

THE COMBINATORIAL ASSIGNMENT PROBLEM: APPROXIMATE COMPETITIVE EQUILIBRIUM FROM EQUAL INCOMES

ERIC BUDISH

ABSTRACT. Most of what is known about the problem of combinatorial assignment – e.g., assigning schedules of courses to students – are impossibility theorems which indicate that there is no perfect mechanism. This paper proposes a new mechanism, *Approximate Competitive Equilibrium from Equal Incomes* (A-CEEI), and shows that it satisfies attractive second-best criteria of efficiency, fairness, and incentives. First, I prove existence of a novel approximation to CEEI: any strictly positive amount of budget inequality is sufficient to guarantee existence of item prices which approximately clear the market. Second, I propose two new criteria of outcome fairness, *maximin share* and *envy bounded by a single good*, which weaken Steinhaus’s (1948) fair share and Foley’s (1967) envy freeness to accommodate indivisibilities and complementarities in a realistic way. Third, I show that A-CEEI satisfies the fairness criteria. Last, I show that A-CEEI is *strategyproof in the large*. I examine A-CEEI’s performance on real data and compare the proposed mechanism to alternatives from theory and practice: all other known mechanisms are either unfair ex-post or manipulable even in the large, and most are both manipulable and unfair.

KEYWORDS: combinatorial assignment, market design, course allocation, CEEI, indivisibilities, fair division, envy freeness, strategyproofness, dictatorship theorems.

Date: August 6, 2010.

University of Chicago Booth School of Business, Chicago IL 60637 (eric.budish@chicagobooth.edu). I am especially grateful to Susan Athey and Al Roth for their guidance and support through 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. 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, Scott Kominers, David Laibson, Robin Lee, Jacob Leshno, Greg Lewis, John Lewis, Paul Milgrom, Markus Mobius, Muriel Niederle, Ariel Pakes, Parag Pathak, David Parkes, Canice Prendergast, Amartya Sen, Neil Sloane, Ross Starr, Adi Sunderam and Glen Weyl, and seminar audiences at AMMA, BQGT, Chicago, Columbia, Harvard, LSE, Microsoft Research, NBER Market Design Workshop, NYU, 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 Project on Justice, Welfare and Economics at Harvard University, and the University of Chicago Booth School of Business.

1. INTRODUCTION

In a combinatorial assignment problem, a set of indivisible objects is to be allocated amongst a set of heterogeneous agents, the agents demand bundles of the objects, and monetary transfers are exogenously prohibited. A motivating example is course allocation at educational institutions: if, due to limits on class size, it is not possible for all students to take their most desired schedule of courses, then how should seats in over-demanded 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 which suggest that there is a particularly stark tension amongst 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, while intuitively sensible and widely used in single-unit assignment (Abdulkadiroğlu and Sönmez, 1998, 1999), seem unreasonably unfair in the multi-unit 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 non-dictatorial mechanisms, often citing fairness as a central design objective: e.g., Wharton’s course allocation system is “designed to achieve an *equitable* and *efficient* allocation of seats in elective courses when demand exceeds supply” (2009,

¹Press coverage and anecdotal evidence suggest that the scarcity problem is particularly acute in higher education, especially at professional schools. See Bartlett (2008), Guernsey (1999), Lehrer (2008), Levitt (2008a, b), and Neil (2008).

²On shifts, leads, players, airports, and scientific laboratories, respectively, see McKesson (2008), incentalign.com; Brams and Straffin (1979) and 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., due to union restrictions). Prendergast and Stole (1999) and Roth (2007) explore foundations for constraints against monetary transfers.

³The seminal reference is Vickrey (1961). See Milgrom (2004) and Cramton et al (2006) for textbook treatments that discuss both theory and applications.

⁴The seminal reference is Gale and Shapley (1962). See Roth and Sotomayor (1990) for a textbook treatment, and Roth (1984) on a well-known application.

⁵The seminal reference is Shapley and Scarf (1974), and some important mechanisms are described in Hylland and Zeckhauser (1979), Abdulkadiroğlu and Sönmez (1998, 1999) and Bogomolnaia and Moulin (2001). Applications to real housing and schooling markets are discussed in Chen and Sönmez (2002), Abdulkadiroğlu and Sönmez (2003) and Abdulkadiroğlu et al (2005a, 2005b, 2009). Another name for this problem is single-unit assignment.

