d_{s,\kappa} - \frac{q_s}{I}\right)\right]\right\} &= z^*_{s,k}, \text{ if } k = k^*(s); \\
\min\left\{\frac{\pi}{2}, \max\left[0, z^*_{s,k} + \left(\sum_{\kappa=1}^{k} d_{s,\kappa} - \frac{q_s}{I}\right)\right]\right\} &= \frac{\pi}{2} = z^*_{s,k}, \text{ if } k > k^*(s).
\end{aligned}$$

Therefore, $Z^* \in H\left[Z^*, G^{(n)}\left(\mathcal{TAN}(Z^*), \left(v_i, V_{-i}^{(n)}\right)\right)\right]$.

Note that not every fixed point of H is a PM price matrix as the proof for Theorem 1 has discussed, while the transformation of any PM price matrix is a fixed point.

Similarly, when agent i reports u_i, $H\left[Z, G^{(n)}\left(\mathcal{TAN}(Z), \left(u_i, V_{-i}^{(n)}\right)\right)\right]$ has the same properties and converges to $H[Z, g(\mathcal{TAN}(Z))]$ uniformly, since $G^{(n)}\left(P, \left(u_i, V_{-i}^{(n)}\right)\right)$ converges to $g(P)$ uniformly. In the same manner, the transformations of all the PM prices can be found as a fixed point of $H\left[Z, G^{(n)}\left(\mathcal{TAN}(Z), \left(u_i, V_{-i}^{(n)}\right)\right)\right]$.

Denote $\mathcal{P}_{v_i}^{(\infty)}$ as the set of PM prices corresponding to the subset of fixed points of $H[Z, g(\mathcal{TAN}(Z))]$ which all have PM price properties (i.e., the structure of priority-specific prices).

(3) Asymptotic Equivalence of $\mathcal{P}_{v_i}^{(\infty)}$ and $\mathcal{P}_{u_i}^{(n)}$.

As a Walrasian equilibrium required by the PM mechanism, some prices may be $+\infty$ for some

s and k. We supplement the definition of Euclidean distance by defining the following for $+\infty$:[33]

$$\begin{aligned} |(+\infty)-(+\infty)| &= 0;\ \sqrt{+\infty}=+\infty;\ (+\infty)^2=+\infty; \\ |(+\infty)-x| &= |x-(+\infty)|=+\infty,\ \forall x\in[0,+\infty); \\ \text{and } (+\infty)+x &= +\infty,\ \forall x\in[0,+\infty]. \end{aligned}$$

For any $\widehat{P}^{(n)} \in \mathcal{P}_{u_i}^{(n)}$, by definition, $\exists \left[d_{s,k}^{(n)}\right]_{s\in\mathcal{S},k\in\mathcal{K}} \in G^{(n)}\left(\widehat{P}^{(n)}, \left(u_i, V_{-i}^{(n)}\right)\right)$, such that $q_s/I = \sum_{\kappa=1}^{\bar{k}} d_{s,\kappa}^{(n)}$, $\forall s$. Since $G^{(n)}\left(P, \left(u_i, V_{-i}^{(n)}\right)\right) \to g(P)$ uniformly as $n\to\infty$,

$$\lim_{n\to\infty} \inf_{[d_{s,k}]_{s\in\mathcal{S},k\in\mathcal{K}} \in g\left(\widehat{P}^{(n)}\right)} \left\| [q_s/I]_{s\in\mathcal{S}} - \left[\sum_{\kappa=1}^{\bar{k}} d_{s,\kappa}\right]_{s\in\mathcal{S}} \right\| = 0,$$

which implies that $Z = \mathcal{TAN}^{-1}\left(\widehat{P}^{(n)}\right)$ has to be a fixed point of $H\left[Z, g\left(\mathcal{TAN}(Z)\right)\right]$ in the limit. Therefore for some $P^* \in \mathcal{P}_{v_i}^{(\infty)}$,

$$\lim_{n\to\infty} \left\| P^* - \widehat{P}^{(n)} \right\| = 0,$$

which means that, more precisely,
(i) when n is large enough, there is $[k^*(s)]_{s\in\mathcal{S}} \in \mathcal{K}^S$ such that $\forall s$, $0 \le p_{s,k^*(s)}^*, \widehat{p}_{s,k^*(s)}^{(n)} < +\infty$; $p_{s,k}^* = \widehat{p}_{s,k}^{(n)} = 0$ if $k < k_s^*$; $p_{s,k}^* = \widehat{p}_{s,k}^{(n)} = +\infty$ if $k > k_s^*$;
(ii) $\lim_{n\to\infty} \left\| \left[p_{s,k^*(s)}^*\right]_{s\in\mathcal{S}} - \left[\widehat{p}_{s,k^*(s)}^{(n)}\right]_{s\in\mathcal{S}} \right\| = 0$.

