The Combinatorial Assignment Problem: Approximate Competitive Equilibrium from Equal Incomes

Eric Budish

University of Chicago

This paper proposes a new mechanism for combinatorial assignment—for example, assigning schedules of courses to students—based on an approximation to competitive equilibrium from equal incomes (CEEI) in which incomes are unequal but arbitrarily close together. The main technical result is an existence theorem for approximate CEEI. The mechanism is approximately efficient, satisfies two new criteria of outcome fairness, and is strategyproof in large markets. Its performance is explored on real data, and it is compared to alternatives from theory and practice: all other known mechanisms are either unfair ex post or manipulable even in large markets, and most are both manipulable and unfair.

I. Introduction

In a combinatorial assignment problem, a set of indivisible objects is to be allocated among a set of heterogeneous agents, the agents demand bundles of the objects, and monetary transfers are exogenously prohib-

I am especially grateful to Susan Athey and Al Roth for their guidance and support throught this project. Sheila Connelly at Harvard Business School and Amy Wright and Lisa Messaglia at Chicago Booth graciously shared data and institutional knowledge of their respective course-allocation systems. Estelle Cantillon and Fuhito Kojima provided extensive comments on an early draft of the paper. I am grateful to the editor and anonymous referees for valuable comments and suggestions. For helpful discussions I thank Attila Ambrus, Eduardo Azevedo, Avinash Dixit, Itay Fainmesser, Michael Faye, Drew Fudenberg, Matt Gentzkow, Michel Goemans, Jerry Green, Richard Holden, Louis Kaplow, Adriaan Ten Kate, Scott Kominers, David Laibson, Robin Lee, Jacob Leshno, Greg Lewis, John Lewis, Paul Milgrom, Markus Mobius, Muriel Niederle, Matt Notowidigdo, Ariel Pakes, Parag Pathak, David Parkes, Canice Prendergast, Amartya Sen, Neil Sloane, Hugo

[*Journal of Political Economy,* 2011, vol. 119, no. 6]
 0022-3808/2011/11906-0002$10.00

ited. A motivating example is course allocation at educational institutions: if, because of limits on class size, it is not possible for all students to take their most desired schedule of courses, then how should seats in overdemanded courses be allocated?[1] Other examples include the assignment of shifts or tasks to interchangeable workers, leads to salespeople, players to sports teams, airport takeoff and landing slots to airlines, and shared scientific resources to scientists.[2]

Combinatorial assignment is one feature removed from several well-known market design problems. It is like a combinatorial auction problem except for the restriction against monetary transfers.[3] It differs from a matching problem in that preferences are one-sided: objects do not have preferences over the agents.[4] It generalizes the house allocation problem, which restricts attention to the case of unit demand.[5]

Yet, despite its similarity to problems that have been so widely studied, progress on combinatorial assignment has remained elusive. The literature consists mostly of impossibility theorems that suggest that there is a particularly stark tension among concerns of efficiency, fairness, and incentive compatibility. The main result is that the only mechanisms that are Pareto efficient and strategyproof are dictatorships,[6] which,

Sonnenschein, Ross Starr, Adi Sunderam, and Glen Weyl and seminar audiences at the First Conference on Auctions, Market Mechanisms, and Their Applications (2009), the Behavioral and Quantitative Game Theory Conference on Future Directions (2010), Chicago, Columbia, Harvard, London School of Economics, Microsoft Research, National Bureau of Economic Research Market Design Workshop, New York University, Rice, Stanford, Texas A&M, Yahoo! Research, and Yale. For financial support, I am grateful to the Division of Research at Harvard Business School, the National Science Foundation Graduate Research Fellowship Program, the Project on Justice, Welfare and Economics at Harvard University, and the University of Chicago Booth School of Business.

[1] Press coverage and prior research suggest that the scarcity problem is particularly acute in higher education, especially at professional schools. See Guernsey (1999), Bartlett (2008), Lehrer (2008), Levitt (2008*a*, 2008*b*), and especially Sönmez and Ünver (2003).

[2] On shifts, leads, players, airports, and scientific "collaboratories," respectively, see McKesson's eShift Web site (http://bit.ly/opqrlH), incentalign.com, Albergotti (2010), Shulman (2008), and Wulf (1993). Whether monetary transfers are permitted often varies by context; for instance, McKesson's nursing shift assignment software, eShift, has both a fixed-price version and an auction version, depending on whether the client hospital has discretion to use flexible wages (e.g., because of union restrictions). Prendergast and Stole (1999) and Roth (2007) explore foundations for constraints against monetary transfers.

[3] See Milgrom (2004) and Cramton, Shoham, and Steinberg (2006) for textbook treatments that discuss both theory and applications.

[4] See Roth and Sotomayor (1990) for a textbook treatment and Roth (1984) on a well-known application.

[5] See Sönmez and Ünver (2011) for a survey treatment and Chen and Sönmez (2002), Abdulkadiroğlu, Pathak, and Roth (2005), and Abdulkadiroğlu et al. (2005) on applications to housing markets and school choice. Another name for this problem is single-unit assignment.

[6] For precise statements, see Klaus and Miyagawa (2001), Pápai (2001), Ehlers and Klaus (2003), and Hatfield (2009); see also n. 21. Sönmez (1999) and Konishi, Quint, and Wako (2001) obtain related negative results under slightly different conditions, including existing endowments. Zhou (1990) and Kojima (2009) obtain related negative results for random mechanisms.

while intuitively sensible and widely used in single-unit assignment (Abdulkadiroğlu and Sönmez 1998, 1999), seem unreasonably unfair in the multiunit case: for any two agents, one gets to choose all her objects before the other gets to choose any. Practitioners have designed a variety of nondictatorial mechanisms, often citing fairness as a central design objective: for example, Wharton's course allocation system is "designed to achieve an *equitable* and *efficient* allocation of seats in elective courses when demand exceeds supply" (2011, 2; emphasis in original).[7] But the mechanisms found in practice have a variety of flaws, most notably with respect to incentives (Sönmez and Ünver 2003, 2010; Krishna and Ünver 2008; Budish and Cantillon, forthcoming).

Missing from both theory and practice is a mechanism that is attractive in all three dimensions of interest: efficiency, fairness, and incentives. This paper proposes such a mechanism. It gets around the impossibility theorems by making several small compromises versus the ideal properties a mechanism should satisfy.

The mechanism is based on an old idea from general equilibrium theory, the competitive equilibrium from equal incomes (CEEI). CEEI itself need not exist in our environment: either indivisibilities or complementarities alone would complicate existence (cf. Varian 1974), and our economy features both. I prove existence of an approximation to CEEI in which (i) agents are given approximately equal instead of exactly equal budgets of an artificial currency, and (ii) the market clears approximately instead of exactly. The first welfare theorem implies that this *approximate CEEI* is Pareto efficient but for the market-clearing approximation; the equal-budgets approximation will play a key role in ensuring fairness. If instead we were to give agents exactly equal budgets, then market-clearing error could be arbitrarily large. At the other extreme, the dictatorships mentioned above can be interpreted as exact competitive equilibria but from arbitrarily unequal budgets.

The second step in the analysis is to articulate what fairness realistically means in this environment: indivisibilities complicate fair division. For instance, if there is a single star professor for whom demand exceeds supply, some ex post unfairness is inevitable. My approach is to weaken Steinhaus's (1948) fair share and Foley's (1967) envy-freeness to accommodate indivisibilities in a realistic and intuitively sensible way. In particular, I want to articulate that if there are two star professors, it is

[7] Here are some additional examples: New York University Law School writes that its system "promotes a fair allocation of coveted classes" (Adler et al. 2008); Massachusetts Institute of Technology's Sloan School writes that its system "establish[es] a 'fair playing field' for access to Sloan classes" (MIT 2008); Harvard Business School has described fairness as its central design objective in numerous conversations with the author regarding the design of its course-allocation system; McKesson advertises its software product for assigning nurses to vacant shifts on the basis of their preferences as "equitable open shift management."

unfair for some students to get both while others who want both get neither. I define an agent's *maximin share* as the most preferred bundle he could guarantee himself as divider in divide-and-choose against adversarial opponents; the *maximin share guarantee* requires that each agent gets a bundle he weakly prefers to his maximin share. I say that an allocation satisfies *envy bounded by a single good* if, whenever some agent i envies another agent i', by removing some single object from i''s bundle we can eliminate i's envy. Dictatorships clearly fail both criteria in combinatorial assignment. Note, though, that dictatorships actually satisfy both criteria in single-unit assignment, for which they are often observed in practice. The criteria's ability to make sense of the empirical pattern of dictatorship usage is a useful external validity check.

The third step asks the logical next question given steps one and two: does approximate CEEI satisfy the fairness criteria? I show that it approximately satisfies the maximin share guarantee and bounds envy by a single good. The key for both of these results is that the existence theorem allows for budget inequality to be arbitrarily small, as long as budgets are not exactly equal. If budgets could be exactly equal without compromising existence, then the allocation would exactly satisfy the maximin share guarantee and be exactly envy-free.

The last step is to formally define the approximate CEEI mechanism (A-CEEI) on the basis of the existence and fairness theorems described above and analyze its incentive properties. While A-CEEI is not strategyproof like a dictatorship, it is straightforward to show that it is *strategyproof in the large*, that is, strategyproof in a limit market in which agents act as price takers. This is in sharp contrast to the course-allocation mechanisms found in practice and also to most fair-division procedures proposed outside of economics (see table 1 below); these mechanisms are manipulable even by the kinds of agents we usually think of as price takers; that is, they are *manipulable in the large*.

Some intuition for both the existence and fairness theorems can be provided by means of a simple example in which there are two agents with identical, additively separable preferences over four objects. Two of the objects are valuable diamonds (big and small), and two of the objects are ordinary rocks (pretty and ugly). Agents require at most two objects each and have the preferences we would expect given the objects' names;[8] think of the diamonds as seats in courses by star professors.

If agents have the same budget, then there is no competitive equilibrium: at any price vector, for each object, either both agents demand the object or neither does. Notice too that discontinuities in aggregate

[8] Specifically, if we label the big diamond as a, the small diamond as b, the pretty rock as c, and the ugly rock as d, each agent has preferences $\succ$: $\{a, b\}$, $\{a, c\}$, $\{a, d\}$, $\{a\}$, $\{b, c\}$, $\{b, d\}$, $\{b\}$, $\{c, d\}$, $\{c\}$, $\{d\}$, $\emptyset$.

demand are "large": any change in price that causes one agent to change his demand causes the other agent to change his demand as well. These discontinuities are the reason why existence is so problematic with equal incomes.

Suppose instead that one agent, chosen at random, is given a slightly larger budget than the other. Now there exist prices that exactly clear the market: set prices such that only the wealthier agent can afford the big diamond, whereas the less wealthy agent, unable to afford the big diamond, instead buys the small diamond and the pretty rock. Furthermore, this allocation satisfies the fairness criteria. The agent who gets {small diamond, pretty rock} may envy the agent who gets {big diamond, ugly rock}, but his envy is bounded by a single good, and he does as well as he could have as divider in divide-and-choose.[9]

Let me make a few further remarks about this example. First, note that it is critical for fairness that budget inequality is sufficiently small. Otherwise, there will exist prices at which the wealthier agent can afford both diamonds whereas the poorer agent can afford neither, leading to the same result as a dictatorship.

Second, it is also critical for fairness that we use item prices, and not the more flexible bundle prices that are commonly used in combinatorial auctions (e.g., Parkes 2006). Otherwise, we can price the bundle {big diamond, small diamond} at the wealthier agent's budget without having to price any bundles that contain just a single diamond at a level affordable by the poorer agent, again leading to the same result as a dictatorship.

Third, another way to achieve the allocation in which one agent receives {big diamond, ugly rock} but the other receives {small diamond, pretty rock} is to use a simple draft procedure in which agents choose objects one at a time and the choosing order reverses each round. Such a draft is used to allocate courses at Harvard Business School. A-CEEI and the draft differ in general; in particular, the draft has incentive problems (Budish and Cantillon, forthcoming). But the two mechanisms are similar in that they both distribute "budgets"—of artificial currency and choosing times, respectively—as equally as possible.

Finally, in this simple example an arbitrarily small amount of budget inequality is enough to ensure exact market clearing. In general, when we allow agents to have heterogeneous and nonadditive preferences, existence requires that we allow for a "small" amount of market-clearing error. I use "small" in two senses. First, the worst-case bound for market-clearing error does not grow with the number of agents or the number

[9] There also exist prices at which the wealthier agent gets {big diamond, pretty rock} whereas the poorer agent gets {small diamond, ugly rock}. This allocation still bounds envy by a single good but, as noted above, only approximately satisfies the maximin share guarantee. See Sec. V.A for further detail.

of copies of each object; so in the limit, worst-case market-clearing error as a fraction of the endowment goes to zero. This is similar in spirit to a famous result of Starr (1969) in the context of divisible-goods exchange economies with continuous but nonconvex preferences. Second, the worst-case bound is economically small for life-sized problems, and of course average case is better than worst case.[10]

A-CEEI compromises among the competing design objectives of efficiency, fairness, and incentive compatibility. To help assess whether it constitutes an attractive compromise, I perform two additional analyses. First, I compare A-CEEI to alternative mechanisms. Table 1 below describes the properties of all previously known mechanisms from both theory and practice. Every other mechanism is either severely unfair ex post or manipulable even in limit markets, and most are both unfair and manipulable. Of special note is the widely used bidding points auction, first studied by Sönmez and Ünver (2003, 2010), which resembles exact CEEI but makes a subtle mistake: it treats fake money as if it were real money that enters the agent's utility function. This mistake can lead to outcomes in which an agent gets zero objects and is implicitly expected to take consolation in a large budget of unspent fake money with no outside use. Such outcomes occur surprisingly frequently in some simple data provided by the University of Chicago's Booth School of Business, one of many educational institutions that use this mechanism.

Second, I examine the performance of A-CEEI on real preference data from Harvard Business School. There are four findings. First, average-case market-clearing error is just a single seat in six courses. Second, students' outcomes always substantially exceed their maximin shares. Third, on average, 99 percent of students have no envy, and for the remainder envy is small in utility terms. Finally, the distribution of students' utilities first-order stochastically dominates that from the actual play of Harvard's own draft mechanism, which Budish and Cantillon (forthcoming) show itself second-order stochastically dominates that from truthful play of the random serial dictatorship. This last finding suggests that a utilitarian social planner, who does not regard fairness as a design objective per se (cf. Kaplow and Shavell 2001, 2007), prefers A-CEEI to both the draft and the dictatorship in this context.

The remainder of the paper is organized as follows. Section II defines the environment. Section III defines approximate CEEI and presents the existence theorem. Section IV proposes the new criteria of outcome

[10] In the specific context of course allocation, a small amount of market-clearing error is not too costly in practice, for reasons discussed in Sec. III.C. In other contexts, such as assigning pilots to planes, market-clearing error is much more costly. Two variants of the proposed mechanism for such contexts are discussed in an earlier version of this paper (Budish 2010).

fairness. Section V provides the two fairness theorems. Section VI formally defines the A-CEEI mechanism and discusses incentives. Section VII compares A-CEEI to alternatives. Section VIII examines A-CEEI's performance on real data. Section IX concludes. Proofs are in the body of the text when both short and instructive; otherwise they are in the appendices. The text also contains a detailed sketch of the proof of the existence theorem.

II. Environment

The combinatorial assignment problem.—A combinatorial assignment problem consists of a set of objects, each with integral capacity, and a set of agents, each with scheduling constraints and preferences. I emphasize the example of course allocation at universities, in which the objects are "courses" and the agents are "students." The elements of a problem $(\mathcal{S}, \mathcal{C}, (q_j)_{j=1}^{M}, (\Psi_i)_{i=1}^{N}, (\succsim_i)_{i=1}^{N})$, also called an *economy*, are defined as follows.

Agents.—There is a set $\mathcal{S}$ of agents (students), with $|\mathcal{S}| = N$.

Objects and capacities.—There is a set $\mathcal{C}$ of object types (courses), with $|\mathcal{C}| = M$. There are q_j copies of object j (seats in course j). There are no other goods in the economy. In particular, there is no divisible numeraire such as money.

Schedules and preferences.—A consumption bundle (schedule) consists of zero or one seat in each course.[11] The set of all possible schedules is the power set of $\mathcal{C}$, that is, $2^{\mathcal{C}}$. Each student i is endowed with a complete, reflexive, transitive preference relation, $\succsim_i$, defined on the set of schedules. I assume that preferences are strict and use $x \succ_i x'$ to mean not $x' \succsim_i x$.