⁶For precise statements see Klaus and Miyagawa (2001), Papai (2001), Ehlers and Klaus (2003) and Hatfield (2009). Konishi et al (2001) and Sönmez (1999) obtain related negative results under slightly different conditions, including existing endowments. Zhou (1990) and Kojima (2009) obtain related negative results for random mechanisms.

emphasis in original).⁷ But the mechanisms found in practice have a variety of flaws, most notably that they ignore incentives (Krishna and Ünver, 2008; Sönmez and Ünver, 2010; cf. Section 8.3). In one case it has been shown empirically that these incentive problems cause substantial inefficiency (Budish and Cantillon, 2009).

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, I show that 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 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

⁷Here are some additional examples: NYU Law School writes that its system “promotes a fair allocation of coveted classes” (Adler, 2008); MIT Sloan writes that its system “establish[es] a ‘fair playing field’ for access to Sloan classes”; 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 (2008) advertises its software product for assigning nurses to vacant shifts based on their preferences as “equitable open shift management.”

in *single*-unit assignment, for which they are often observed in practice (Abdulkadiroğlu and Sönmez, 1998, 1999). 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, so 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) based on 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*, i.e., 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); these mechanisms are manipulable even by the kinds of agents we usually think of as price takers, i.e., they are *manipulable in the large*.⁸

Some intuition for both the existence and fairness theorems can be provided by means of a simple example. Suppose there are two agents and four indivisible objects: two valuable diamonds (big and small) and two ordinary rocks (pretty and ugly). Agents require at most two objects each, and have the preferences we would expect given the objects’ names;⁹ think of the diamonds as seats in courses by star professors. If agents have the same budget then the market does not clear: at any price vector, for each object, either both agents demand the object or neither does. Notice too that discontinuities in aggregate demand are “large”: any change in price that causes one agent to change their demand causes the other agent to change their 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

⁸Note that in a variety of market-design contexts it has been shown that the limit incentive properties of a mechanism are a good approximation for the incentive properties of the mechanism in large finite markets. Examples of non-strategyproof mechanisms that are strategyproof in the large and thought to have attractive incentive properties in practice include double auctions (Rustichini et al, 1994; Perry and Reny, 2006) and deferred-acceptance algorithms (Roth and Peranson, 1999; Kojima and Pathak, 2009). Examples of mechanisms that are manipulable in the large and known to have important incentive problems in practice include non-stable matching algorithms (cf. Roth, 2002), the Boston mechanism for school choice (Abdulkadiroğlu et al, 2005a), and discriminatory-price multi-unit auctions (e.g., Friedman, 1991). The relationship between incentives in the limit and incentives in the finite is discussed further in Section 7.3.

⁹Specifically, each agent i has additive-separable preferences satisfying (i) $u_i(\{\text{big diamond}\}) > u_i(\{\text{small diamond}\}) > u_i(\{\text{pretty rock}\}) > u_i(\{\text{ugly rock}\})$; (ii) $u_i(\{\text{big diamond, ugly rock}\}) > u_i(\{\text{small diamond, pretty rock}\})$; and (iii) $u_i(\{\text{small diamond}\}) > u_i(\{\text{pretty rock, ugly rock}\})$.

that only the wealthier agent can afford the big diamond, while 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.¹⁰

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 while 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} while 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 is typically simple to manipulate (Budish and Cantillon, 2009). 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 the example an arbitrarily small amount of budget inequality is enough to ensure exact market clearing. In general, the existence theorem allows for at worst 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 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 non-convex preferences.¹¹ Second, the worst-case bound is economically small for life-sized problems, and of course average case is superior to worst case.¹²

¹⁰There also exist prices at which the wealthier agent gets {big diamond, pretty rock} while 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 Section 5.1 for further detail.