Since this is true $\forall \widehat{P}^{(n)} \in \mathcal{P}_{u_i}^{(n)}$,

$$\lim_{n\to\infty} \sup_{\widehat{P}^{(n)}\in\mathcal{P}_{u_i}^{(n)}} \inf_{P^*\in\mathcal{P}_{v_i}^{(\infty)}} \left\| P^* - \widehat{P}^{(n)} \right\| = 0. \tag{6}$$

On the other hand, for any $P^* \in \mathcal{P}_{v_i}^{(\infty)}$, by definition, $\exists [d_{s,k}]_{s\in\mathcal{S},k\in\mathcal{K}} \in g(P^*)$, such that $q_s/I = \sum_{\kappa=1}^{\bar{k}} d_{s,\kappa}$, $\forall s$. Since $G^{(n)}\left(P, \left(u_i, V_{-i}^{(n)}\right)\right)$ converges to $g(P)$ uniformly,

$$\lim_{n\to\infty} \inf_{\left[d_{s,k}^{(n)}\right]_{s\in\mathcal{S},k\in\mathcal{K}} \in G^{(n)}\left(P^*, \left(u_i, V_{-i}^{(n)}\right)\right)} \left\| [q_s/I]_{s\in\mathcal{S}} - \left[\sum_{\kappa=1}^{\bar{k}} d_{s,\kappa}^{(n)}\right]_{s\in\mathcal{S}} \right\| = 0,$$

which implies that P^* is an asymptotic PM price matrix for $\left(u_i, V_{-i}^{(n)}\right)$, i.e.,

$$\lim_{n\to\infty} \inf_{\widehat{P}^{(n)}\in\mathcal{P}_{u_i}^{(n)}} \left\| P^* - \widehat{P}^{(n)} \right\| = 0.$$

[33] $p_{s,k} = +\infty$ means that there is no supply for the preference group k at school s. It therefore makes sense to define the distance between $+\infty$ and $+\infty$ as 0.

Thus the above two properties (i) and (ii) are satisfied. Since this is true for all $P^* \in \mathcal{P}_{v_i}^{(\infty)}$, therefore

$$\lim_{n\to\infty} \sup_{P^*\in\mathcal{P}_{v_i}^{(\infty)}} \inf_{\widehat{P}^{(n)}\in\mathcal{P}_{u_i}^{(n)}} \left\| P^* - \widehat{P}^{(n)} \right\| = 0. \tag{7}$$

Combining (6) and (7), we have $\lim_{n\to\infty} d_H\left(\mathcal{P}_{v_i}^{(\infty)}, \mathcal{P}_{u_i}^{(n)}\right) = 0$, $\forall u_i \in [0,1]^S$ and for any i in all $\mathcal{I}^{(n)}$.

Furthermore, $\lim_{n\to\infty} d_H\left(\mathcal{P}_{v_i}^{(\infty)}, \mathcal{P}_{v_i}^{(n)}\right) = 0$ and, therefore, $\lim_{n\to\infty} d_H\left(\mathcal{P}_{v_i}^{(n)}, \mathcal{P}_{u_i}^{(n)}\right) = 0$, $\forall u_i \in [0,1]^S$ and for any i in all $\mathcal{I}^{(n)}$. ■

Proposition C2 *Suppose i is in every economy of the sequence $\left\{\Gamma^{(n)}\right\}_{n\in\mathbb{N}}$. The PM mechanism is limiting individually incentive compatible for i.*

Proof. By Lemma C2, for any $\xi > 0$, there exists n^* such that for $n > n^*$ and for every price in $P_{u_i} \in \mathcal{P}_{u_i}^{(n)}$ there exists a price $P_{v_i} \in \mathcal{P}_{v_i}^{(n)}$ such that $|P_{v_i} - P_{u_i}| < \xi$.

We define the indirect utility function $W_{u_i}(P)$ as the expected utility (with respect to true preferences v_i) that i can obtain when reporting u_i given price P. By the maximum theorem, i's utility maximization problem implies that $W_{u_i}(P)$ is continuous in P. Moreover, the utility from manipulation, $W_{u_i}(P)$, is always bounded above by $W_{v_i}(P)$. Therefore, $W_{u_i}(P_{u_i}) \leq W_{v_i}(P_{u_i})$.