It is convenient to treat schedules as both sets and vectors. In set form, a schedule x is a subset of $\mathcal{C}$; in vector form, a schedule x is an element of $\{0, 1\}^M$. I adopt $\succ_i : x, x', \ldots$ as notational shorthand for $x \succ_i x' \succ_i x''$ for all $x \in 2^{\mathcal{C}} \backslash \{x, x'\}$. I use the operator $\max_{(\succsim)} X$ to pick out the element of set $X \subseteq 2^{\mathcal{C}}$ that is maximal for $\succsim$; that is, $x = \max_{(\succsim)} X$ means $x \in X$ and $x \succ x'$ for all $x' \in X$, $x' \neq x$. The operator $\min_{(\succsim)} X$ is defined analogously.

In practical applications, each student will have a limited set of permissible schedules. For instance, students take at most a certain number

[11] The restriction that agents consume at most one of each type of object is technically without loss of generality. Any economy in which this assumption does not hold can be transformed into one in which it does by giving each copy of each object its own serial number (as in, e.g., Ostrovsky [2008]). However, the market-clearing bound of theorem 1 will be most compelling in environments in which individual agents' consumptions are small relative to the goods endowment. The role this issue plays in the proof of theorem 1 is described in Sec. III.D.1.

of courses per term, they cannot take two courses that meet at the same time, and some courses may have prerequisites. I assume that the market administrator endows each student i with a set $\Psi_i \subseteq 2^{\mathcal{C}}$ of permissible schedules and that $\emptyset \succ_i x$ for $x \notin \Psi_i$. When designing the language by which students report their preferences, it may be useful to exploit the market administrators' knowledge of the Ψ_i's (see, e.g., Othman, Budish, and Sandholm 2010).

No other restrictions are placed on preferences: in particular, students are free to regard courses as complements and substitutes. This is the reason the assignment problem is called "combinatorial" as opposed to the multiunit assignment problem studied by Sönmez and Ünver (2003, 2010) and Budish and Cantillon (forthcoming).

Allocations, mechanisms, and efficiency.—An *allocation* $\mathbf{x} = (x_i)_{i \in \mathcal{S}}$ assigns a schedule x_i to each agent $i \in \mathcal{S}$. Allocation $\mathbf{x}$ is *feasible* if $\sum_i x_{ij} \leq q_j$ for each object $j \in \mathcal{C}$. A *mechanism* is a systematic procedure, possibly with an element of randomness, that selects an allocation for each economy.

A feasible allocation is (*ex post*) *Pareto efficient* if there is no other feasible allocation weakly preferred by all agents, with at least one strict preference. The mechanism developed in this paper will be approximately feasible and approximately ex post efficient in a sense that will be made clear in Section III. Two variants of this mechanism that are exactly feasible and exactly ex post efficient are described in an earlier version of the paper (Budish 2010); these variants are less attractive with respect to fairness and incentives.

III. The Approximate Competitive Equilibrium from Equal Incomes

CEEI is well known to be an attractive solution to the problem of efficient and fair division of divisible goods.[12] Arnsperger (1994, 161) writes that "essentially, to many economists, [CEEI is] the description of perfect justice." CEEI's appeal extends beyond economics. The philosopher Ronald Dworkin (1981, 2000) argues extensively that CEEI is fair and uses CEEI as the motivation for an important theory of fairness (see, e.g., Sen 1979, 2009).

CEEI would be an attractive solution to our problem of combinatorial assignment as well were it not for existence problems. Either indivisibilities or complementarities alone can cause existence problems, and our economy features both. In order to recover existence, we will approximate both the "CE" and the "EI" of CEEI.

[12] See Foley (1967), Varian (1974), and several other seminal references summarized in Thomson and Varian (1985).

A. *Definition of Approximate CEEI*

DEFINITION 1. Fix an economy $(\mathcal{S}, \mathcal{C}, (q_j)_{j=1}^M, (\Psi_i)_{i=1}^N, (\succsim_i)_{i=1}^N)$. The allocation $\mathbf{x}^* = (x_1^*, \ldots, x_N^*)$, budgets $\mathbf{b}^* = (b_1^*, \ldots, b_N^*)$, and item prices $\mathbf{p}^* = (p_1^*, \ldots, p_M^*)$ constitute an (α, β)–*approximate competitive equilibrium from equal incomes* ((α, β)–CEEI) if

i. $x_i^* = \max_{(\succsim_i)}\{x' \in 2^{\mathcal{C}} : \mathbf{p}^* \cdot x' \leq b_i^*\}$ for all $i \in \mathcal{S}$;

ii. $\|\mathbf{z}^*\|_2 \leq \alpha$, where $\mathbf{z}^* = (z_1^*, \ldots, z_M^*)$ and
 a. $z_j^* \equiv \sum_i x_{ij}^* - q_j$ if $p_j^* > 0$,
 b. $z_j^* \equiv \max(\sum_i x_{ij}^* - q_j, 0)$ if $p_j^* = 0$;

iii. $1 \leq \min_i (b_i^*) \leq \max_i (b_i^*) \leq 1 + \beta$.

Condition i indicates that, at the competitive equilibrium prices and budgets, each agent chooses her most preferred schedule that costs weakly less than her budget. Condition ii is where I approximate "CE." The market is allowed to clear with some error, α, calculated as the Euclidean distance of the excess demand vector, $\mathbf{z}^*$. This market-clearing error will be discussed in detail in Section III.C. Condition iii is where I approximate "EI." The largest budget can be no more than β proportion larger than the smallest budget. The parameter β will play a key role in the fairness theorems.

If $\alpha = \beta = 0$, then we have an exact CEEI. This version of exact CEEI is stated a bit differently from the classical version (e.g., Varian 1974) because agents have equal incomes of an artificial currency rather than equal shares of a divisible-goods endowment. The currency-endowment formulation of competitive equilibrium is sometimes called the "Fisher model" after Irving Fisher (see, e.g., Brainard and Scarf 2005; Vazirani 2007).

B. *The Existence Theorem*

THEOREM 1. Fix an economy $(\mathcal{S}, \mathcal{C}, (q_j)_{j=1}^M, (\Psi_i)_{i=1}^N, (\succsim_i)_{i=1}^N)$. Let $k = \max_{i \in \mathcal{S}} \max_{x \in \Psi_i} |x|$ denote the maximum number of objects in a permissible schedule, and let $\sigma = \min(2k, M)$.

1. For any $\beta > 0$, there exists a $(\sqrt{\sigma M}/2, \beta)$–CEEI.
2. Moreover, for any $\beta > 0$, any budget vector $\mathbf{b}'$ that satisfies $1 \leq \min_i (b_i') \leq \max_i (b_i') \leq 1 + \beta$, and any $\varepsilon > 0$, there exists a $(\sqrt{\sigma M}/2, \beta)$–CEEI with budgets of $\mathbf{b}^*$ that satisfy $|b_i^* - b_i'| < \varepsilon$ for all $i \in \mathcal{S}$.

The dictatorship mechanisms extensively studied in the prior literature on combinatorial assignment correspond to an (α, β)–CEEI with no market-clearing error ($\alpha = 0$) but substantial budget inequality β. Specifically, if there are at most k objects in a permissible schedule, then budgets of 0, $(1+k)$, $(1+k)^2$, …, $(1+k)^{N-1}$ implement the dictatorship.

On the other hand, if we seek an (α, β)–CEEI with exactly equal incomes ($\beta = 0$), then it is not possible to provide a meaningful guarantee on market-clearing error α. Consider the case in which all agents have the same preferences: at any price vector, all agents demand the same bundle, so demand for each object is always either zero or N, irrespective of supply.

Theorem 1 indicates that any strictly positive amount of budget inequality is enough to ensure that there is a price vector whose market-clearing error is at worst $\sqrt{\sigma M}/2$. Part 2 of the theorem statement indicates that the market administrator can assign these close but unequal budgets to agents however she likes, subject to an ε perturbation that can be made arbitrarily small. Two natural choices are (i) assign budgets uniform randomly and (ii) assign budgets on the basis of some preexisting priority order, such as seniority or grade point average.

C. *Discussion of Market-Clearing Error*

There are two senses in which $\sqrt{\sigma M}/2$ is "small." First, $\sqrt{\sigma M}/2$ does not grow with either N (the number of agents) or $(q_j)_{j=1}^{M}$ (object quantities). This means that as N, $(q_j)_{j=1}^{M} \rightarrow \infty$, we converge toward exact market clearing, in the sense that error goes to zero as a fraction of the endowment. This notion of approximate market clearing was often emphasized in the prior literature on general equilibrium with nonconvexities (e.g., Starr 1969; Arrow and Hahn 1971; Dierker 1971).

Second, $\sqrt{\sigma M}/2$ is actually a small number for practical problems, especially as a worst-case bound. For instance, in a semester at Harvard Business School, students require five courses each and there are about 50 courses overall, so $\sqrt{\sigma M}/2 \approx 11$. Furthermore, if some courses are known to be in substantial excess supply, we can reformulate the problem as one of allocating only the potentially scarce courses.[13] In the HBS data described in Section VIII.A, only about 20 courses per semester are ever scarce, so in the reformulated problem the bound becomes $\sqrt{\sigma M}/2 \approx 7$. This corresponds to a maximum market-clearing error of seven seats in one class, or of two seats in each of 12 classes (since $\sqrt{12 \cdot 2^2} \approx 7$), and so on, as compared with about 4,500 course seats allocated each semester.

I show below that the $\sqrt{\sigma M}/2$ bound is tight. For comparison, Starr's (1969) bound, developed in the context of a divisible-goods exchange

[13] Write $\mathcal{C} = \mathcal{C}^{\text{scarce}} \cup \mathcal{C}^{\text{unscarce}}$. In the reformulated problem, students' preferences are defined over subsets of $\mathcal{C}^{\text{scarce}}$, with student i's preference for scarce-course bundle $x \subseteq \mathcal{C}^{\text{scarce}}$ based on the best bundle he can form by adding a subset of $\mathcal{C}^{\text{unscarce}}$ to x.

economy with continuous but nonconvex preferences, would be $M/2$ if it applied to this environment. Dierker's (1971) bound, developed in the context of an indivisible-goods exchange economy, would be $(M-1)\sqrt{M}$ if it applied here. The substantive reason why the Starr and Dierker bounds cannot be applied or adapted to this environment is that approximately equal incomes need not be well defined in exchange economies with indivisibilities: Starr's economy allows for equal endowments but not indivisibilities; Dierker's allows for indivisibilities but not approximately equal endowments. That is why I use a Fisher economy, in which agents are directly endowed with budgets of the artificial currency.

While perfect market clearing would obviously be preferable, there are at least two reasons to think that a small amount of error is not especially costly in the context of course allocation.[14] First, a course's capacity should trade off the benefits and costs of allowing in additional students: more students get to enjoy the class, but all students get less attention from the professor. An envelope theorem argument suggests that at the optimal capacity the social costs of adding or removing a marginal student are small. Second, most universities allow students to adjust their schedules during the first week or so of classes in an "add-drop" period. A small amount of market-clearing error in the primary market can be corrected in this secondary market.

D. *Sketch of Proof of Theorem 1*

The proof of theorem 1 is contained in Appendix A. Here I provide a detailed sketch. Some readers may wish to skip directly to Section IV.

1. Demand Discontinuities

The basic difficulty for existence is that agents' demands are discontinuous with respect to price; standard Arrow-Debreu-McKenzie existence results assume that preferences are continuous. The role of the parameter σ is that $\sqrt{\sigma}$ is an upper bound on the magnitude of any discontinuity in any single agent's demand. At worst, a small change in price can cause an agent's demand to change from one bundle of objects to an entirely disjoint bundle of objects. Since there are at most k objects in a permissible bundle and M types of objects overall, this discontinuity

[14] For other contexts in which market-clearing error is more costly, see Budish (2010) for two variants of the proposed mechanism that clear the market without error. One approach is to increase budget inequality until an exact competitive equilibrium is found. A second approach is to find an approximate competitive equilibrium with excess supply but not excess demand and then execute Pareto-improving trades to eliminate the excess supply.

involves at most $\min(2k, M) \equiv \sigma$ objects and has Euclidean distance of at most $\sqrt{\sigma}$.[15]

Consider the diamonds and rocks example from the introduction. A small increase in the price of the big diamond might cause an agent who no longer can afford the bundle {big diamond, ugly rock} to instead demand the bundle {small diamond, pretty rock}.[16]

2. The Role of Unequal Budgets

The role of unequal budgets is to mitigate how individual agent demand discontinuities aggregate up into aggregate demand discontinuities.

If agents have the same budgets, their demand discontinuities occur at the same points in price space. It is possible that the magnitude of the discontinuity in aggregate demand is as large as $N\sqrt{\sigma}$. For instance, imagine that some small change in price causes all N agents to change their demands simultaneously from {big diamond, ugly rock} to {small diamond, pretty rock}.

If agents have distinct budgets, then it becomes possible to change one agent's choice set without changing all agents' choice sets. This is the basic intuition for why even an arbitrarily small amount of budget inequality is so useful.

The formal way to represent how choice sets change as prices change is using what I will call *budget constraint hyperplanes*. Let $H(i, x) = \{\mathbf{p} : \mathbf{p} \cdot x = b_i\}$ denote the hyperplane in M-dimensional price space along which agent i can exactly afford bundle x. As prices cross $H(i, x)$ from below, bundle x goes from being affordable for i to unaffordable for i.

Importantly, the number of such hyperplanes is finite because the number of agents and the number of bundles are finite. This is an advantage of having only indivisible goods. As long as each agent has

[15] More generally, we can define

$$\sqrt{\sigma} \equiv \sup_{i,\mathbf{p}} \lim_{\delta \to 0^+} \sup_{\mathbf{p}' \in B_\delta(\mathbf{p})} \| x_i^*(\mathbf{p}) - x_i^*(\mathbf{p}') \|_2,$$

with

$$x_i^*(\mathbf{p}) \equiv \max_{\succsim_i} \{x' \in 2^C : \mathbf{p} \cdot x' \leq 1\}.$$

In some contexts $\sqrt{\sigma}$ defined this way will be strictly less than $\min(2k, M)$.

[16] Observe too that in this example the big diamond and ugly rock are complements: increasing the price of one reduces demand for the other. This complementarity is intrinsic to allocation problems with both indivisibilities and budget constraints (be they of fake money or real money) and is the reason that we are unable to use the monotone price path techniques that have been successful at establishing the existence of market-clearing item prices in certain combinatorial auction environments. See Parkes (2006) for a survey of monotone price path techniques, Milgrom (2000) for a well-known example, and Mongell and Roth (1986) on the relationship between budget constraints and complementarities.

a unique budget, the number of agents' hyperplanes intersecting at any one point is generically at most M, the dimensionality of the price space.[17] Now, the maximum discontinuity in aggregate demand with respect to price is $M\sqrt{\sigma}$, which no longer grows with N (see step 1 of the formal proof).

3. A Fixed Point of Convexified Excess Demand

The next step is to artificially smooth out the (mitigated) discontinuities, enabling application of a fixed-point theorem to artificially convexified aggregate demand. Consider a traditional tâtonnement price adjustment function of the form

$$f(\mathbf{p}) = \mathbf{p} + \mathbf{z}(\mathbf{p}), \quad (1)$$

where $\mathbf{z}(\mathbf{p})$ indicates excess demand. If $f(\cdot)$ had a fixed point, this point would be a competitive equilibrium price vector (step 2). Next consider the following convexification of $f(\cdot)$:

$$F(\mathbf{p}) = co\{y : \exists \text{ a sequence } \mathbf{p}^w \to \mathbf{p},\ \mathbf{p}^w \neq \mathbf{p} \text{ such that } f(\mathbf{p}^w) \to y\}, \quad (2)$$

where co denotes the convex hull. The correspondence $F(\cdot)$ smooths out discontinuities at budget constraint hyperplanes. If aggregate demand is $\mathbf{x}'$ on one side of a discontinuity and $\mathbf{x}''$ on the other, then on the point of discontinuity itself $F(\cdot)$ maps to the set of convex combinations of $\mathbf{x}'$ and $\mathbf{x}''$.