¹¹Starr’s (1969) result does not apply here, both for the technical reason that preferences in my environment are not continuous due to indivisibilities, and the substantive reason that approximately equal budgets are not well defined in exchange economies with indivisibilities. These differences necessitate a new proof technique, which ends up having the payoff of strengthening Starr’s bound. I am also able to prove that my bound is tight. Dierker’s (1971) economy accommodates indivisibilities, though not approximately equal wealth; his bound is substantially weaker than Starr’s. See further discussion in Section 3.3.

¹²In the specific context of course allocation, a small amount of market-clearing error is not too costly in practice, for reasons discussed in Section 3.3. In other contexts, such as assigning pilots to planes, market-clearing error is much more

A-CEEI compromises amongst 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 describes the properties of all other 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, 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 uses 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% of students have no envy, and for the remainder envy is small in utility terms. Last, the distribution of students’ utilities first-order stochastically dominates that from the actual play of Harvard’s own draft mechanism, which Budish and Cantillon (2009) 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 2 defines the environment. Section 3 defines Approximate CEEI and presents the existence theorem. Section 4 proposes the new criteria of outcome fairness. Section 5 provides the two fairness theorems. Section 6 formally defines the Approximate CEEI mechanism (A-CEEI). Section 7 discusses A-CEEI’s incentive properties. Section 8 compares A-CEEI to alternatives. Section 9 examines A-CEEI’s performance on real data. The paper concludes with open questions and a note on methodology. Proofs are in the body of the text when both short and instructive; otherwise they are in appendices. The text also contains a detailed sketch of the proof of the existence theorem.

costly. Two variants of the proposed mechanism for such contexts are discussed in an appendix. Both variants involve greater budget inequality, introduced in slightly different ways.

2. 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 $(\mathcal{S}, \mathcal{C}, (q_j)_{j=1}^M, (\Psi_i)_{i=1}^N, (u_i)_{i=1}^N)$, also called an *economy*, are defined as follows.

Agents. There is a set of N agents (students), $\mathcal{S} = \{1, \dots, i, \dots, N\}$.

Objects and Capacities. There is a set of M object types (courses), $\mathcal{C} = \{1, \dots, j, \dots, M\}$. There are $q_j \in \mathbb{Z}_+$ copies of object j (seats in course j). There are no other goods in the economy. In particular, there is no divisible numeraire like money.

Schedules and Preferences. A consumption bundle (schedule) consists of 0 or 1 seats in each course. The set of all possible schedules is the powerset of \mathcal{C} , i.e., $2^{\mathcal{C}}$. Each student i is endowed with a von-Neumann Morgenstern utility function that indicates her utility from each schedule of courses: $u_i : 2^{\mathcal{C}} \rightarrow \mathbb{R}_+$.¹³

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 $\{0, 1\}^M$.

In practical applications, each student will have a limited set of permissible schedules from which they generate positive utility. For instance, students take at most a certain number of courses per term, 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 a student’s utility from an impermissible schedule is zero, i.e., $u_i(x) = 0$ 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).

Most of the analysis works only with students’ ordinal preferences over schedules, and ignores the additional information in u_i about i ’s preferences over lotteries. I assume that students’ ordinal preferences over permissible schedules are strict, i.e., $u_i(x) \neq u_i(x')$ for any $x \neq x' \in \Psi_i$. I adopt $u_i : x, x', \dots$ as notational shorthand for $u_i(x) > u_i(x') > u_i(x'')$ for all $x'' \in 2^{\mathcal{C}} \setminus \{x, x'\}$.

No other restrictions are placed on the utility function: in particular, students are free to regard courses as complements and substitutes. This is the reason the assignment

¹³The 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 Section 3.4.1.

problem is called “combinatorial” as opposed to “multi-unit” (e.g., Budish and Cantillon, 2009). Note though that the formulation $u_i : 2^{\mathcal{C}} \rightarrow \mathbb{R}_+$ implicitly rules out peer effects.

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. A probability distribution over feasible allocations is *ex-ante Pareto efficient* if there is no other such lottery weakly preferred by all agents, with at least one strict.

The primary mechanism developed in this paper will be approximately feasible and approximately ex-post efficient in a sense that will be made clear in Section 3. Two variants that are exactly feasible and exactly ex-post efficient are also described, in an appendix; these variants are less attractive with respect to fairness and incentives. For more on the relationship between ex-post and ex-ante efficiency in this environment see Budish and Cantillon (2009) and Section 9.5.