When ξ is set small enough, the continuity of $W_{v_i}(.)$ implies that we can find $P_{v_i} \in \mathcal{P}_{v_i}^{(n)}$ in all large enough economies ($n > n^*$) such that:

$$|W_{v_i}(P_{u_i}) - W_{v_i}(P_{v_i})| < \varepsilon.$$

Therefore,

$$W_{u_i}(P_{u_i}) \leq W_{v_i}(P_{u_i}) < W_{v_i}(P_{v_i}) + \varepsilon.$$

Or equivalently,

$$W_{v_i}(P_{v_i}) > W_{u_i}(P_{u_i}) - \varepsilon,$$

which proves that the PM mechanism is limiting individually incentive compatible for i. ■

D An Extension to Multi-Unit Demand

Our PM mechanism can be naturally extended to multi-unit allocations with priorities, such as course allocation in colleges in which students $i \in \mathcal{I}$ are given priority for courses $s \in \mathcal{S}$. Students' have cardinal utility over course bundles which are additive in their utility of individual courses, denoted $u_i = (u_{i,1}, ..., u_{i,S}) \in \mathbb{R}^S$.[34] An individual course assignment for agent i is a vector $\pi_i = (\pi_{i,1}, ..., \pi_{i,S})$ such that $\sum_{s=1}^{S} \pi_{i,s} = C$, where $C \geq 0$ is the maximum number of courses an agent might take. The utility from a probabilistic bundle π_i is the scalar product $u_i \pi_i$. We may allow

[34] In introducing additivity, we follow Budish et al. (2013); in their setup there are no priorities and PM prices and assignment exist.

dummy courses with guaranteed excess supply to accommodate the possibility that an agent takes less than C real courses.

We define the multi-unit PM mechanism as in the unit-demand setting. In particular, for each course s, a cut-off priority group $k^*(s)$ would face the market price $p_s^* \in [0, \infty)$. Higher priority groups would face zero price and lower priority groups would face infinite price. The existence of multi-unit PM assignment is now obtained in the same way as in the the unit-demand allocation problem.

As before, an assignment is **ex-ante stable** if it does not cause ex-ante justified envy. An assignment Π causes **ex-ante justified envy** of $i \in \mathcal{I}$ toward $j \in \mathcal{I} \setminus \{i\}$ if $\exists s, s' \in \mathcal{S}$ such that $v_{i,s} > v_{i,s'}$, $k_{s,i} < k_{s,j}$, $\pi_{j,s} > 0$, and $\pi_{i,s'} > 0$. That is, agent i who has higher-priority at s than another agent j has ex-ante justified envy towards j if j has positive probability of obtaining object s, while with positive probability i obtains an object less preferable than s. As before, if an assignment causes ex-ante justified envy, then its every implementation with positive probability generates deterministic assignments that are not justified-envy-free, or not stable, in the sense of Abdulkadiroglu & Sonmez (2003).

In multi-unit demand settings, our PM mechanism remains ex-ante stable. Indeed, if j has positive probability of course s then agent i who has higher priority than j at s would face zero price for course s.

Furthermore, we can also easily accommodate the natural constraint that each agent can consume at most one unit of any given course. Given such a constraint, an assignment Π causes ex-ante justified envy of $i \in \mathcal{I}$ toward $j \in \mathcal{I} \setminus \{i\}$ if $\exists s, s' \in \mathcal{S}$ such that $v_{i,s} > v_{i,s'}$, $k_{s,i} < k_{s,j}$, $\pi_{j,s} > 0$, $\pi_{i,s'} > 0$, and $\pi_{i,s} = 0$. That is, agent i who has higher-priority at s than another agent j has ex-ante justified envy towards j if j has positive probability of obtaining object s, while i has probability zero of s and strictly higher probability of an object worse than s.[35] An assignment is ex-ante stable if it does not cause ex-ante justified envy. Our PM mechanism remains ex-ante stable also in this environment.

[35] We are adding the restriction that $\pi_{i,s} = 0$ because of the constraint that each agent can consume at most one unit of any given course. Consider, for instance, agents who—under this constraint—pick two courses from among two courses, s_1 and s_2, each of which has at least two units. Suppose agent i has priority at s_1 over other agents and that i strictly prefers s_1 over s_2. Feasibility then implies that i obtains at most one unit of s_1. Under the unconstrained version of the stability definition, no agent other than i could obtain positive probability of s_1, which is not a reasonable restriction. The above definition resolves this issue.