Cromme and Diener (1991, lemma 2.4) show that for any map $f(\cdot)$ on a compact and convex set, correspondences of the form (2) are upper hemicontinuous. To apply the Cromme and Diener result, we need to specify the set on which $f(\cdot)$ is defined. In a traditional Arrow-Debreu-McKenzie setting, this would be difficult because excess demand for a good goes to infinity as its price approaches zero. However, in our setting excess demand is bounded because each student requires at most one seat in each course; formally, our setting violates the standard nonsatiation assumption. In fact, the main boundary concern in the proof arises in the case in which excess demand for good j is negative at prices near zero; hence $f_j(\mathbf{p})$ can be negative. The proof handles these issues by defining two price spaces: a legal price space $\mathcal{P} = [0, 1 + \beta + \varepsilon]^M$ and an auxiliary enlargement of this space, which is also convex and compact, on which $f(\cdot)$ and $F(\cdot)$ are defined. For the remainder of the proof sketch, we will ignore the distinction between the legal and the auxiliary price spaces.

Once we have upper hemicontinuity of $F(\cdot)$ on a convex and compact

[17] A perturbation scheme ensures that there are in fact no L-way intersections with $L > M$. This perturbation scheme is the reason theorem 1 allows $\mathbf{b}^*$ to differ from $\mathbf{b}'$ pointwise by $\varepsilon > 0$.

set, we can apply Kakutani's fixed-point theorem (the other conditions are trivially satisfied): there exists $\mathbf{p}^*$ such that $\mathbf{p}^* \in F(\mathbf{p}^*)$ (step 3). In words, $\mathbf{p}^* \in F(\mathbf{p}^*)$ tells us that there exists a set of prices arbitrarily close to $\mathbf{p}^*$ such that a convex combination of their demands exactly clears the market.

At this point we could apply theorem 2.1 of Cromme and Diener (1991) to obtain a price vector that clears the market to within error of $M\sqrt{\sigma}$. The purpose of the remainder of the proof is to tighten the bound to $\sqrt{\sigma M}/2$.

4. Mapping from Price Space to Demand Space Near $\mathbf{p}^*$

This step maps from an arbitrarily small neighborhood of $\mathbf{p}^*$ in price space to the actual excess demands associated with these prices in excess demand space. This map is the key to tightening the bound.

Because agents' demands change only when price crosses one of their budget constraint hyperplanes, we can put a lot of structure on demands in a neighborhood of $\mathbf{p}^*$. If $\mathbf{p}^*$ is not on any of the hyperplanes, then in a small enough neighborhood of $\mathbf{p}^*$ demand is unchanging, and so $\mathbf{p}^* \in F(\mathbf{p}^*)$ actually implies $\mathbf{p}^* = f(\mathbf{p}^*)$, and we are done (step 4). Suppose instead that $\mathbf{p}^*$ is on $L \leq M$ of the hyperplanes.

The two key ideas for building the map are as follows. First, for any price $\mathbf{p}'$ in a small enough neighborhood of $\mathbf{p}^*$, demand at $\mathbf{p}'$ is entirely determined by which side of the L hyperplanes $\mathbf{p}'$ is on: the affordable side or the unaffordable side. That is, out of a whole neighborhood of $\mathbf{p}^*$, we can limit attention to a finite set of at most 2^L points (steps 5 and 6).

Second, for each of the L agents corresponding to the L hyperplanes, their demand depends only on which side of their own hyperplane price is on. For each of the L agents we can define a change-in-demand vector $v_l \in \{-1, 0, 1\}^M$ that describes how their demand changes as price crosses from the affordable to the unaffordable side of their hyperplane.[18] Thus, demand in a neighborhood of $\mathbf{p}^*$ is described by at most 2^L points, which themselves are described by at most L vectors. Since $\mathbf{p}^*$ itself is on the affordable side of each hyperplane (weakly), the set of feasible demands in an arbitrarily small neighborhood of $\mathbf{p}^*$ can be written as (step 7)

[18] There are two exceptions to this statement that are handled in the proof. The first exception is if $\mathbf{p}^*$ is on the boundary of legal price space. In this case we may need to perturb budgets a tiny bit more in order to cross certain combinations of hyperplanes. The second exception is if multiple intersecting hyperplanes belong to a single agent. Then the agent's change in demand close to $\mathbf{p}^*$ is a bit more complicated than can be described by a single change-in-demand vector, which is bad for the bound. But, there will be fewer total agents to worry about, which is good for the bound. The latter effect dominates.

$$\left\{(a_1, \ldots, a_L) \in \{0, 1\}^L : \mathbf{z}(\mathbf{p}^*) + \sum_{l=1}^{L} a_l v_l\right\}. \tag{3}$$

5. Obtaining the Bound

Now $\mathbf{p}^* \in F(\mathbf{p}^*)$ tells us something very useful: perfect market clearing is in the convex hull of (3). Our market-clearing error is the maximum-minimum distance between a vertex of (3)—one of the feasible demands near $\mathbf{p}^*$—and a point in the convex hull of (3). This worst-case distance occurs when (3) is an M-dimensional hypercube of side length $\sqrt{\sigma}$, and the perfect market-clearing ideal is equidistant from all 2^M vertices of (3). Half the diagonal length of such a hypercube is $\sqrt{\sigma M}/2$ (steps 8 and 9).

E. Tightness of Theorem 1

The bound of theorem 1 is tight in the following sense.

PROPOSITION 1. For any M' there exists an economy with $M \geq M'$ object types such that, for some $\beta > 0$, there does not exist an (α, β)–CEEI for any $\alpha < \sqrt{\sigma M}/2$.

The proof, contained in Appendix B, presents an example for which the bound is tight and then describes how to construct arbitrarily large versions of the example.

F. Analogue to the First Welfare Theorem

For completeness I provide the analogue to the first welfare theorem for approximate as opposed to exact competitive equilibria.

PROPOSITION 2. Let $[\mathbf{x}^*, \mathbf{b}^*, \mathbf{p}^*]$ be an (α, β)–CEEI of economy $(\mathcal{S}, \mathcal{C}, (q_j)_{j=1}^M, (\Psi_i)_{i=1}^N, (\succsim_i)_{i=1}^N)$. For each $j \in \mathcal{C}$, let $q_j^* = \sum_i x_{ij}^*$ if $p_j^* > 0$ and $q_j^* = \max(\sum_i x_{ij}^*, q_j)$ if $p_j^* = 0$. The allocation $\mathbf{x}^*$ is Pareto efficient in economy $(\mathcal{S}, \mathcal{C}, (q_j^*)_{j=1}^M, (\Psi_i)_{i=1}^N, (\succsim_i)_{i=1}^N)$.

A practical implication of proposition 2 is that the allocation induced by an approximate CEEI will not admit any Pareto-improving trades among the agents but may admit Pareto-improving trades among sets of agents and the market administrator.

IV. Criteria of Outcome Fairness

Indivisibilities complicate fair division. If there are two agents and two indivisible objects—say a valuable diamond and an ordinary rock—then one of the agents will be left with just the rock. Some outcome unfairness is inevitable.[19]

There have been several previous approaches to outcome fairness in the presence of indivisibilities. First, many authors consider a simpler problem in which monetary transfers are permitted (Svensson 1983; Maskin 1987; Alkan, Demange, and Gale 1991). This makes it possible to run an auction in which the high bidder gets the diamond but pays a transfer to the low bidder to ensure outcome fairness. A second approach is that of Brams and Taylor (1999), who assume that indivisible objects are actually divisible in a pinch; this may be a reasonable assumption in the context of complex multi-issue negotiations. Of course, it is easier to divide a diamond and a rock if the diamond can be cut in half without loss of value. A third approach is to assess criteria of outcome fairness at an interim stage, after preferences have been reported but before the resolution of some randomness (Hylland and Zeckhauser 1979; Bogomolnaia and Moulin 2001). If we award each agent the lottery in which he receives each object with probability one-half, then neither agent envies the other agent's lottery.

The common thread in all these approaches is that by modifying either the problem or the time at which fairness is assessed, it becomes possible to use traditional criteria of outcome fairness.[20] I take a different approach. I keep my problem as is and assess outcome fairness ex post, but I weaken the criteria themselves to accommodate indivisibilities in a realistic way. Specifically, I weaken the fair-share guarantee (Steinhaus 1948) and envy-freeness (Foley 1967)—which Moulin (1995, 166) describes as "the two most important tests of equity"—and propose the *maximin share guarantee* and *envy bounded by a single good*, respectively.

A. *The Maximin Share Guarantee*

In a divisible-goods fair-division problem, an agent is said to receive his *fair share* if he receives a bundle he likes at least as well as his per capita share of the endowment. Formally, if $\mathbf{q} \in \mathbb{R}_+^M$ is an endowment of di-

[19] Procedural fairness is not similarly problematic: tossing a fair coin to determine which agent gets the diamond satisfies the standard procedural fairness requirement of anonymity, also known as symmetry or equal treatment of equals. A mechanism violates anonymity if its treatment of agents depends on not only their reports but their identities; see Moulin (1995) for a formal definition. See also Moulin (2004, chap. 1) for more on the distinction between outcome fairness and procedural fairness, which are sometimes called "end state justice" and "procedural justice."

[20] Another alternative is to ignore outcome fairness altogether and look solely to procedural fairness. Ehlers and Klaus (2003) take this approach to argue that dictatorships are fair for combinatorial assignment: "Dictatorships can be considered to be 'fair' if the ordering of agents is fairly determined" (266).

visible goods, an allocation $\mathbf{x}$ satisfies the fair-share guarantee if $x_i \succsim_i \mathbf{q}/N$ for all i (Steinhaus 1948).

With indivisibilities, fair share is not well defined: $\mathbf{q}/N$, with $\mathbf{q} \in \mathbb{Z}_+^M$, may not be a valid consumption bundle. I propose the following weakening of the fair-share ideal.

Definition 2. Agent i's *maximin share* is the consumption bundle

$$\max_{(\succsim_i)} \left\{ \min_{(\succsim_i)} \{x_1, \ \ldots, \ x_N\} \right\}, \tag{4}$$

where the $\max_{(\succsim_i)}\{\cdot\}$ is taken over all allocations $(x_l)_{l=1}^N$ such that $x_l \in 2^C$ for all l and $\sum_l x_{lj} \leq q_j$ for all j. A *maximin split* of agent i is any allocation that maximizes (4). Any allocation in which all N agents get a bundle they weakly prefer to their own maximin share is said to satisfy the *maximin share guarantee.*

The maximin share can be interpreted as the outcome of a divide-and-choose procedure against adversaries, in which an agent divides the endowment into N bundles such that his least favorite bundle is as attractive as possible (cf. Crawford 1977). The maximin share can also be interpreted as a Rawlsian guarantee from behind what Moulin (1991) calls a "thin veil of ignorance." The agent knows his own preferences and knows what resources are available to be divided (this is what makes the veil "thin") but does not know other agents' preferences.

Maximin shares are equivalent to fair shares in divisible-goods economies with convex and monotonic preferences; for a formal proof, see Budish (2010).

B. *Envy Bounded by a Single Good*

An allocation $\mathbf{x}$ is said to be *envy-free* if $x_i \succsim_i x_{i'}$ for all $i, i' \in \mathcal{S}$ (Foley 1967). In words, envy-freeness requires that each agent likes his own bundle weakly better than anyone else's.

In contrast to the fair-share guarantee, envy-freeness is perfectly well defined in the presence of indivisibilities. Its difficulty is that it is unrealistic: if there is a single diamond, then whichever agent receives it will be envied by the other. But we can take advantage of the fact that bundles of indivisible objects are somewhat divisible. If there are two diamonds, then an allocation in which some agent gets both creates more envy than is necessary given the level of indivisibility in the economy. I propose the following weakening of the envy-free test.

Definition 3. An allocation $\mathbf{x}$ satisfies *envy bounded by a single good* if, for any $i, i' \in \mathcal{S}$, either (i) $x_i \succsim_i x_{i'}$ or (ii) there exists some good $j \in x_{i'}$ such that $x_i \succsim_i (x_{i'} \backslash \{j\})$.

In words, if agent i envies agent i', we require that by removing some single good from i''s bundle we can eliminate i's envy.

C. *Dictatorships and Fairness*

In combinatorial assignment problems, dictatorships violate the maximin share guarantee and envy bounded by a single good. This observation, together with theorem 1 of Klaus and Miyagawa (2001),[21] directly yields the following impossibility result.

Proposition 3. There is no combinatorial assignment mechanism that is strategyproof, is ex post efficient, and satisfies either the maximin share guarantee or envy bounded by a single good.

Note that in single-unit assignment problems, for which dictatorships are commonly observed in practice (see the references in n. 5), dictatorships actually satisfy both fairness criteria.

V. Fairness Theorems for Approximate CEEI

I provide two fairness theorems for the approximate CEEI guaranteed to exist by theorem 1. The first indicates that for small enough β, an (α, β)–CEEI guarantees an approximation to maximin shares; the second indicates that for small enough β, an (α, β)–CEEI guarantees that envy is bounded by a single good.

A. *Theorem 2: Approximate CEEI Guarantees Approximate Maximin Shares*

I begin by showing that exact CEEIs guarantee exact maximin shares. The proof is short and helps build intuition for theorem 2.

Proposition 4. If $[\mathbf{x}^*, \mathbf{b}^*, \mathbf{p}^*]$ is a (0, 0)–CEEI, then $\mathbf{x}^*$ satisfies the maximin share guarantee.

Proof. Let $\mathbf{x}^{MS}$ denote a maximin split for agent i. Suppose that $x_l^{MS} \succ_i x_i^*$ for each $x_l^{MS} \in \mathbf{x}^{MS}$. By conditions i and iii of definition 1 we have $\mathbf{p}^* \cdot x_l^{MS} > b_i^*$ for each $x_l^{MS} \in \mathbf{x}^{MS}$ and $\mathbf{p}^* \cdot x_l^* \leq b_i^*$ for each $x_l^* \in \mathbf{x}^*$, respectively. But by condition ii of definition 1, any object that has positive price under $\mathbf{p}^*$ is at full capacity under $\mathbf{x}^*$, so $\mathbf{x}^{MS}$ cannot cost more in total than $\mathbf{x}^*$. This yields a contradiction:

[21] Theorem 1 of Klaus and Miyagawa (2001) says that the serial dictatorship is the only mechanism that is strategyproof and ex post efficient for the case of $N = 2$ agents, which is enough to yield proposition 3 as stated. If we restrict attention to larger economies (i.e., $N > 2$), then we can obtain a slightly weaker statement than proposition 3 by using either proposition 1 of Pápai (2001) or theorem 1 of Ehlers and Klaus (2003), each of which characterizes dictatorships in terms of strategyproofness, ex post efficiency, and a mild additional property (nonbossiness and coalitional strategyproofness, respectively).

$$Nb_i^* \geq \sum_{x_l^* \in \mathbf{x}^*} \mathbf{p}^* \cdot x_l^* \geq \sum_{x_l^{MS} \in \mathbf{x}^{MS}} \mathbf{p}^* \cdot x_l^{MS} > Nb_i^*.$$

QED

The proof of proposition 4 relies on two facts about (0, 0)–CEEIs: (i) $\beta = 0$ means that each agent has $1/N$ of the budget endowment, and (ii) $\alpha = 0$ means that at price vector $\mathbf{p}^*$ the goods endowment costs weakly less than the budget endowment.

The approximate CEEI jeopardizes both of these properties. Setting $\beta > 0$ but sufficiently small minimizes the harm from violating condition i. Issue ii is handled with the following approximation parameter and minor extension of theorem 1.

DEFINITION 4. For $\delta \geq 0$ and budgets $\mathbf{b}$, the set $\mathcal{P}(\delta, \mathbf{b})$ is defined to be the set of price vectors at which the goods endowment costs at most δ proportion more than the budget endowment. Formally,

$$\mathcal{P}(\delta, \mathbf{b}) = \left\{\mathbf{p} \in [0, \max_i (b_i)]^M : \sum_j p_j q_j \leq \sum_i b_i(1 + \delta)\right\}.$$

LEMMA 1. For any $\delta > 0$ and any set of target budgets $\mathbf{b}'$, there exists an (α, β)–CEEI $[\mathbf{x}^*, \mathbf{b}^*, \mathbf{p}^*]$ that satisfies all of the conditions of theorem 1 and additionally $\mathbf{p}^* \in \mathcal{P}(\delta, \mathbf{b}^*)$.

By choosing β, δ small enough, we can ensure that each agent's budget is at least $1/(N+1)$ of the cost of the endowment at the approximate CEEI price vector $\mathbf{p}^*$. This guarantees that the approximate CEEI approximately satisfies the maximin share guarantee.

DEFINITION 5. Agent i's $N+1$ *maximin share* is the consumption bundle

$$\max_{(\succsim_i)} \left\{\min_{(\succsim_i)} \{x_1, \ldots, x_N, x_{N+1}\}\right\},$$

where the $\max_{(\succsim_i)} \{\cdot\}$ is taken over all allocations $(x_l)_{l=1}^{N+1}$ such that $x_l \in 2^C$ for all l and $\sum_l x_{lj} \leq q_j$ for all j.