3. THE APPROXIMATE COMPETITIVE EQUILIBRIUM FROM EQUAL INCOMES

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.¹⁴ Arnsperger (1994) writes “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).

Were it not for existence problems, CEEI would be an attractive solution to our problem of combinatorial assignment as well. When it exists, a CEEI is Pareto efficient by the first welfare theorem, is envy free because all agents have the same budget and face the same prices, and satisfies the maximin share guarantee (see Proposition 7). Further, a CEEI mechanism could be defined to satisfy the procedural fairness requirement of symmetry, and the approximate incentives criterion of strategyproof in the large (defined formally in Section 7).

Unfortunately in our economy CEEI need not exist. Either indivisibilities or complementarities alone can make existence problematic, and our economy features both. In order to recover existence we will approximate both the “CE” and the “EI” of CEEI.

¹⁴See Foley (1967), Varian (1974), and several other seminal references summarized in Thomson and Varian (1985).

3.1. Definition of Approximate CEEI.

Definition 1. Fix an economy. The allocation $\mathbf{x}^* = (x_1^*, \dots, x_N^*)$, budgets $\mathbf{b}^* = (b_1^*, \dots, b_N^*)$, and item prices $\mathbf{p}^* = (p_1^*, \dots, p_M^*)$ constitute an **(α, β) -Approximate Competitive Equilibrium from Equal Incomes $((\alpha, \beta)$ -CEEI)** if:

- (i) $x_i^* = \arg \max_{x' \in 2^c} [u_i(x') : \mathbf{p}^* \cdot x' \leq b_i^*]$ for all $i \in \mathcal{S}$
- (ii) $\|(z'_1(\mathbf{p}^*), \dots, z'_M(\mathbf{p}^*))\|_2 \leq \alpha$ where
 - (a) $z'_j(\mathbf{p}^*) \equiv \sum_i x_{ij}^* - q_j$ if $p_j^* > 0$
 - (b) $z'_j(\mathbf{p}^*) \equiv \max(\sum_i x_{ij}^* - q_j, 0)$ if $p_j^* = 0$
- (iii) $\min_i(b_i^*) = 1 \leq \max_i(b_i^*) \leq 1 + \beta$

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.¹⁵ Observe that agents consume sure bundles rather than probability shares of objects as in Hylland and Zeckhauser (1979); see Section 8.2 for a discussion of why I do not adopt the probability-shares approach. Also see Section 8.3 for a discussion of what goes wrong when the maximand in (i) is replaced by $[u_i(x') - \mathbf{p}^* \cdot x']$, as occurs in a widely used mechanism studied by Sönmez and Ünver (2010).

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. This market-clearing error will be discussed in detail in Section 3.3.

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, 2002; Vazirani, 2007).

¹⁵In the event that i cannot afford any of her permissible schedules Ψ_i at \mathbf{p}^* , break the tie by assigning i the empty set. Equivalently, we can assume that the empty set is in Ψ_i .

3.2. The Existence Theorem. Our main existence theorem is:

Theorem 1. *Fix an economy. 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 $(\frac{\sqrt{\sigma M}}{2}, \beta)$ -CEEI.*
- (2) *Moreover, for any $\beta > 0$, any budget vector \mathbf{b}' that satisfies $\min_i(b'_i) = 1 \leq \max_i(b'_i) \leq 1 + \beta$, and any $\varepsilon > 0$, there exists a $(\frac{\sqrt{\sigma M}}{2}, \beta)$ -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 $0, (1 + k), (1 + k)^2, \dots, (1 + k)^{N-1}$ implement the dictatorship (see further discussion in Section 8.1).

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, for each object, either all agents demand it or none do; demand is entirely unrelated to 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 $\frac{\sqrt{\sigma M}}{2}$. Part (2) of the theorem statement indicates that the market administrator can assign these close but unequal budgets to agents however she likes, but for an ε perturbation which can be made arbitrarily small. Two natural choices are (i) assign budgets uniform randomly; and (ii) assign budgets based on some pre-existing priority order, like seniority or grade-point average.

3.3. Discussion of Market Clearing Error. There are two senses in which $\frac{\sqrt{\sigma M}}{2}$ is “small”.

First, $\frac{\sqrt{\sigma M}}{2}$ 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 towards exact market clearing, in the sense that error goes to zero as a fraction of the endowment. This notion of approximate market clearing was emphasized in the literature on general equilibrium with non-convexities (e.g., Starr 1969, Dierker 1971, Arrow and Hahn 1971). But, see also Anderson et al (1982) and Geller (1986) for bounds that grow with market size.

Second, $\frac{\sqrt{\sigma M}}{2}$ 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 5 courses each and there are about 50 courses overall, so $\frac{\sqrt{\sigma M}}{2} \approx 11$. Furthermore, if we have preference data we may be able to determine that some courses are in substantial excess

supply, and we can reformulate the problem as one of allocating only the potentially scarce courses.¹⁶ In the HBS data described in Section 9, only about 20 courses per semester are ever scarce, so in the reformulated problem the bound becomes $\frac{\sqrt{\sigma M}}{2} \approx 7$. This corresponds to a maximum market-clearing error of 7 seats in one class, or of 2 seats in each of 12 classes (since $\sqrt{12 \cdot 2^2} \approx 7$), etc., as compared with about 4500 course seats allocated each semester.

I show below that the $\frac{\sqrt{\sigma M}}{2}$ bound is tight. For comparison, Starr’s (1969) bound, developed in the context of a divisible-goods exchange economy with continuous but non-convex preferences, would be $\frac{M}{2}$ 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.

As we will see in the proof sketch, if I endowed agents with exactly equal incomes market-clearing error might be as large as $N\sqrt{\sigma}$. The simple idea of introducing a small amount of budget inequality reduces market-clearing error to $M\sqrt{\sigma}$. The bound $M\sqrt{\sigma}$ does not grow with the number of agents, but it likely is not a compelling worst-case guarantee for practice. For instance, in an HBS-sized problem $M\sqrt{\sigma} \approx 158$, which corresponds to 22 seats in each of the 50 classes. The hard part of the proof is in further reducing the bound to $\frac{\sqrt{\sigma M}}{2}$.

While perfect market clearing would obviously be preferable, a small amount of error may not be especially costly in the specific context of course allocation. 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.

Two variants of the proposed mechanism that clear the market without error are discussed in Appendix E. Each variant involves greater budget inequality, introduced in slightly different ways.

¹⁶Write $\mathcal{C} = \mathcal{C}^{scarce} \cup \mathcal{C}^{unscarce}$. In the reformulated problem, students’ preferences are defined over subsets of \mathcal{C}^{scarce} , with student i ’s utility from scarce-course bundle $x \subseteq \mathcal{C}^{scarce}$ equal to i ’s utility in the original problem from the best bundle he can form by adding a subset of $\mathcal{C}^{unscarce}$ to x .

3.4. 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 proceed directly to Section 4.

3.4.1. Demand Discontinuities. The basic difficulty for existence is that agents' demands are discontinuous with respect to price. 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 involves at most $\min(2k, M) \equiv \sigma$ objects, and has Euclidean distance of at most $\sqrt{\sigma}$.¹⁷

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}.

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 existence of market-clearing item prices in certain combinatorial auction environments.¹⁸

3.4.2. 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 helpful.

The story is a bit more complicated because our economy uses M item prices, not 2^M bundle prices. Let $H(i, x) = \{\mathbf{p} : \mathbf{p} \cdot x = b_i\}$ denote the hyperplane in M -dimensional nominal price space along which agent i can exactly afford bundle x . Every time price

¹⁷More generally, we can define $\sqrt{\sigma} \equiv \sup_{i, \mathbf{p}} \lim_{\delta \rightarrow 0^+} \sup_{\mathbf{p}' \in B_\delta(\mathbf{p})} \|x_i^*(\mathbf{p}) - x_i^*(\mathbf{p}')\|_2$, with $x_i^*(\mathbf{p}) \equiv \arg \max_{x' \in 2^{\mathcal{C}}} [u_i(x') : \mathbf{p} \cdot x' \leq$

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

¹⁸Parkes (2006) provides a survey of the use of monotone price path techniques; see also Milgrom (2000). Mongell and Roth (1986) observe that budget constraints induce complementarities in the context of auctions.

crosses such a *budget-constraint hyperplane* (b-c-h), some agent's choice set changes, and hence their demand might change.