THEOREM 2. *If* $[\mathbf{x}^*, \mathbf{b}^*, \mathbf{p}^*]$ *is an* (α, β)–*CEEI where, for some* $\delta \geq 0$, $\mathbf{p}^* \in \mathcal{P}(\delta, \mathbf{b}^*)$ *and* $\beta < (1 - \delta N)/N(1 + \delta)$, *then* $\mathbf{x}^*$ *satisfies the* $(N+1)$–*maximin share guarantee.*

Theorem 2 indicates that approximate CEEI guarantees each agent his maximin share based on a hypothetical economy in which there are $N+1$ instead of N total participants. In fact, a slightly better approximation to maximin shares can be provided in case $\alpha = 0$. In particular, in the diamonds and rocks example from the introduction we can guarantee that each of the two agents gets a diamond; the approximation is that the agent who gets the small diamond may also get the ugly rock. For details, see Budish (2010).

B. *Theorem 3: Approximate CEEI Guarantees Envy Bounded by a Single Good*

Exact CEEIs are envy-free because all agents have the same choice set. When agents have unequal incomes ($\beta > 0$), they have different choice sets, so envy-freeness cannot be assured. The following result shows that if the inequality in budgets is sufficiently small, we can bound the degree of envy.

THEOREM 3. If $[\mathbf{x}^*, \mathbf{b}^*, \mathbf{p}^*]$ is an (α, β)–CEEI with $\beta \leq 1/(k-1)$, where k is the maximum number of objects per agent as defined in theorem 1, then $\mathbf{x}^*$ satisfies envy bounded by a single good.

Proof. Suppose for a contradiction that i envies i' and that this envy is not bounded by a single good. Let $k' \leq k$ denote the number of objects in the envied bundle $x^*_{i'}$, and number these objects $j_1, \ldots, j_{k'}$. Condition i of definition 1 indicates that i cannot afford any of the k' bundles formed by removing a single object from $x^*_{i'}$:

$$\mathbf{p}^* \cdot (x^*_{i'} \backslash \{j_1\}) > b^*_i,$$
$$\vdots$$
$$\mathbf{p}^* \cdot (x^*_{i'} \backslash \{j_{k'}\}) > b^*_i.$$

Since $\{j_1\} \cup \{j_2\} \cup \cdots \cup \{j_{k'}\} = x^*_{i'}$, we can sum these inequalities to obtain

$$(k'-1)(\mathbf{p}^* \cdot x^*_{i'}) > k' b^*_i,$$

which, since $b^*_{i'} \geq (\mathbf{p}^* \cdot x^*_{i'})$ by the fact that i' can afford her bundle, gives

$$(k'-1) b^*_{i'} > k' b^*_i.$$

Since $k' \leq k$, we have

$$\frac{b^*_{i'}}{b^*_i} > \frac{k'}{k'-1} \geq \frac{k}{k-1} \geq 1 + \beta,$$

which contradicts condition iii of definition 1. QED

VI. The Approximate CEEI Mechanism

On the basis of the efficiency and fairness results of Sections III–V and with an eye toward incentives as discussed below, I propose the approximate CEEI mechanism.

MECHANISM 1 (Approximate CEEI mechanism).

1. Each student i reports her preferences $\succ_i$ over her permissible schedules Ψ_i. The market administrator sets $\emptyset \succ_i x$ for $x \notin \Psi_i$.

2. Assign each student i a budget b_i^* that is a uniform random draw from $[1, 1 + \beta]$, with

$$0 < \beta < \min\left(\frac{1}{N}, \frac{1}{k-1}\right).$$

3. Compute a set of prices $(p_1^*, \ldots, p_M^*)$ and allocations $(x_1^*, \ldots, x_N^*)$, in an anonymous manner, such that (a) each student's allocation maximizes her utility, on the basis of her reported preferences, subject to her budget constraint (formally, condition i of definition 1 is satisfied); (b) the magnitude of market-clearing error, as defined in condition ii of definition 1, is smaller than the theorem 1 bound of $\sqrt{\sigma M}/2$.[22]
4. Announce $[\mathbf{x}^*, \mathbf{b}^*, \mathbf{p}^*]$.

In step 1, students report their preferences over permissible schedules. In practice, the number of possible schedules can be quite large, so the market administrator must provide a language by which students can express their preferences concisely. Othman et al. (2010) propose one such language, the starting point of which is the observation that if i's preferences are additively separable but for scheduling constraints, M numbers can be used to describe ordinal preferences over the entire set Ψ_i.

In step 2, students are randomly assigned approximately equal budgets. In practice, the market administrator may wish to assign budgets nonrandomly, for example, on the basis of seniority or grade point average.

In step 3, the mechanism computes approximate market-clearing prices. Othman et al. (2010) propose one such computational procedure that appears to work well in life-size problems; in particular, in both the data in Section VIII and simulation economies, it typically finds prices that clear the market with error that is much smaller than the theorem 1 bound. However, it remains an important open question what is the best such computational procedure, both in terms of minimizing market-clearing error and in terms of computational efficiency.

Finally, in step 4, allocations, budgets, and prices are announced publicly. This step relates to the issue of "transparency," which is an important if imprecisely defined concern in practical market design. The computation of A-CEEI prices is necessarily somewhat opaque since it involves calculating approximate fixed points. But, once computed, the prices can then be announced publicly. The assignment of random budgets can be public as well. Perhaps the most important sense in which A-CEEI is transparent is that students can verify that they received

[22] Technically, it may be necessary to perturb budgets by some $\varepsilon > 0$ to guarantee an error smaller than $\sqrt{\sigma M}/2$. One other technicality is that prices should satisfy $\mathbf{p}^* \in \mathcal{P}((1 - N\beta)/(N + N\beta), \mathbf{b}^*)$ as per lemma 1 and theorem 2.

their most preferred affordable schedules at the publicly announced prices.

Incentive properties of the approximate CEEI mechanism.—A-CEEI is not strategyproof because agents' reports might affect prices. But, in a large-market limit economy in which agents are price takers, it is a dominant strategy to report truthfully. Formally, A-CEEI has the property of being *strategyproof in the large*, as defined in Budish (2010).[23] The intuition is straightforward: by reporting truthfully, the agent is sure to receive her most preferred affordable bundle at the realized prices; if she cannot affect prices, she might as well tell the truth.

By contrast, the course-allocation mechanisms currently found in practice not only are not strategyproof but also fail to be strategyproof in the large. In the bidding points auction studied by Sönmez and Ünver (2003, 2010), students should misreport their preferences even if they cannot affect prices. In the HBS draft mechanism studied by Budish and Cantillon (forthcoming), students should misreport their preferences even if they cannot affect course run-out times, which are that mechanism's analogue of prices.

Empirical evidence suggests that large-market strategyproofness is important in practice. In addition to the two course-allocation mechanisms mentioned above, other mechanisms that fail to be strategyproof in the large and that have been shown to have important incentive problems in practice include nonstable matching algorithms (cf. Roth 2002), the Boston mechanism for school choice (Abdulkadiroğlu, Pathak, and Roth 2005), and discriminatory price multiunit auctions (e.g., Friedman 1991). Examples of mechanisms that are strategyproof in the large and that are thought to have attractive incentive properties in practice include deferred acceptance algorithms and double auctions.[24] To the best of my knowledge, there are no empirical examples of market designs that are strategyproof in the large but have been shown to be harmfully manipulated in large finite markets, nor are there empirical examples of market designs that are manipulable in the large but are thought to be truthful in large finite markets.

[23] The Budish (2010) definition of strategyproof in the large is based on the continuum replication of a given finite economy. See also Azevedo and Budish (2012) for a definition based on the limit of a sequence of finite economies.

[24] For both deferred acceptance and double auctions (e.g., Perry and Reny 2006; Kojima and Pathak 2009) and Walrasian mechanisms with perfectly divisible goods (e.g., Roberts and Postlewaite 1976), we have a formal theoretical understanding of how truthful behavior converges to optimal behavior as the market grows large. Unfortunately, the kinds of proof techniques that have been developed for these other contexts are not readily applicable here because the relationship between agents' reports and realized prices is nonconstructive and discontinuous, as well as somewhat random. Developing a better understanding of A-CEEI's incentive properties away from the limit is a natural topic for future research.

VII. Comparison to Alternate Mechanisms

Impossibility theorems indicate that there is no perfect mechanism for combinatorial assignment. The approximate CEEI mechanism developed in Sections III–VI offers a compromise of competing design objectives. It is approximately ex post efficient, guarantees approximate maximin shares, bounds envy by a single good, and is strategyproof in the large. Additionally, it satisfies the procedural fairness requirement of anonymity.

One way to assess whether A-CEEI constitutes an attractive compromise is to compare its properties with those of alternatives. Table 1 describes the efficiency, fairness, and incentives properties of every prior combinatorial assignment mechanism I am aware of from either theory or practice.

Every other mechanism but for A-CEEI is either severely unfair ex post or manipulable in the large, and most are both unfair and manipulable. Many of the mechanisms are ex post Pareto efficient under truthful play, whereas A-CEEI is only approximately efficient. However, most of these mechanisms are manipulable even in large markets, and most also restrict the preference information agents can report. So it is difficult to say which will be more efficient in practice. In the one case we are able to test with data (in Sec. VIII), A-CEEI is the more efficient mechanism, both ex ante and ex post.

Comparison to the bidding points auction.[25]—Of special note among alternate mechanisms is the widely used bidding points auction (BPA) since it resembles exact CEEI.[26] The BPA works roughly as follows: Each student submits integer bids for individual classes, the sum of their bids not to exceed some fixed budget amount (say 10,000 points). If course j has q_j seats, the q_j highest bidders for it get a seat, with ties broken randomly.[27] Student i's bids can be interpreted as a report of an additively separable utility function, $u_i(\cdot)$. The q_jth-highest bid for course j is frequently interpreted as the "clearing price" p_j for course j and the allocation itself as a "market equilibrium" (e.g., Wharton 2011).

[25] See Budish (2010) for two additional detailed comparisons, to the random serial dictatorship and the Hylland-Zeckhauser pseudomarket mechanism.

[26] Variants of the BPA are used at Berkeley, Chicago, Columbia, MIT, Michigan, NYU, Northwestern, Penn, Princeton, and Yale. See Sönmez and Ünver (2003, 2010), Adler et al. (2008), MIT (2008), https://ibid.chicagobooth.edu/registrar-student/Home.tap, and Wharton (2011) for details.

[27] More precisely, bids for all courses are sorted in descending order and are either filled or rejected one at a time depending on whether (i) the course still has capacity for the student and (ii) the student still has capacity for the course. Because of condition ii, a student whose bid for course j is among the q_j highest might not get it whereas some other student who bids less does. Strategic issues aside, condition ii can lead to inefficient allocations. Sönmez and Ünver (2003, 2010) and Krishna and Ünver (2008) propose a mechanism that eliminates the inefficiencies that arise from this specific aspect of the BPA.

TABLE 1
COMPARISON OF ALTERNATIVE MECHANISMS

Mechanism	Efficiency (Truthful Play)	Fairness (Truthful Play)	Incentives	Preference Language
A-CEEI mechanism	Pareto efficient with respect to allocated goods Worst-case allocation error is small for practice and goes to zero in the limit	$N+1$—maximin share guaranteed Envy bounded by a single good	Strategyproof in the large	Ordinal over schedules
Mechanisms from practice:				
Bidding points auction (Sonmez and Ünver 2003)	If preferences are additive separable, Pareto efficient but for quota issues	Worst case: get zero goods	Manipulable in the large	Cardinal over items
HBS draft mechanism (Budish and Cantillon, forthcoming)	If preferences are responsive, Pareto efficient with respect to the reported information	If preferences are responsive and $k = 2$, maximin share guaranteed If preferences are responsive, envy bounded by a single good	Manipulable in the large	Ordinal over items
Univ. Chicago primal-dual linear program mechanism (Graves, Schrage, and Sankaran 1993)	Pareto efficient when preference-reporting limits do not bind	Worst case: get zero goods	Manipulable in the large	Cardinal over a limited number of schedules
Mechanisms from prior theory:				
Adjusted winner (Brams and Taylor 1996)	If preferences are additive-separable, Pareto efficient	Worst case: get zero goods	Manipulable in the large	Cardinal over items

Descending demand procedure (Herreiner and Puppe 2002)	Pareto efficient	Does not satisfy maximin share guarantee or envy bounded by a single good	Manipulable in the large	Ordinal over schedules
Gale-Shapley enhancement to the BPA (Sönmez and Ünver 2003)	If preferences are additive separable, Pareto efficient	Worst case: get zero goods	Bidding phase: manipulable in the large	Bidding phase: cardinal over items
			Allocation phase: strategyproof in the large	Allocation phase: ordinal over items
Geometric prices mechanism (Pratt 2007)	If von Neumann–Morgenstern preferences are additive separable, Pareto efficient	Worst case: get zero goods	If von Neumann–Morgenstern preferences are additive separable, strategyproof in the large	Cardinal over items
Minimize envy algorithm (Lipton et al. 2004)	Algorithm ignores efficiency	If preferences are additive separable, envy bounded by a single good	Manipulable in the large	Cardinal over schedules
Serial/sequential dictatorship (cf. Pápai 2001)	Pareto efficient	Worst case: get *k* worst goods	Strategyproof	Ordinal over schedules
Other solution concepts:				
Maximin utility	Pareto efficient	Worst case: get approximately zero goods (if a hedonist and all other agents are depressives)	Manipulable in the large	Cardinal over schedules
Utilitarian solution	Pareto efficient	Worst case: get zero goods (if a depressive and all other agents are hedonists)	Manipulable in the large	Cardinal over schedules

To the casual observer, the BPA looks like an equal-incomes competitive equilibrium procedure. Yet it turns out that the BPA makes a subtle mistake: *the BPA treats fake money as if it were real money that enters the utility function*. That is, it treats a general equilibrium theory problem as if it were an auction theory problem.

Suppose that student i has a budget of b_i and faces a price vector of $\mathbf{p}$, and let $B_i = \{x' \in 2^C : \mathbf{p} \cdot x \leq b_i\}$. With fake money, student i's correct demand is

$$x_i^* = \arg\max_{x' \in B_i} [u_i(x')], \tag{5}$$

which is what student i would receive at these prices under A-CEEI. The BPA instead allocates student i the bundle

$$x_i^{\text{BPA}} = \arg\max_{x' \in B_i} [u_i(x') - \mathbf{p} \cdot x'], \tag{6}$$

that is, the bundle that maximizes the difference between reported utility and expenditure of fake money. This can lead to situations in which a student gets zero of the courses she bids on. Suppose that Alice bids 7,000, 2,000, and 1,000 for courses a, b, and c, but their prices are 8,000, 3,000, and 1,500. Then, since each of Alice's bids is less than the course's price, she gets none of them.

If Alice bids differently—say, submits bids of 8,001, 0, and 1,501—then she can trick the demand function (6) into behaving like the demand function (5). That is, she will receive the bundle $\{a, c\}$ that maximizes (5) for $u_{\text{Alice}} = (7{,}000, 2{,}000, 1{,}000)$ and that maximizes (6) for $\hat{u}_{\text{Alice}} = (8{,}001, 0, 1{,}501)$. This seems to solve the problem with the BPA, but in fact it just pushes the problem onto some other student: when Alice obtains a for 8,001, she causes the unlucky student who bid 8,000 for a (say Bob) now no longer to get it. Results in Budish (2010) formalize that the BPA is not a backdoor method for implementing true competitive equilibria and that the BPA can lead to outcomes, both under truthful play and in Nash equilibrium, in which some students get zero courses.

The zero courses issue is not just a theoretical curiosity but manifests in some simple data provided by the University of Chicago's Booth School of Business, which recently adopted a BPA. During the four quarters from summer 2009 to spring 2010, the numbers of students allocated zero courses in the main round of bidding have been 53, 37, 64, and 17.[28] In Chicago Booth's BPA, budget that is unspent in one quarter carries over to future quarters, so these numbers should be interpreted with some caution. The cleanest evidence comes from fo-

[28] Students allocated zero courses in the main round can fill their schedules in subsequent rounds of bidding, typically with courses that were in excess supply in the main round.

cusing on students who get zero courses in their last term. Of the 17 students who got zero courses in spring 2010, five were full-time master of business administration students about to graduate. One student bears an uncanny resemblance to Alice: he bid 5,466, 5,000, 1,500, and 1 for courses that had prices of 5,741, 5,104, 2,023, and 721. Another case that is instructive is a student who bid 11,354, 3, 3, 3, and 2 for courses that then had prices of 13,266, 2,023, 1,502, 1,300, and 103. This student used essentially all of his budget in a futile attempt to get the single most expensive course. Not only did he not get the "big diamond," but he also did not get a small diamond or even any rocks.[29] Theorems 2 and 3 ensure that such outcomes never occur under A-CEEI.