Importantly, the number of b-c-h's is finite, because the number of agents (N) and the number of bundles (2^M) are finite. This is an advantage of having only indivisible goods. As long as each agent has a unique budget, the number of agents' b-c-h's intersecting at any one point is generically at most M , the dimensionality of the price space; a perturbation scheme ensures that there are in fact no L -way intersections with $L > M$.¹⁹ Now, the maximum discontinuity in demand with respect to price is $M\sqrt{\sigma}$, which no longer grows with N . (*See Step 1 of the formal proof*)

3.4.3. *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

$$(3.1) \quad f(\mathbf{p}) = \mathbf{p} + \mathbf{z}(\mathbf{p})$$

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 $f(\cdot)$:

$$(3.2) \quad F(\mathbf{p}) = co\{y : \exists \text{ a sequence } \mathbf{p}^w \rightarrow \mathbf{p}, \mathbf{p}^w \neq \mathbf{p} \text{ such that } f(\mathbf{p}^w) \rightarrow y\}$$

where co denotes the convex hull. The correspondence $F(\cdot)$ smoothes 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 (3.2) are upper-hemicontinuous. To apply the Cromme and Diener result, we need to specify the set on which $f(\cdot)$ is defined. This is subtle because we need the set of legal prices to respect a lower bound of zero, but we also need $f(\mathbf{p})$ to stay within the price space even when some object's price is small yet its excess demand is negative. This is handled in the proof by defining two price spaces: a legal price space $\mathcal{P} = [0, 1 + \beta + \varepsilon]^M$, and an auxiliary enlargement of this space, on which $f(\cdot)$ and $F(\cdot)$ are defined. For the remainder of the proof sketch, we will ignore the distinction between the actual and the auxiliary price space.

Once we have upper-hemicontinuity of $F(\cdot)$ we can apply Kakutani's fixed-point theorem (the other conditions are trivially satisfied): there exists \mathbf{p}^* such that $\mathbf{p}^* \in F(\mathbf{p}^*)$

¹⁹This perturbation scheme is the reason Theorem 1 allows \mathbf{b}^* to differ from \mathbf{b}' pointwise by $\varepsilon > 0$, as opposed to simply requiring that no two budgets in \mathbf{b}' are equal and then setting $\mathbf{b}^* = \mathbf{b}'$.

(*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 $\frac{\sqrt{\sigma M}}{2}$.

3.4.4. *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 b-c-h, 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 $L \leq M$ b-c-h's.

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 5-7*).

Second, for each of the L agents corresponding to the L hyperplanes, their demand depends only on which side of their *own* b-c-h 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 b-c-h.²⁰ 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 b-c-h (weakly), the set of feasible demands in an arbitrarily small neighborhood of \mathbf{p}^* can be written as: (*Step 8*)

$$(3.3) \quad \left\{ (a_1, \dots, a_L) \in \{0, 1\}^L : \mathbf{z}(\mathbf{p}^*) + \sum_{l=1}^L a_l v_l \right\}$$

3.4.5. *Obtaining the Bound*. Now $\mathbf{p}^* \in F(\mathbf{p}^*)$ tells us something very useful: perfect market clearing is in the convex hull of (3.3). Our market-clearing error is the maximum-minimum distance between a vertex of (3.3) – one of the feasible demands near \mathbf{p}^* – and a point in the convex hull of (3.3). Either a probabilistic method argument or the Shapley-Folkman theorem can be used to verify that the worst case occurs when (3.3) is

²⁰There 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 of the 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.

an M -dimensional hypercube of side length $\sqrt{\sigma}$, and the perfect market clearing ideal is equidistant from all 2^M vertices of (3.3). Half the diagonal length of such a hypercube is $\frac{\sqrt{\sigma M}}{2}$ (Steps 9-10).

3.5. 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 any $\alpha < \frac{\sqrt{\sigma M}}{2}$ and $\beta \gtrsim 0$, there does not exist an (α, β) -CEEI.*

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

3.6. 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 $(\mathcal{S}, \mathcal{C}, (q_j)_{j=1}^M, (\Psi_i)_{i=1}^N, (u_i)_{i=1}^N)$. 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 $(\mathcal{S}, \mathcal{C}, (q_j^*)_{j=1}^M, (\Psi_i)_{i=1}^N, (u_i)_{i=1}^N)$.*

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

4. 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.

There have been several previous approaches to defining outcome fairness in the presence of indivisibilities.²¹ First, many authors consider a simpler problem in which monetary transfers are permitted (Svensson, 1983; Maskin, 1987; Alkan et al, 1991; Moulin, 1995; Brams and Kilgour, 2001). 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

²¹Procedural fairness is not similarly problematic: tossing a fair coin to determine which agent gets the diamond satisfies the standard procedural fairness requirement of symmetry, also known as anonymity or equal treatment of equals. A mechanism violates symmetry if its treatment of agents depends not only on their reports but on their identities; see e.g. Moulin (1995) for a formal definition. My proposed mechanism is symmetric, as are most mechanisms found in practice.

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; Pratt, 2007). If we award each agent the lottery in which they receive each object with probability one-half, then neither agent envies the other agent's *lottery*.

The common thread in all of 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.²² 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) describes as “the two most important tests of equity” – and propose the *maximin share guarantee* and *envy bounded by a single good*, respectively.

The proposed criteria allow for complementarities as well, and in particular accommodate cases where different agents regard different objects as complementary. The extant fairness literature typically assumes that agents' preferences are additive separable or convex.

4.1. 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 divisible goods, an allocation \mathbf{x} satisfies the *fair share guarantee* if $u_i(x_i) \geq u_i(\frac{\mathbf{q}}{N})$ for all i . Early papers on the cake-cutting problem (Steinhaus 1948, Dubins and Spanier 1961) actually defined fairness itself as this guarantee. The appeal of the fair share guarantee is that it expresses the ideal of common ownership of the goods which are to be divided. Agents might expect to do better than this ideal, due to heterogeneity in preferences, but certainly they should not do worse.

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

²²Another 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.”

Definition 2. Fix an economy. Agent i 's maximin share utility is

$$(4.1) \quad \max_{(x_l)_{l=1}^N} [\min[u_i(x_1), \dots, u_i(x_N)]] \quad \text{subject to}$$

$$x_l \in 2^{\mathcal{C}} \quad \text{for all } l = 1, \dots, N$$

$$\sum_l x_{lj} \leq q_j \quad \text{for all } j \in \mathcal{C}.$$

If her maximin share utility is strictly positive, then the bundle which gives her this utility is her **maximin share**; else, her maximin share is \emptyset . Any allocation that solves i 's program (4.1) is said to be a **maximin split** for agent i . Any allocation in which all N agents get an allocation they weakly prefer to their own maximin share is said to satisfy the **maximin share guarantee**.

Maximin shares can be interpreted as the outcome of a certain divide-and-choose procedure; two-player divide-and-choose is perhaps the oldest known method of fair division, with accounts of its use appearing in the old testament and in Greek mythology.²³ Imagine that agent i is to divide the set of indivisible goods into N bundles. Each of the other $N - 1$ agents, in turn, chooses one of the bundles. Agent i then receives the one remaining bundle. His maximin share is the outcome he obtains from this procedure in Nash equilibrium when either (i) the other agents' preferences are identical to his own; or (ii) the other agents are adversaries who seek to leave i with the worst possible outcome. In either of these environments, the agent will propose a division such that his least favorite of the N bundles is as attractive as possible.

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.