VIII. Performance of A-CEEI in an Empirical Environment

This section examines the performance of the approximate CEEI mechanism in a specific course-allocation environment. I use Budish and Cantillon's (forthcoming) data from course allocation at Harvard Business School (HBS) and Othman et al.'s (2010) computational procedure for A-CEEI.

A. Data and Key Assumptions

The HBS data consist of 456 students' true and stated ordinal preferences over 50 fall semester courses and 48 spring semester courses, as well as these courses' capacities, for academic year 2005–6. For details, see Budish and Cantillon (forthcoming).

To convert preferences over courses into preferences over bundles, I follow Budish and Cantillon (forthcoming) and assume that students compare bundles on the basis of the "average rank" of the courses in each bundle. For instance, a student prefers the bundle consisting of her second- and third-favorite courses to that consisting of her first and fifth because 2.5 is a lower average rank than 3.0. Ties are broken randomly. The average-rank assumption seems reasonable for handling the data incompleteness problem for two reasons: first, the HBS elective year curriculum is designed to avoid complementarities and overlap between courses; second, in the HBS draft mechanism, students are

[29] A second implication of the BPA's treatment of fake money as if it entered the utility function is that some students will graduate with large budgets of unspent fake money. Among full-time MBA students graduating in Spring 2010, the median student graduated with a budget of 6,601 unspent points, which is nearly a full quarter's budget (8,000 points). The 90th-percentile student graduated with 17,547 unspent points and the 99th with 26,675 unspent points.

unable to express the intensity of their preference for individual courses beyond ordinal rank.

I assume that students report their preferences truthfully under A-CEEI.[30] For computational reasons, I treat each semester's allocation problem separately.[31]

B. *Market-Clearing Error*

Theorem 1 indicates that there exist A-CEEI prices that clear the HBS course-allocation market to within market-clearing error of $\sqrt{2kM}/2$, where k is the number of courses per student and M is the number of courses. Here, $k = 5$ and $M = 50$, 48 for the fall and spring semesters, respectively.

I run A-CEEI 100 times for each semester and record its actual market-clearing error. The actual error is meaningfully smaller than the bound implied by theorem 1. The maximum (mean) observed error in Euclidean distance is $\sqrt{13}$ ($\sqrt{5.61}$) in the fall and $\sqrt{22}$ ($\sqrt{6.12}$) in the spring, as compared with the theorem 1 bounds of $\sqrt{125}$ and $\sqrt{120}$, respectively. In terms of seats, the maximum (mean) observed error is 11 (5.46) seats in the fall and 14 (5.96) seats in the spring.

Part of the explanation for the low amount of market-clearing error is that only a subset of courses ever have a strictly positive price: 21 in the fall and 23 in the winter. If we reformulated the problem as one of allocating only the potentially scarce courses (see Sec. III.C), this would reduce the bounds to $\sqrt{52.5}$ and $\sqrt{57.5}$, respectively.

C. *Outcome Fairness*

Theorem 2 indicates that A-CEEI guarantees students an approximation to their maximin share that is based on adding one more student to the economy. In the HBS economy, students' outcomes always exceed their exact maximin shares by a large margin. The worst outcome any student receives is an average rank of around 8, whereas students' maximin shares have an average rank of around 18–20.

[30] Unfortunately, I have no formal way of assessing whether an HBS-size economy is "large," i.e., whether truthful reporting is an approximate equilibrium. There are 50 courses per semester, and each student ranks about 15 courses per semester. So there are about $50!/(50-15)! \approx 3 \times 10^{24}$ possible reports for each student, even within the restricted class of average-rank preferences. I have no theoretically motivated way to restrict attention to some subset of these potential manipulations, in contrast to, e.g., Roth and Peranson (1999).

[31] The Othman et al. (2010) computational procedure can solve problems the size of a single semester at HBS, in which there are roughly 50 courses and $\binom{50}{5} \approx 10^6$ schedules, in around 20 minutes in Matlab on a standard workstation. It can solve problems the size of a full year at HBS, in which there are roughly 100 courses and $\binom{100}{10} \approx 10^{13}$ schedules, in around 11 hours in a more sophisticated computing environment.

Theorem 3 indicates that A-CEEI bounds each student's envy by a single good. In the data, over 100 runs of A-CEEI, around 99 percent of students have no envy; that is, they weakly prefer their own allocation to any other student's allocation. For the 1 percent of students who do envy, the degree of envy is small. The worst observed case is envy of seven course ranks, for example, a student who receives her second- to fifth- and eighth-favorite courses whereas someone else receives the student's first- to fifth-favorite courses.

D. *Ex Ante Welfare Comparison versus the HBS Draft Mechanism*

In each semester, the distribution of average ranks under A-CEEI first-order stochastically dominates that under the actual play of the HBS draft. First-order stochastic dominance is an especially strong comparison relation: we do not need to make any further assumptions on how von Neumann–Morgenstern utility responds to average rank to reach an ex ante welfare comparison.[32] Additionally, the magnitude of the improvement seems economically meaningful. The mean average ranks under A-CEEI are 4.20 in the fall and 4.32 in the spring versus 4.56 and 4.47 for HBS (lower is better; 3.00 is bliss). Thus, on average, the quality of a student's schedule improves by 0.25 ranks per course.

IX. Conclusion

Most of what is known about the combinatorial assignment problem is a series of impossibility theorems that indicate that there is no perfect solution. This paper gets around the impossibility theorems by seeking second-best approximations of the ideal properties a combinatorial assignment mechanism should satisfy. Ideally, a mechanism would be exactly Pareto efficient, both ex post and ex ante. A-CEEI is approximately ex post efficient in theory and has attractive ex ante efficiency performance in a specific empirical environment. Ideally, a mechanism would satisfy the outcome fairness criteria of envy-freeness and the maximin share guarantee. A-CEEI approximates these two ideals in theory and gives exact maximin shares and is 99 percent envy-free in the data. Ideally, a mechanism would be strategyproof. A-CEEI is strategyproof in the large, whereas the mechanisms found in practice are simple to manipulate even in large markets.

The computational analysis raises two interesting questions for future research. First, market-clearing error in the data is considerably smaller

[32] I have assumed that students' ordinal preferences over bundles are based on the average rank of the courses contained in each bundle. I have not made any additional assumption about how their von Neumann–Morgenstern utilities depend on average rank.

than the theorem 1 worst-case bound, consisting, on average, of just a single seat in six courses. Can we improve the bound of theorem 1 if we restrict attention to certain classes of preferences or make assumptions about the degree of preference heterogeneity? Second, envy in the data is exceedingly rare, whereas we know that in the worst case all agents but one will have envy. What are the features of an environment that make average-case envy small?

Two other interesting questions are raised by considering combinatorial assignment's relationship to other well-known market design problems. First, there may be an interesting hybrid problem combining combinatorial assignment with two-sided matching, just as the school choice problem is often formulated as a hybrid between single-unit assignment and two-sided matching (Abdulkadiroğlu and Sönmez 2003). For instance, in the context of course allocation, schools may wish to give course-specific priority to students who need a certain course to fulfill a requirement or who performed well in a related prerequisite. It would be interesting to see if the competitive equilibrium approach can be adapted to such environments. Second, there may be an interesting hybrid problem combining combinatorial assignment with combinatorial auctions. In a sense, combinatorial assignment is like a combinatorial auction in which all participants have a real-money budget constraint of zero, so it becomes important to use an artificial currency instead. It would be interesting to ask whether there are useful ways to combine real-money market designs and fake-money market designs in environments in which budget constraints are nonzero but often bind, or in which monetary transfers are restricted in other ways.

I close on a methodological note. Practical market design problems often prompt the development of new theory that enhances and extends old ideas. To give a prominent example, the elegant matching model of Gale and Shapley (1962) was not able to accommodate several complexities found in the practical design problem of matching medical students to residency positions. This problem prompted the development of substantial new theory (summarized in Roth [2002]) and a new market design described in Roth and Peranson (1999). Similarly, the beautiful theory of competitive equilibrium from equal incomes developed by Foley (1967), Varian (1974), and others is too simple for practice because it assumes perfect divisibility. This paper proposes a richer theory that accommodates indivisibilities and develops a market design based on this richer theory. I hope that, just as a concrete application renewed interest in Gale and Shapley's remarkable deferred-acceptance algorithm, this paper and its motivating application will renew interest in CEEI as a framework for market design.

Appendix A

Proof of Theorem 1

Preliminaries.—Fix an economy $(\mathcal{S}, \mathcal{C}, (q_j)_{j=1}^M, (\Psi_i)_{i=1}^N, (\succsim_i)_{i=1}^N)$, and fix $\beta > 0$ and $\varepsilon > 0$. Fix a budget vector $\mathbf{b}' = (b'_1, \ldots, b'_N)$ that satisfies $1 \leq \min_i (b'_i) \leq \max_i (b'_i) \leq 1 + \beta$.

Let $\bar{b} = 1 + \beta + \varepsilon$. Define an M-dimensional price space $\mathcal{P} = [0, \bar{b}]^M$ and an auxiliary enlargement of this space $\tilde{\mathcal{P}} = [-1, \bar{b} + 1]^M$. Define a truncation function $t: \tilde{\mathcal{P}} \to \mathcal{P}$ that takes any price vector in $\tilde{\mathcal{P}}$ and truncates each of its components to be within $[0, \bar{b}]$. Formally, for $\tilde{\mathbf{p}} \in \tilde{\mathcal{P}}$, $t(\tilde{\mathbf{p}}) = (\min[\bar{b}, \max(0, \tilde{p}_1)], \ldots, \min[\bar{b}, \max(0, \tilde{p}_M)])$.

In step 1 we will assign to each agent-bundle pair a small reverse tax $\tau_{ix} \in (-\varepsilon, \varepsilon)$ that affects agent i's cost of purchasing bundle x: at prices $\tilde{\mathbf{p}}$, i's total cost is $\tilde{\mathbf{p}} \cdot x - \tau_{ix}$.

Demand and excess demand are defined on all prices in $\tilde{\mathcal{P}}$. Agent i's demand $d_i(\cdot)$ depends on prices $\tilde{\mathbf{p}}$, her budget b_i, and her set of taxes $\tau_i \equiv (\tau_{ix})_{x \in 2^{\mathcal{C}}}$:

$$d_i(\tilde{\mathbf{p}}; b_i, \tau_i) = \max_{(\succsim_i)} \{x' \in 2^{\mathcal{C}} : \tilde{\mathbf{p}} \cdot x' \leq b_i + \tau_{ix'}\}. \tag{A1}$$

Demand is a function rather than a correspondence because of the assumption that preferences are strict over schedules. Let $\boldsymbol{\tau} \equiv (\tau_i)_{i \in \mathcal{S}}$. Excess demand $\mathbf{z}(\cdot)$ is defined by

$$\mathbf{z}(\tilde{\mathbf{p}}; \mathbf{b}, \boldsymbol{\tau}) = \left[\sum_{i=1}^{N} d_i(\tilde{\mathbf{p}}; b_i, \tau_i)\right] - \mathbf{q}. \tag{A2}$$

Note the slight difference between $\mathbf{z}(\cdot)$ as defined here and the $\mathbf{z}^*$ defined in definition 1.ii.*b*, which is that $\mathbf{z}(\cdot)$ does not distinguish between goods that have prices of zero and goods with a strictly positive price. We will suppress the $\mathbf{b}$ and $\boldsymbol{\tau}$ arguments from $d_i(\cdot)$ and $\mathbf{z}(\cdot)$ when their values are clear from the context.

Since each agent i consumes either 0 or 1 of each object j, it is without loss of generality to assume $q_j \in \{1, \ldots, N\}$, so $-N \leq z_j \leq N - 1$ for all $j \in \mathcal{C}$. The fact that excess demand is bounded is an important advantage of our environment relative to the traditional Arrow-Debreu-McKenzie environment.

For each agent $i \in \mathcal{S}$ and schedule $x \in 2^{\mathcal{C}}$, define the budget constraint hyperplane $H(i, x)$ by $H(i, x) \equiv \{\tilde{\mathbf{p}} \in \tilde{\mathcal{P}} : \tilde{\mathbf{p}} \cdot x = b_i + \tau_{ix}\}$.

Both the taxes and the enlarged price space play a role that is entirely internal to the proof. At the end we will have a price vector in $\mathcal{P}$ and set all of the taxes to zero.

Step 1: Choose a set of taxes $(\tau'_{ix})_{i \in \mathcal{S}, x \in 2^{\mathcal{C}}}$ such that

i. $-\varepsilon < \tau'_{ix} < \varepsilon$ (taxes are small);
ii. $\tau'_{ix} > \tau'_{ix'}$ if $x \succ_i x'$ (taxes favor more preferred bundles);
iii. $1 \leq \min_{i,x}(b'_i + \tau'_{ix}) \leq \max_{i,x}(b'_i + \tau'_{ix}) \leq 1 + \beta$ (inequality bound is preserved);
iv. $b'_i + \tau'_{ix} \neq b'_{i'} + \tau'_{i'x'}$ for any $(i, x) \neq (i', x')$ (no two perturbed budgets are equal);

v. there is no price $\tilde{\mathbf{p}} \in \tilde{\mathcal{P}}$ at which more than M perturbed budget constraint hyperplanes intersect.

Existence of a set of taxes $(\tau'_{ix})_{i \in \mathcal{S}, x \in 2^{\mathcal{C}}}$ satisfying parts i–v follows from the fact that the number of agents, number of schedules, and number of budget constraint hyperplanes are all finite.

Step 2: Define a tâtonnement price adjustment function f on $\tilde{\mathcal{P}}$. If f has a fixed point $\tilde{\mathbf{p}}^* = f(\tilde{\mathbf{p}}^*)$, then its truncation $\mathbf{p}^* = t(\tilde{\mathbf{p}}^*)$ is an exact competitive equilibrium price vector.

Let $\gamma \in (0, 1/N)$ be a small positive constant. Given budgets $\mathbf{b}$ and taxes $\boldsymbol{\tau}$, define $f: \tilde{\mathcal{P}} \rightarrow \tilde{\mathcal{P}}$ by

$$f(\tilde{\mathbf{p}}) = t(\tilde{\mathbf{p}}) + \gamma \mathbf{z}(t(\tilde{\mathbf{p}}); \mathbf{b}, \boldsymbol{\tau}). \tag{A3}$$

The reason we impose $\gamma < 1/N$ is to ensure that the image of f lies in $\tilde{\mathcal{P}}$.

Suppose, for budgets of $\mathbf{b}'$ and taxes of $\boldsymbol{\tau}$, that f has a fixed point $\tilde{\mathbf{p}}^* = f(\tilde{\mathbf{p}}^*)$. Then its truncation $\mathbf{p}^* = t(\tilde{\mathbf{p}}^*)$ is an exact competitive equilibrium price vector for the allocation $\mathbf{x}^*$ given by $x_i^* = d_i(\mathbf{p}^*; b_i, \tau_i)$ for all $i \in \mathcal{S}$ and budgets of $\mathbf{b}^*$ given by $b_i^* = b_i' + \tau_{ix_i^*}$ for all $i \in \mathcal{S}$. First, note that at any fixed point no individual price $p_j^* \geq \bar{b}$. Given the definition of $\bar{b}$, no agent can afford a seat in object j at price $\bar{b}$. So $\tilde{p}_j^* \geq \bar{b}$ implies $z_j(\mathbf{p}^*; \mathbf{b}', \boldsymbol{\tau}) = 0 - q_j < 0$, which contradicts $\tilde{p}_j^* \geq \bar{b}$ being part of a fixed point. Second, $p_j^* \in (0, \bar{b})$ implies that $z_j(\mathbf{p}^*; \mathbf{b}', \boldsymbol{\tau}) = 0$. Third, $p_j^* = 0$ implies that $z_j(\mathbf{p}^*; \mathbf{b}', \boldsymbol{\tau}) \leq 0$. Finally, revealed preference and requirement i of step 1 together imply that any bundle that i prefers to x_i^* costs strictly more than $b_i' + \tau_{ix_i^*}$, so each agent's demand at the budgets $\mathbf{b}^*$ with no taxes is the same as his demand at the budgets $\mathbf{b}'$ with taxes $\boldsymbol{\tau}$. Thus $z_j(\mathbf{p}^*; \mathbf{b}^*, \mathbf{0}) \leq 0$ and $z_j(\mathbf{p}^*; \mathbf{b}^*, \mathbf{0}) < 0 \Rightarrow p_j^* = 0$, as required for competitive equilibrium.

Step 3: Define an upper hemicontinuous set-valued correspondence F, which is a "convexification" of f and is guaranteed to have a fixed point by Kakutani's theorem. Let $\tilde{\mathbf{p}}^* \in F(\tilde{\mathbf{p}}^*)$ denote the fixed point and let $\mathbf{p}^* = t(\tilde{\mathbf{p}}^*)$ denote its truncation.

Fix budgets to $\mathbf{b}'$ and taxes to $\boldsymbol{\tau}'$ as described in step 1. Create the correspondence $F: \tilde{\mathcal{P}} \rightarrow \tilde{\mathcal{P}}$ as follows:

$$F(\mathbf{p}) = co\{\mathbf{y} : \exists \text{ a sequence } \mathbf{p}^w \rightarrow \mathbf{p}, \ \mathbf{p} \neq \mathbf{p}^w \in \tilde{\mathcal{P}} \text{ such that } f(\mathbf{p}^w) \rightarrow \mathbf{y}\}, \tag{A4}$$

where co denotes the convex hull. Cromme and Diener (1991, lemma 2.4) show that for any map f, the correspondence F constructed according to (A4) is upper hemicontinuous and hence has a fixed point (the other conditions for Kakutani's fixed-point theorem—F is nonempty, $\tilde{\mathcal{P}}$ is compact and convex, and $F(\mathbf{p})$ is convex—are trivially satisfied).

So there exists $\tilde{\mathbf{p}}^* \in F(\tilde{\mathbf{p}}^*)$. Let $\mathbf{p}^* = t(\tilde{\mathbf{p}}^*)$ denote its truncation. Fix $\mathbf{p}^*$ and $\tilde{\mathbf{p}}^*$ for the remainder of the proof.

Step 4: If the price vector $\mathbf{p}^*$ is not on any budget constraint hyperplane, then it is an exact competitive equilibrium price vector and we are done.

If $\mathbf{p}^*$ is not on any budget constraint hyperplane, then in a small enough neighborhood of $\mathbf{p}^*$, every agent's choice set is unchanging in price. Hence, every agent's demand is unchanging in price near $\mathbf{p}^*$, and $f(\cdot)$ is continuous at $\mathbf{p}^*$. From the construction of $F(\cdot)$ in (A4), this means that $F(\mathbf{p}^*) = f(\mathbf{p}^*)$.

If $\mathbf{p}^* = \tilde{\mathbf{p}}^*$, that is, if the fixed point lies within the legal price space $\mathcal{P}$ and so the truncation is meaningless, then we have

$$\mathbf{p}^* = \tilde{\mathbf{p}}^* \in F(\tilde{\mathbf{p}}^*) = F(\mathbf{p}^*) = f(\mathbf{p}^*),$$

and so step 2 implies that $\mathbf{p}^*$ is an exact competitive equilibrium price vector. For $\mathbf{p}^* \neq \tilde{\mathbf{p}}^*$, that is, for cases in which the fixed point lies in $\tilde{\mathcal{P}}\backslash\mathcal{P}$, we need the following simple lemma.

LEMMA 2. For any $\tilde{\mathbf{p}} \in \tilde{\mathcal{P}}\backslash\mathcal{P}$, (i) $f(\tilde{\mathbf{p}}) = f(t(\tilde{\mathbf{p}}))$ and (ii) $F(\tilde{\mathbf{p}}) \subseteq F(t(\tilde{\mathbf{p}}))$.

Proof. (i) Follows immediately from (A3). (ii) Consider a $\mathbf{y}$ for which there exists a sequence $\tilde{\mathbf{p}}^w \to \tilde{\mathbf{p}}$, $\tilde{\mathbf{p}}^w \neq \tilde{\mathbf{p}}$ such that $f(\tilde{\mathbf{p}}^w) \to \mathbf{y}$. Now consider the sequence $t(\tilde{\mathbf{p}}^w)$. By the continuity of $t(\cdot)$, this sequence converges to $t(\tilde{\mathbf{p}})$, and from part i of the lemma, $f(t(\tilde{\mathbf{p}}^w))$ converges to $\mathbf{y}$. So $\mathbf{y} \in F(\tilde{\mathbf{p}}) \Rightarrow \mathbf{y} \in F(t(\tilde{\mathbf{p}}))$, and the desired result follows. QED

Combining lemma 2 with $F(\mathbf{p}^*) = f(\mathbf{p}^*)$ from above and $\tilde{\mathbf{p}}^* \in F(\tilde{\mathbf{p}}^*)$ from step 3 yields

$$\tilde{\mathbf{p}}^* \in F(\tilde{\mathbf{p}}^*) \subseteq F(\mathbf{p}^*) = f(\mathbf{p}^*) = f(\tilde{\mathbf{p}}^*),$$

so $\tilde{\mathbf{p}}^* = f(\tilde{\mathbf{p}}^*)$, and step 2 implies that $\mathbf{p}^*$ is an exact competitive equilibrium price vector.

Step 5: Suppose that $\mathbf{p}^*$ is on $L \geq 1$ budget constraint hyperplanes. From step 1 we know that $L \leq M$. Let $\Phi = \{0, 1\}^L$. Define a set of 2^L price vectors $\{\mathbf{p}^\phi\}_{\phi\in\Phi}$ satisfying the following conditions:

i. Each $\mathbf{p}^\phi$ is close enough to $\mathbf{p}^*$ that there is a path from $\mathbf{p}^\phi$ to $\mathbf{p}^*$ that does not cross any budget constraint hyperplane (until the moment it reaches $\mathbf{p}^*$).
ii. Each $\mathbf{p}^\phi$ is on the "affordable" side of the lth hyperplane if $\phi_l = 0$ and is on the "unaffordable" side if $\phi_l = 1$.

That is, each $\phi \in \Phi$ "labels" a region of price space close to $\mathbf{p}^*$.

Each of the L intersecting budget constraint hyperplanes defines two half spaces. Let $H_l^0 = \{\mathbf{p} \in \tilde{\mathcal{P}}: \mathbf{p} \cdot x_l \leq b_{i_l} + \tau_{i_l x_l}\}$ denote the closed half space in which the agent named on the lth hyperplane, i_l, can weakly afford the bundle named on the lth hyperplane, x_l. Let $H_l^1 = \{\mathbf{p} \in \tilde{\mathcal{P}}: \mathbf{p} \cdot x_l > b_{i_l} + \tau_{i_l x_l}\}$ denote the open half space in which agent i_l cannot afford bundle x_l.

We label combinations of half spaces as follows. Let $\Phi = \{0, 1\}^L$, with each label $\phi = (\phi_1, \ldots, \phi_L) \in \Phi$ an L-dimensional vector of zeros and ones. The convex polytope $\pi(\phi) := \cap_{l=1}^L H_l^{\phi_l}$ denotes the set of points in $\tilde{\mathcal{P}}$ that belong to the intersection of half spaces indexed by ϕ.

Let $\mathcal{H}$ denote the finite set of all hyperplanes formed by any i, x: $\mathcal{H} = \{H(i, x)_{i\in\mathcal{S}, x_i\in 2^{\mathcal{C}}}\}$. Let

$$\delta < \inf_{\mathbf{p}''\in\tilde{\mathcal{P}}, H\in\mathcal{H}} \{\|(\mathbf{p}^* - \mathbf{p}'')\|_2 : \mathbf{p}'' \in H, \mathbf{p}^* \notin H\}.$$

That is, any hyperplane to which $\mathbf{p}^*$ does not belong is strictly further than δ away from $\mathbf{p}^*$ in Euclidean distance. Let $B_\delta(\mathbf{p}^*)$ denote a δ-ball of $\mathbf{p}^*$.

We can now define a set $\{\mathbf{p}^\phi\}_{\phi\in\Phi}$ satisfying the requirements above: each $\mathbf{p}^\phi$ is an arbitrary element of $\pi(\phi) \cap B_\delta(\mathbf{p}^*)$.

Step 6: Consider the set of excess demands $\{\mathbf{z}(\mathbf{p}^\phi)\}_{\phi\in\Phi}$ corresponding to the prices $\{\mathbf{p}^\phi\}_{\phi\in\Phi}$ defined in step 5. A perfect market-clearing excess demand vector, ζ, lies in the convex hull of $\{\mathbf{z}(\mathbf{p}^\phi)\}_{\phi\in\Phi}$.

We need the following lemma.

LEMMA 3. For any $\mathbf{y} \in F(\mathbf{p}^*)$ there exist nonnegative weights $\{\lambda^\phi\}_{\phi\in\Phi}$ with $\sum_{\phi\in\Phi}\lambda^\phi = 1$ such that $\mathbf{p}^* + \sum_{\phi\in\Phi}\lambda^\phi\gamma z(\mathbf{p}^\phi) = \mathbf{y}$.

Proof. Consider an arbitrary ϕ and consider any two prices $\mathbf{p}', \mathbf{p}'' \in \pi(\phi) \cap B_\delta(\mathbf{p}^*)$. Since both prices are in $\pi(\phi)$, they are on the same side of each of the L hyperplanes that intersect at $\mathbf{p}^*$. Since both prices are in $B_\delta(\mathbf{p}^*)$, by the way we chose δ in step 5, for any other hyperplane in $\mathcal{H}$, $\mathbf{p}'$ and $\mathbf{p}''$ are on the same side. Together, this means that every agent has the same choice set at $\mathbf{p}'$ as at $\mathbf{p}''$. Since we chose $\mathbf{p}', \mathbf{p}''$ arbitrarily, demand at any price vector in $\pi(\phi) \cap B_\delta(\mathbf{p}^*)$ is equal to demand at $\mathbf{p}^\phi$.

Consider any sequence of prices $\mathbf{p}^{w,\phi} \to \mathbf{p}^*$, with each $\mathbf{p}^{w,\phi} \in \pi(\phi) \cap B_\delta(\mathbf{p}^*)$. The preceding argument implies

$$f(\mathbf{p}^{w,\phi}) \to \mathbf{p}^* + \gamma\mathbf{z}(\mathbf{p}^\phi). \tag{A5}$$

Note too that any sequence $\mathbf{p}^w \to \mathbf{p}^*$ for which $f(\mathbf{p}^w)$ converges must converge to $\mathbf{p}^* + \mathbf{z}(\mathbf{p}^{\phi'})$ for some $\phi' \in \Phi$ because $\cup_{\phi\in\Phi}\pi(\phi) \cap B_\delta(\mathbf{p}^*) = \tilde{\mathcal{P}} \cap B_\delta(\mathbf{p}^*)$. This observation, (A5), and the way that we constructed $F(\cdot)$ (i.e., [A4]) together imply that if $\mathbf{y} \in F(\mathbf{p}^*)$, there exist nonnegative weights $\{\lambda^\phi\}_{\phi\in\Phi}$, with $\sum_{\phi\in\Phi}\lambda^\phi = 1$, such that $\mathbf{p}^* + \sum_{\phi\in\Phi}\lambda^\phi\gamma\mathbf{z}(\mathbf{p}^\phi) = \mathbf{y}$. QED

From step 3 we have $\tilde{\mathbf{p}}^* \in F(\tilde{\mathbf{p}}^*)$ and from lemma 2 in step 4 we have $F(\tilde{\mathbf{p}}^*) \subseteq F(\mathbf{p}^*)$, so $\tilde{\mathbf{p}}^* \in F(\mathbf{p}^*)$. Thus we can apply lemma 3 to $\tilde{\mathbf{p}}^* \in F(\mathbf{p}^*)$ to obtain that there exist nonnegative weights $\{\lambda^\phi\}_{\phi\in\Phi}$, with $\sum_{\phi\in\Phi}\lambda^\phi = 1$ such that

$$\mathbf{p}^* + \sum_{\phi\in\Phi}\lambda^\phi\gamma\mathbf{z}(\mathbf{p}^\phi) = \tilde{\mathbf{p}}^*.$$

This in turn implies (using the same λ's)

$$\sum_{\phi\in\Phi}\lambda^\phi\mathbf{z}(\mathbf{p}^\phi) = \frac{\tilde{\mathbf{p}}^* - \mathbf{p}^*}{\gamma}.$$

By the same argument as in step 2, demand for any object j must be zero at price $\bar{b}$ or higher, so $\tilde{p}_j^* \in [-1, \bar{b})$ for all $j \in \mathcal{C}$. So, for all j, either $\tilde{p}_j^* = p_j^*$ or $\tilde{p}_j^* < 0 = p_j^*$. So we have that $\sum_{\phi\in\Phi}\lambda^\phi\mathbf{z}(\mathbf{p}^\phi) \le 0$ with $\sum_{\phi\in\Phi}\lambda^\phi z_j(\mathbf{p}^\phi) < 0 \Rightarrow p_j^* = 0$.

Define $\boldsymbol{\zeta} \equiv \sum_{\phi\in\Phi}\lambda^\phi\mathbf{z}(\mathbf{p}^\phi)$. The vector $\boldsymbol{\zeta}$ is in the convex hull of $\{\mathbf{z}(\mathbf{p}^\phi)\}_{\phi\in\Phi}$ by construction, and it is a perfect market-clearing ideal at prices $\mathbf{p}^*$ since $\boldsymbol{\zeta} \le 0$ with $\zeta_j < 0 \Rightarrow p_j^* = 0$.

Step 7: The set of excess demands $\{\mathbf{z}(\mathbf{p}^\phi)\}_{\phi\in\Phi}$ has a special geometric structure.

The L budget constraint hyperplanes that intersect at $\mathbf{p}^*$ name $L' \le L$ distinct agents; renumber the agents in $\mathcal{S}$ so that agents $\{1, \ldots, L'\}$ are the ones so named. Denote by w_i the number of intersecting hyperplanes that name agent $i \in \{1, \ldots, L'\}$, and let $x_i^1, \ldots, x_i^{w_i}$ denote the bundles pertaining to i's hyperplanes, numbered so that $x_i^1 \succ_i \cdots \succ_i x_i^{w_i}$. Note that $\sum_{i=1}^{L'} w_i = L$.

The following argument illustrates that agent $i \in \{1, \ldots, L'\}$ purchases at most $w_i + 1$ distinct bundles at prices near to $\mathbf{p}^*$. In the half space $H^0(i, x_i^1)$ he can afford x_i^1, his favorite bundle whose affordability is in question near to $\mathbf{p}^*$, and so it does not matter which side of $H(i, x_i^2), \ldots, H(i, x_i^{w_i})$ price is on. Let d_i^0 denote his demand at prices in $H^0(i, x_i^1) \cap B_\delta(\mathbf{p}^*) \cap \tilde{\mathcal{P}}$. If price is in $H^1(i, x_i^1) \cap H^0(i, x_i^2)$, then i cannot afford x_i^1 but can afford x_i^2, his second-favorite bundle whose affordability is in question. So it does not matter which side of $H(i, x_i^3), \ldots, H(i, x_i^{w_i})$ price is on. Let d_i^1 denote his demand at prices in $H^1(i,$

$x_i^1) \cap H^0(i, x_i^2) \cap B_\delta(\mathbf{p}^*) \cap \tilde{\mathcal{P}}$. Continuing in this manner, define $d_i^2, \ldots, d_i^{w_i}$. The process ends when we have crossed to the unaffordable side of all w_i of i's budget constraint hyperplanes, so i cannot afford any of $x_i^1, \ldots, x_i^{w_i}$.

The demand of any agents other than the L' just discussed is unchanging near $\mathbf{p}^*$. Call the total demand of such agents

$$d_{\mathcal{S}\setminus\{1,\ldots,L'\}}(\mathbf{p}^*) = \sum_{i=L'+1}^{N} d_i(\mathbf{p}^*;\, b_i',\, \tau_i'),$$

and let

$$z_{\mathcal{S}\setminus\{1,\ldots,L'\}}(\mathbf{p}^*) = d_{\mathcal{S}\setminus\{1,\ldots,L'\}}(\mathbf{p}^*) - \mathbf{q}.$$

We can now characterize the set $\{\mathbf{z}(\mathbf{p}^\phi)\}_{\phi\in\Phi}$ in terms of the demands of the L' individual agents near $\mathbf{p}^*$:

$$\{\mathbf{z}(\mathbf{p}^\phi)\}_{\phi\in\Phi} = \left\{z_{\mathcal{S}\setminus\{1,\ldots,L'\}}(\mathbf{p}^*) + \sum_{i=1}^{L'}\sum_{f=0}^{w_i} a_i^f d_i^f\right\} \tag{A6}$$

subject to

$$a_i^f \in \{0,\, 1\} \quad \text{for all } i = 1, \ldots, L',\, f = 0, \ldots, w_i,$$

$$\sum_{f=0}^{w_i} a_i^f = 1 \quad \text{for all } i = 1, \ldots, L'.$$

At any price vector near to $\mathbf{p}^*$, each agent $i = 1, \ldots, L'$ demands exactly one of his $w_i + 1$ demand bundles. Over the set $\Phi = \{0, 1\}^L$, every combination of the L' agents' demands is possible. The set (A6) has a particularly intuitive structure in case $L' = L$ (and so $w_i = 1$ for $i = 1, \ldots, L'$); see (3).

Step 6 tells us that there exists a market-clearing excess demand vector in the convex hull of (A6). This convex hull can be written as

$$\{\mathbf{z}(\mathbf{p}^\phi)\}_{\phi\in\Phi} = \left\{z_{\mathcal{S}\setminus\{1,\ldots,L'\}}(\mathbf{p}^*) + \sum_{i=1}^{L'}\sum_{f=0}^{w_i} a_i^f d_i^f\right\} \tag{A7}$$

subject to

$$a_i^f \in [0,\, 1] \quad \text{for all } i = 1, \ldots, L',\, f = 0, \ldots, w_i,$$

$$\sum_{f=0}^{w_i} a_i^f = 1 \quad \text{for all } i = 1, \ldots, L'.$$

Step 8: There exists a vertex of the geometric structure from step 7, (A6), that is within $\sqrt{\sigma M}/2$ distance of the perfect market-clearing excess demand vector, ζ, found in step 6. That is, for some $\mathbf{z}(\mathbf{p}^{\phi'}) \in \{\mathbf{z}(\mathbf{p}^\phi)\}_{\phi\in\Phi}$, $\|\mathbf{z}(\mathbf{p}^{\phi'}) - \mathbf{z}\|_2 \le \sqrt{\sigma M}/2$.

We are interested in bounding the distance between an element of (A6) and an element of its convex hull (A7), which we know contains ζ.[33]

Fix an arbitrary point of (A7); that is, fix a set of $a_i^f \in [0, 1]$ that satisfy $\sum_{f=0}^{w_i} a_i^f = 1$ for all $i = 1, \ldots, L'$. For each $i = 1, \ldots, L'$ define a random vector

[33] The proof technique for this step closely follows that of theorem 2.4.2 in Alon and Spencer (2000). I am grateful to Michel Goemans for the pointer. Another choice would be to use the Shapley Folkman theorem (Starr 1969).

$\Theta_i = (\Theta_i^0, \ldots, \Theta_i^{w_i})$, where the support of each Θ_i^f is $\{0, 1\}$, $E(\Theta_i^f) = a_i^f$ for all $f = 1, \ldots, w_i$, and in any realization θ_i, $\sum_{f=0}^{w_i} \theta_i^f = 1$. Suppose that the Θ_i's are independent. Let

$$\rho^2 = \mathbb{E}_{\Theta_1,\ldots,\Theta_{L'}}\left[\left\| \sum_{i=1}^{L'} \sum_{f=0}^{w_i} [(a_i^f - \theta_i^f) d_i^f] \right\|_2\right]^2.$$

Linearity of expectations yields

$$\rho^2 = \sum_{i=1}^{L'} \mathbb{E}_{\Theta_i}\left[\left\| \sum_{f=0}^{w_i} [(a_i^f - \theta_i^f)] d_i^f \right\|_2\right]^2 + \sum_{j\neq i} \sum_{f=0}^{w_i} \sum_{g=0}^{wj} \mathbb{E}_{\Theta_i^f \Theta_j^g}[(a_i^f - \theta_i^f)(a_j^g - \theta_j^g)](d_i^f \cdot d_j^g). \tag{A8}$$

Independence yields

$$\mathbb{E}_{\Theta_i^f \Theta_j^g}[(a_i^f - \theta_i^f)(a_j^g - \theta_j^g)] = \mathbb{E}_{\Theta_i^f}[a_i^f - \theta_i^f]\mathbb{E}_{\Theta_j^g}[a_j^g - \theta_j^g] = 0 \tag{A9}$$

since the random vectors are independent across agents and $\mathbb{E}_{\Theta_i^f}\theta_i^f = a_i^f$ for all i, f.

Lemma 4. For each $i = 1, \ldots, L'$,

$$\mathbb{E}_{\Theta_i}\left[\left\| \sum_{f=0}^{w_i} [(a_i^f - \theta_i^f)] d_i^f \right\|_2\right]^2 \leq \frac{\sigma w_i}{4}.$$

Proof. Fix i. For any d_i^f, $d_i^{f'}$, the vector $d_i^f - d_i^{f'} \in \{-1, 0, 1\}^M$ has at most $\sigma = \min(2k, M)$ nonzero elements, where k is the maximum number of objects in a permissible bundle and M is the number of object types. Thus $\|d_i^f - d_i^{f'}\|_2 \leq \sqrt{\sigma}$. Let $\bar{d}_i = \sum_{f=0}^{w_i} a_i^f d_i^f$. Now rewrite

$$\mathbb{E}_{\Theta_i}\left[\left\| \sum_{f=0}^{w_i} [(a_i^f - \theta_i^f)] d_i^f \right\|_2\right]^2 = \sum_{f=0}^{w_i} a_i^f (\|\bar{d}_i - d_i^f\|_2)^2. \tag{A10}$$

If $w_i = 1$, then (A10) is largest when $\|d_i^1 - d_i^0\|_2 = \sqrt{\sigma}$ and $a_i^0 = a_i^1 = \frac{1}{2}$; this maximum value is $\sigma/4$, which is equal to the bound. If $w_i = 2$, then (A10) is largest when $\{d_i^0, d_i^1, d_i^2\}$ forms an equilateral triangle of side length $\sqrt{\sigma}$ and $a_i^0 = a_i^1 = a_i^2 = \frac{1}{3}$; this maximum value is $\sigma/3$, which is strictly lower than the bound of $\sigma/2$. If $w_i = 3$, then (A10) is largest when $\{d_i^0, d_i^1, d_i^2, d_i^3\}$ forms a triangular pyramid of side length $\sqrt{\sigma}$ and $a_i^0 = a_i^1 = a_i^2 = a_i^3 = \frac{1}{4}$; this maximum value is $3\sigma/8$, which is strictly lower than the bound of $3\sigma/4$. For $w_i \geq 4$, the bound can be obtained by observing that there exists some sphere of diameter $\sqrt{\sigma}$ that contains the convex hull of $\{d_i^f\}_{f=0}^{w_i}$, so the expected squared distance is less than or equal to σ whereas the right-hand side of the bound $\sigma w_i/4 \geq \sigma$. QED

Combining lemma 4, (A8), and (A9) yields

$$\rho^2 = \sum_{i=1}^{L'} \mathbb{E}_{\Theta_i}\left[\left\| \sum_{f=0}^{w_i} [(a_i^f - \theta_i^f) d_i^f] \right\|_2\right]^2 \leq \sum_{i=1}^{L'} \frac{\sigma w_i}{4} = \frac{\sigma L}{4} \leq \frac{\sigma M}{4}.$$

This means that there must exist at least one realization of ϕ such that

$$\left\| \sum_{i=1}^{L'} \sum_{f=0}^{w_i} [(a_i^f - \theta_i^f)] d_i^f \right\|_2 \leq \frac{\sqrt{\sigma M}}{2}.$$

Since we chose the a_i^f's arbitrarily, there exists such a realization for any interior point of (A7), in particular for weights $\tilde{a}_i^f$ such that

$$z_{\mathcal{S}\setminus\{1,\ldots,L'\}}(\mathbf{p}^*) + \sum_{i=1}^{L'}\sum_{f=0}^{w_i} \tilde{a}_i^f d_i^f = \zeta.$$

Call this realization $\tilde{\theta}$. This realization points us to an element of $\{\mathbf{z}(\mathbf{p}^{\phi})\}_{\phi\in\Phi}$, namely

$$z_{\mathcal{S}\setminus\{1,\ldots,L'\}}(\mathbf{p}^*) + \sum_{i=1}^{L'}\sum_{f=0}^{w_i} \tilde{\theta}_i^f d_i^f,$$

that is within $\sqrt{\sigma M}/2$ Euclidean distance of the perfect market-clearing ideal point, ζ. Call this element $\mathbf{z}(\mathbf{p}^{\phi'})$.

Step 9: Use the vertex found in step 8, $\mathbf{z}(\mathbf{p}^{\phi'})$, to produce prices, budgets, and an allocation that satisfy the statement of theorem 1.

There is no guarantee that $\mathbf{p}^{\phi'} \in \mathcal{P}$; in particular, if $p_j^* = 0$, it is possible that $p_j^{\phi'}$ is strictly negative, so $\mathbf{p}^{\phi'} \in \tilde{\mathcal{P}}\setminus\mathcal{P}$. So we will use the prices $\mathbf{p}^*$, which are guaranteed to be in $\mathcal{P}$, and perturb budgets in a way that generates excess demand at $\mathbf{p}^*$ equal to $\mathbf{z}(\mathbf{p}^{\phi'})$ from step 8.

If agent $i \in \mathcal{S}$ is not named on any of the L budget constraint hyperplanes of step 5, then his consumption is $x_i^* = d_i(\mathbf{p}^*, b_i', \tau_i')$ and we set $b_i^* = b_i' + \tau'_{ix_i^*}$. Requirement i of step 1 implies that any bundle he prefers to x_i^* costs strictly more than $b_i' + \tau'_{ix_i^*}$; else he would demand it at prices $\mathbf{p}^*$, budget b_i', and taxes of τ_i'.

If agent $i \in \mathcal{S}$ is named on some of the budget constraint hyperplanes, then we will use the information in ϕ' to perturb his taxes and ultimately his budget. For $f = 1, \ldots, w_i$, if $\mathbf{p}^{\phi'} \in H^1(i, x^f)$, that is, x^f is unaffordable for i at $\mathbf{p}^{\phi'}$, then set $\tau''_{ix^f} = \tau'_{ix^f} - \delta_2$ for $\delta_2 > 0$ but small enough to preserve conditions i–iii of step 1. For all other bundles, including bundles not named on any hyperplane, set $\tau''_{ix} = \tau'_{ix}$. Consider $d_i(\mathbf{p}^*, b_i', \tau_i'')$: this is simply i's demand at the original budget and taxes but at prices $\mathbf{p}^{\phi'}$, that is, $d_i(\mathbf{p}^*, b_i', \tau_i'') = d_i(\mathbf{p}^{\phi'}, b_i', \tau_i')$.

Set $x_i^* = d_i(\mathbf{p}^*, b_i', \tau_i'')$ and set $b_i^* = b_i' + \tau''_{ix_i^*}$ for all $i \in \mathcal{S}$. Now set all taxes equal to zero. Since we set δ_2 small enough to ensure that requirement ii of step 1 still obtains, x_i^* remains optimal for i at prices $\mathbf{p}^*$ and a budget of b_i^*. Similarly, we have preserved the original level of budget inequality and the ε bounds, by requirements iii and i, respectively, of step 1. Approximate market clearing is ensured by step 8. So budgets of $\mathbf{b}^*$, prices of $\mathbf{p}^*$, and the allocation $\mathbf{x}^*$ satisfy all of the requirements of theorem 1. QED

Appendix B

Other Omitted Proofs

Proof of Proposition 1

For the case $M = 4$, consider the following example.

EXAMPLE 1. There are four objects, $\mathcal{C} = \{a, b, c, d\}$, each with capacity 2. There are four agents, $\mathcal{S} = \{i_1, i_2, i_3, i_4\}$, whose preferences are $\succ_{i_1} : \{a, b, c\}, \{d\}, \ldots, \succ_{i_2} : \{a, b, d\}, \{c\}, \ldots, \succ_{i_3} : \{a, c, d\}, \{b\}, \ldots$, and $\succ_{i_4} : \{b, c, d\}, \{a\}, \ldots$.

Fix $0 < \beta < \frac{1}{2}$, and consider an arbitrary budget vector $\mathbf{b}^* = (1 + \beta_1, 1 + \beta_2,$

$1+\beta_3$, $1+\beta_4$) for $\beta_1, \beta_2, \beta_3, \beta_4 < \beta$. The unique fixed point of correspondence (A4) is $\mathbf{p}^*$ as given by

$$p_a^* = \frac{1+\beta_1+\beta_2+\beta_3-2\beta_4}{3},$$

$$p_b^* = \frac{1+\beta_1+\beta_2-2\beta_3+\beta_4}{3},$$

$$p_c^* = \frac{1+\beta_1-2\beta_2+\beta_3+\beta_4}{3},$$

$$p_d^* = \frac{1-2\beta_1+\beta_2+\beta_3+\beta_4}{3}.$$

At $\mathbf{p}^*$, each agent can exactly afford her most preferred bundle and can strictly afford her second most preferred bundle, so in arbitrary sequences $\mathbf{p}^w \rightarrow \mathbf{p}^*$, each agent's demand converges to either of her two most preferred bundles. The convex combination in which each agent receives each bundle with probability one-half exactly clears the market (and is unique in this respect).

Every feasible demand in a neighborhood of $\mathbf{p}^*$ is Euclidean distance $\sqrt{\sigma M}/2 = 2$ from the perfect market-clearing demand of $\mathbf{q} = (2, 2, 2, 2)$.

To see why example 1 obtains the theorem 1 bound, consider the matrix that is formed by stacking the four agents' change-in-demand vectors at $\mathbf{p}^*$ (see Sec. III.D.4):

$$\begin{pmatrix} -1 & -1 & -1 & +1 \\ -1 & -1 & +1 & -1 \\ -1 & +1 & -1 & -1 \\ +1 & -1 & -1 & -1 \end{pmatrix}. \qquad \text{(B1)}$$

This is a *Hadamard matrix*: all of its entries are ± 1, and its rows are mutually orthogonal (Wallis, Street, and Wallis 1972). Whenever the change-in-demand matrix at $\mathbf{p}^*$ is a Hadamard matrix, aggregate demand in a neighborhood of $\mathbf{p}^*$ forms a hypercube with sides of length $\sqrt{M}$ (here, $\sigma = M$).

The Hadamard matrix (B1) has an additional feature, *regularity*, which requires that each row has the same number of +1's. Neil Sloane has shown that regular Hadamard matrices exist for all powers of 4.[34] We can use these regular Hadamard matrices to construct examples that are exactly analogous to example 1 for $M = 16, 64, 256, \ldots$. QED

[34] Here is Sloane's proof. Let A be the matrix defined in (B1). The tensor product of two Hadamard matrices is itself a Hadamard matrix, and the tensor product preserves the "same number of +1's per row" property. So $A \otimes A$ is a 16-dimensional Hadamard matrix with the same number of +1's per row, $A \otimes (A \otimes A)$ is a 64-dimensional example, and so forth. QED. It has been conjectured that there exist regular Hadamard matrices of order $(2n)^2$ for any integer n. Useful references are http://www2.research.att.com/~njas/hadamard/index.html and http://oeis.org/A016742.

Proof of Proposition 2

Suppose that $\mathbf{x}'$ Pareto improves on $\mathbf{x}^*$ in economy $(\mathcal{S}, \mathcal{C}, (q_j^*)_{j=1}^M, (\Psi_i)_{i=1}^N, (\succsim_i)_{i=1}^N)$. By condition i of definition 1 and strict preferences, if $x_i' \neq x_i^*$, then $\mathbf{p}^* \cdot x_i' > \mathbf{p}^* \cdot x_i^*$. This implies that $\sum_i \mathbf{p}^* \cdot x_i' > \sum_i \mathbf{p}^* \cdot x_i^*$, a contradiction since prices are nonnegative and $\mathbf{x}^*$ allocates all units of positive-priced goods. QED

Proof of Lemma 1

Step 6 of the proof of theorem 1 shows that $\sum_{\phi \in \Phi} \lambda^\phi \mathbf{p}^* \cdot \mathbf{z}(\mathbf{p}^\phi) = 0$ for a set of nonnegative weights $(\lambda^\phi)_{\phi \in \Phi}$ summing to one and a set of prices $(\mathbf{p}^\phi)_{\phi \in \Phi}$ arbitrarily close to $\mathbf{p}^*$. Choose $\varepsilon_2 > 0$ as an upper bound on the distance in the L_1 norm between $\mathbf{p}^*$ and each $\mathbf{p}^\phi$; for any bundle x and any ϕ, this means $\mathbf{p}^\phi \cdot x \geq \mathbf{p}^* \cdot x - \varepsilon_2$. Summing over all $\mathbf{p}^\phi$, we have

$$\begin{aligned} \sum_{\phi \in \Phi} \lambda^\phi \mathbf{p}^\phi \cdot \left[\sum_{i=1}^N d_i(\mathbf{p}^\phi, b_i', \tau_i')\right] &\geq \sum_{\phi \in \Phi} \lambda^\phi \mathbf{p}^* \cdot \left[\sum_{i=1}^N d_i(\mathbf{p}^\phi, b_i', \tau_i')\right] - N\varepsilon_2 \\ &= \mathbf{p}^* \cdot \mathbf{q} - N\varepsilon_2, \end{aligned} \tag{B2}$$

where the equality follows from noting that

$$\sum_{i=1}^N d_i(\mathbf{p}^\phi, b_i', \tau_i') = \mathbf{z}(\mathbf{p}^\phi) + \mathbf{q}$$

by definition (A2). Recall that $\varepsilon > 0$ is an upper bound for both τ_{ix}' and the discrepancy between each b_i' and b_i^*. Thus at each price $\mathbf{p}^\phi$, each agent i's expenditure is weakly less than $b_i' + \varepsilon$, which itself is weakly less than $b_i^* + 2\varepsilon$, so $\mathbf{p}^\phi \cdot d_i(\mathbf{p}^\phi, b_i', \tau_i') \leq b_i^* + 2\varepsilon$. Summing over all i and ϕ gives

$$\sum_{\phi \in \Phi} \lambda^\phi \mathbf{p}^\phi \cdot \left[\sum_{i=1}^N d_i(\mathbf{p}^\phi, b_i', \tau_i')\right] \leq \sum_{i=1}^N b_i^* + 2N\varepsilon.$$

Together with (B2), this implies

$$\mathbf{p}^* \cdot \mathbf{q} \leq \sum_{i=1}^N b_i^* + 2N\varepsilon + N\varepsilon_2.$$

Since ε and ε_2 can be set arbitrarily small, the desired result follows. QED

Proof of Theorem 2

Since $\mathbf{b}^*$ and $\mathbf{p}^*$ are part of an (α, β)–CEEI with $\mathbf{p}^* \in \mathcal{P}(\delta, \mathbf{b}^*)$,

$$N(1+\beta)(1+\delta) \geq \sum_{i=1}^N b_i^*(1+\delta) \geq \mathbf{p}^* \cdot \mathbf{q}.$$

Let $\mathbf{x}^{MS^i}$ denote an $(N+1)$–maximin split for agent i. Suppose that i cannot afford any bundle in $\mathbf{x}^{MS^i}$ at $\mathbf{p}^*$. Then $\mathbf{p}^* \cdot x_l^{MS^i} > b_i^* \geq 1$ for all $l = 1, \ldots, N, N+1$. By the definition of the $(N+1)$–maximin split, we have $\sum_l \mathbf{p}^* \cdot x_l^{MS^i} \leq \mathbf{p}^* \cdot \mathbf{q}$. Putting this all together gives

$$N(1+\beta)(1+\delta) \geq \mathbf{p}^* \cdot \mathbf{q} \geq \sum_{l} \mathbf{p}^* \cdot x_l^{MS'} > N+1,$$

which contradicts $\beta < (1 - N\delta)/N(1+\delta)$. QED

References

Abdulkadiroğlu, Atila, Parag Pathak, and Alvin E. Roth. 2005. "The Boston Public School Match." *A.E.R. Papers and Proc.* 95 (2): 368–71.

Abdulkadiroğlu, Atila, Parag Pathak, Alvin E. Roth, and Tayfun Sönmez. 2005. "The New York City High School Match." *A.E.R. Papers and Proc.* 95 (2): 364–67.

Abdulkadiroğlu, Atila, and Tayfun Sönmez. 1998. "Random Serial Dictatorship and the Core from Random Endowments in House Allocation Problems." *Econometrica* 66 (3): 689–702.

———. 1999. "House Allocation with Existing Tenants." *J. Econ. Theory* 88 (2): 233–60.

———. 2003. "School Choice: A Mechanism Design Approach." *A.E.R.* 93 (June): 729–47.

Adler, Barry, Thomas Delaney, Michelle Kirkland, and Liam Murphy. 2008. "Memorandum re: Changes to the Registration System." New York Univ., Law School, April 25, http://bit.ly/n6FmQp.

Albergotti, Reed. 2010. "Why the NFL Draft Drives Economists Crazy." *Wall Street J.*, April 22.

Alkan, Ahmet, Gabrielle Demange, and David Gale. 1991. "Fair Allocation of Indivisible Goods and Criteria of Justice." *Econometrica* 59 (4): 1023–39.

Alon, Noga, and Joel Spencer. 2000. *The Probabilistic Method.* 2nd ed. New York: Wiley.

Arnsperger, Christian. 1994. "Envy-Freeness and Distributive Justice." *J. Econ. Surveys* 8 (2): 155–86.

Arrow, Kenneth, and F. H. Hahn. 1971. *General Competitive Analysis.* San Francisco: Holden-Day.

Azevedo, Eduardo, and Eric Budish. 2012. "Strategyproofness in the Large as a Desideratum for Market Design." Working paper, Univ. Chicago, http://faculty.chicagobooth.edu/eric.budish/research/Azevedo-Budish-SPL-Jan2012.pdf.

Bartlett, Thomas. 2008. "Class Warfare: When Getting In Is the Hardest Part." *Chronicle Higher Educ.* 54 (February 15): A1.

Bogomolnaia, Anna, and Hervé Moulin. 2001. "A New Solution to the Random Assignment Problem." *J. Econ. Theory* 100:295–328.

Brainard, William C., and Herbert E. Scarf. 2005. "How to Compute Equilibrium Prices in 1891." *American J. Econ. and Sociology* 64 (1): 57–83.

Brams, Steven J., and Alan D. Taylor. 1996. *Fair Division: From Cake Cutting to Dispute Resolution.* Cambridge: Cambridge Univ. Press

———. 1999. *The Win-Win Solution: Guaranteeing Fair Shares to Everybody.* New York: Norton.

Budish, Eric. 2010. "The Combinatorial Assignment Problem: Approximate Competitive Equilibrium from Equal Incomes." Applied Theory Initiative Working Paper no. ATI-2010-15, Univ. Chicago, Booth School.

Budish, Eric, and Estelle Cantillon. Forthcoming. "The Multi-unit Assignment Problem: Theory and Evidence from Course Allocation at Harvard." *A.E.R.*

Chen, Yan, and Tayfun Sönmez. 2002. "Improving Efficiency of On-Campus Housing: An Experimental Study." *A.E.R.* 92 (5): 1669–86.

Cramton, Peter, Yoav Shoham, and Richard Steinberg, eds. 2006. *Combinatorial Auctions.* Cambridge, MA: MIT Press.

Crawford, Vincent P. 1977. "A Game of Fair Division." *Rev. Econ. Studies* 44 (2): 235–47.

Cromme, Ludwig J., and Immo Diener. 1991. "Fixed Point Theorems for Discontinuous Mapping." *Math. Programming* 51:257–67.

Dierker, Egbert. 1971. "Equilibrium Analysis of Exchange Economies with Indivisible Commodities." *Econometrica* 39 (6): 997–1008.

Dworkin, Ronald. 1981. "What Is Equality? Part 2: Equality of Resources." *Philosophy and Public Affairs* 10 (4): 283–345.

———. 2000. *Sovereign Virtue: The Theory and Practice of Equality.* Cambridge, MA: Harvard Univ. Press.

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

Foley, Duncan. 1967. "Resource Allocation and the Public Sector." *Yale Econ. Essays* 7:45–98.

Friedman, Milton. 1991. "How to Sell Government Securities." *Wall Street J.*, August 28.

Gale, David, and Lloyd Shapley. 1962. "College Admissions and the Stability of Marriage." *American Math. Monthly* 69 (January): 9–15.

Graves, Robert L., Linus Schrage, and Jayaram Sankaran. 1993. "An Auction Method for Course Registration." *Interfaces* 23 (5): 81–92.

Guernsey, Lisa. 1999. "Business School Puts Courses in Hands of an On-Line Market." *New York Times*, September 9.

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

Herreiner, Dorothea, and Clemens Puppe. 2002. "A Simple Procedure for Finding Equitable Allocations of Indivisible Goods." *Soc. Choice and Welfare* 19:415–30.

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

Kaplow, Louis, and Steven Shavell. 2001. "Any Non-welfarist Method of Policy Assessment Violates the Pareto Principle." *J.P.E.* 109 (2): 281–86.

———. 2007. "Moral Rules, the Moral Sentiments, and Behavior: Toward a Theory of an Optimal Moral System." *J.P.E.* 115 (3): 494–514.

Klaus, Bettina, and Eiichi Miyagawa. 2001. "Strategy-Proofness, Solidarity, and Consistency for Multiple Assignment Problems." *Internat. J. Game Theory* 30: 421–35.

Kojima, Fuhito. 2009. "Random Assignment of Multiple Indivisible Objects." *Math. Soc. Sci.* 57:134–42.

Kojima, Fuhito, and Parag Pathak. 2009. "Incentives and Stability in Large Two-Sided Matching Markets." *A.E.R.* 99:608–27.

Konishi, Hideo, Thomas Quint, and Jun Wako. 2001. "On the Shapley-Scarf Economy: The Case of Multiple Types of Indivisible Goods." *J. Math. Econ.* 35 (1): 1–15.

Krishna, Aradhna, and Utku Ünver. 2008. "Improving the Efficiency of Course Bidding at Business Schools: Field and Laboratory Studies." *Management Sci.* 27 (March/April): 262–82.

Lehrer, Jim. 2008. "College Students Squeezed by Rising Costs, Less Aid: Straining the System." *PBS NewsHour*, December 9, http://to.pbs.org/qfNgXJ.

Levitt, Steven. 2008*a*. “The Coase Theorem Rules at NYU Law.” In *Freakonomics Blog, New York Times*, August 11.

———. 2008*b*. “What Do I Have in Common with Hannah Montana?” In *Freakonomics Blog, New York Times*, January 8.

Lipton, Richard, Evangelos Markakis, Elchanan Mossel, and Amin Saberi. 2004. “On Approximately Fair Allocations of Indivisible Goods.” In *EC ’04: Proceedings of the 5th ACM Conference on Electronic Commerce*. New York: ACM.

Maskin, Eric S. 1987. “On the Fair Allocation of Indivisible Goods.” In *Arrow and the Foundations of the Theory of Economic Policy*, edited by George R. Feiwel. New York: New York Univ. Press.

Milgrom, Paul. 2000. “Putting Auction Theory to Work: The Simultaneous Ascending Auction.” *J.P.E.* 108 (2): 245–72.

———. 2004. *Putting Auction Theory to Work*. Cambridge: Cambridge Univ. Press.

MIT (Massachusetts Institute of Technology). 2008. “Sloan Course Bidding Information.” Massachusetts Inst. Tech., Sloan School, https://sloanbid.mit.edu.

Mongell, Susan J., and Alvin E. Roth. 1986. “A Note on Job Matching with Budget Constraints.” *Econ. Letters* 21:135–38.

Moulin, Hervé. 1991. “Welfare Bounds in the Fair Division Problem.” *J. Econ. Theory* 54:321–37.

———. 1995. *Cooperative Microeconomics*. London: Prentice Hall.

———. 2004. *Fair Division and Collective Welfare*. Cambridge, MA: MIT Press.

Ostrovsky, Michael. 2008. “Stability in Supply Chain Networks.” *A.E.R.* 98 (June): 897–923.

Othman, Abe, Eric Budish, and Tuomas Sandholm. 2010. “Finding Approximate Competitive Equilibria: Efficient and Fair Course Allocation.” In *Proceedings of 9th International Conference on Autonomous Agents and Multiagent Systems.*. New York: ACM.

Pápai, Szilvia. 2001. “Strategyproof and Nonbossy Multiple Assignments.” *J. Public Econ. Theory* 3 (3): 257–71.

Parkes, David. 2006. “Iterative Combinatorial Auctions.” In *Combinatorial Auctions*, edited by Peter Cramton, Yoav Shoman, and Richard Steinberg. Cambridge, MA: MIT Press.

Perry, Motty, and Philip J. Reny. 2006. “Toward a Strategic Foundation for Rational Expectations Equilibrium.” *Econometrica* 74:1231–69.

Pratt, John W. 2007. “Fair (and Not So Fair) Division.” *J. Risk and Uncertainty* 35 (3): 203–36.

Prendergast, Canice, and Lars Stole. 1999. “Restricting the Means of Exchange within Organizations.” *European Econ. Rev.* 43:1007–19.

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

Roth, Alvin E. 1984. “The Evolution of the Labor Market for Medical Interns and Residents: A Case Study in Game Theory.” *J.P.E.* 92:991–1016.

———. 2002. “The Economist as Engineer: Game Theory, Experimentation, and Computation as Tools for Design Economics.” *Econometrica* 70 (4): 1341–78.

———. 2007. “Repugnance as a Constraint on Markets.” *J. Econ. Perspectives* 21 (3): 37–58.

Roth, Alvin E., and Elliott Peranson. 1999. “The Redesign of the Matching Market for American Physicians: Some Engineering Aspects of Economic Design.” *A.E.R.* 89 (September): 748–80.

Roth, Alvin E., and Marilda Sotomayor. 1990. *Two-Sided Matching*. Cambridge: Cambridge Univ. Press.

Sen, Amartya. 1979. "Equality of What?" Tanner Lecture on Human Values. Repr. in *Choice, Welfare and Measurement*. Cambridge, MA: Harvard Univ. Press, 1982.

———. 2009. *The Idea of Justice*. Cambridge, MA: Harvard Univ. Press.

Shulman, Robin. 2008. "Contention over Proposal for N.Y. Airports: Transportation Dept. Wants to Auction Off Takeoff, Landing Slots to Ease Delays." *Washington Post*, September 1.

Sönmez, Tayfun. 1999. "Strategy-Proofness and Essentially Single-Valued Cores." *Econometrica* 67 (3): 677–89.

Sönmez, Tayfun, and Utku Ünver. 2003. "Course Bidding at Business Schools." Working paper, http://ideas.repec.org/p/wpa/wuwpga/0306001.html.

———. 2010. "Course Bidding at Business Schools." *Internat. Econ. Rev.* 51 (1): 99–123.

———. 2011. "Matching, Allocation and Exchange of Discrete Resources." In *Handbook of Social Economics*, vol. 1A, edited by Jess Benhabib, Alberto Bisin, and Matthew O. Jackson, 781–852. Amsterdam: North-Holland.

Starr, Ross M. 1969. "Quasi-Equilibria in Markets with Non-convex Preferences." *Econometrica* 37 (1): 25–38.

Steinhaus, Hugo. 1948. "The Problem of Fair Division." *Econometrica* 16 (1): 101–4.

Svensson, Lars. 1983. "Large Indivisibles: An Analysis with Respect to Price Equilibrium and Fairness." *Econometrica* 51 (4): 939–54.

Thomson, William, and Hal R. Varian. 1985. "Theories of Justice Based on Symmetry." In *Social Goals and Social Organization: Essays in Memory of Elisha Pazner*, edited by Leonid Hurwicz, David Schmeidler, and Hugo Sonnenschein. Cambridge: Cambridge Univ. Press.

Varian, Hal. 1974. "Equity, Envy and Efficiency." *J. Econ. Theory* 29 (2): 217–44.

Vazirani, Vijay V. 2007. "Combinatorial Algorithms for Market Equilibria." In *Algorithmic Game Theory*, edited by Noam Nisan, Tim Roughgarden, Éva Tardos, and Vijay V. Vazirani. Cambridge: Cambridge Univ. Press.

Wallis, W. D., Anne Penfold Street, and Jennifer Seberry Wallis. 1972. *Combinatorics: Room Squares, Sum-Free Sets, Hadamard Matrices.* Lecture Notes in Mathematics, vol. 292. Berlin: Springer-Verlag.

Wharton. 2011. "The Course Registration Auction for MBA Electives." Version 18.0. Univ. Pennsylvania, Wharton School, Grad. Div., MBA Program Office, http://spike.wharton.upenn.edu/mbaprogram/auction/auction_description_v18.pdf.

Wulf, William A. 1993. "The Collaboratory Opportunity." *Science* 261 (5123): 854–55.

Zhou, Lin. 1990. "On a Conjecture by Gale about One-Sided Matching Problems." *J. Econ. Theory* 52 (1): 123–35.