Agents with different preferences will have different maximin shares, whereas, in divisible-goods problems, every agent's fair share is $\frac{1}{N}$ of the endowment, regardless of their preferences. Despite this apparent difference the concepts are closely related, and the original fair share guarantee can also be interpreted in terms of divide-and-choose:

Proposition 3. Consider a divisible-goods version of the combinatorial assignment problem in which $\Psi_i = [0, q_1] \times \dots \times [0, q_M]$ and $u_i : [0, q_1] \times \dots \times [0, q_M] \rightarrow \mathbb{R}_+$, for all $i \in \mathcal{S}$. Modify Definition 2 to require that each $x_l \in [0, q_1] \times \dots \times [0, q_M]$ rather than $x_l \in 2^{\mathcal{C}}$. If preferences are convex and monotonic, then each agent's fair share is a maximin share, uniquely so if preferences are strictly convex.

²³See Crawford (1977) and Brams and Taylor (1996); see also Moulin (1991, 1995).

Suppose that two agents are to divide two diamonds and two rocks, and that each agent consumes at most two objects. If the two diamonds are identical and the two rocks are identical, then each agent's maximin share will be the bundle {diamond, rock}, exactly his fair share. Suppose instead that the diamonds and rocks are heterogeneous: say, a big and a small diamond, and a pretty and an ugly rock. Now, each agent's maximin share will be their less preferred of bundles {big diamond, ugly rock} and {small diamond, pretty rock}. Some degree of ex-post unfairness is inevitable given the indivisibilities in this economy; only one agent can get the lone big diamond. The maximin share guarantee ensures that the agent who does not get the big diamond at least gets his favorite of what remains.

Complementarities can make it technologically infeasible to allocate all agents at least their maximin share.²⁴ Let an agent's N' -maximin share be her exact maximin share in a hypothetical economy with the same objects and capacities, but with $N' \geq N$ total agents.

In a dictatorship, if there are N agents choosing exactly k objects each, the last agent to choose potentially gets his k least favorite out of Nk objects. This can be an outcome worse than what he could guarantee himself as divider in a $k(N - 1)$ -way divide-and-choose procedure. This observation, together with Theorem 1 of Klaus and Miyagawa (2001),²⁵ directly yields the following impossibility result.

Proposition 4. *There is no combinatorial assignment mechanism that is strategyproof, ex-post efficient, and which always yields allocations that satisfy the $k(N - 1)$ -maximin share guarantee, with k the maximum number of objects per agent as defined in Theorem 1.*

Any procedure that yields a better fairness guarantee than the dictatorship must compromise on other attractive criteria. Theorem 2 shows that the Approximate CEEI mechanism always satisfies the $(N + 1)$ -maximin share guarantee.

4.2. Envy Bounded by a Single Good. An allocation \mathbf{x} is said to be *envy free* if $u_i(x_i) \geq u_i(x_{i'})$ 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.

²⁴For instance, suppose that there are two agents, $\mathcal{S} = \{i_1, i_2\}$, and four objects, $\mathcal{C} = \{a, b, c, d\}$, each in unit supply, and that i_1 's preferences are $u_{i_1} : \{a, b\}, \{c, d\}, \dots$, while i_2 's preferences are $u_{i_2} : \{a, d\}, \{b, c\}, \dots$. If i_1 gets a bundle he likes at least as well as his maximin share $\{c, d\}$ then it is not possible for i_2 to get a bundle he likes at least as well as his maximin share $\{b, c\}$.

²⁵Theorem 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 4 as stated. If we restrict attention to larger economies (i.e., $N > 2$), then we can obtain a slightly weaker statement than Proposition 4 by using either Proposition 1 of Papai (2001) or Theorem 1 of Ehlers and Klaus (2003), each of which characterize dictatorships in terms of strategyproofness, ex-post efficiency, and a mild additional property (non-bossiness and coalitional strategyproofness, respectively).

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) $u_i(x_i) \geq u_i(x_{i'})$ or
- (ii) *There exists some good $j \in x_{i'}$ such that $u_i(x_i) \geq u_i(x_{i'} \setminus \{j\})$*

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. This bounds the magnitude of any agent's envy by their maximum marginal utility for a single object, i.e., by $\max_{x \in 2^{\mathcal{C}}, j \in x} [u(x) - u(x \setminus \{j\})]$. Note that if a set of objects is highly complementary – e.g., pieces of a puzzle – then the marginal utility for any single object will be similar to the total utility of the set. As the maximum marginal utility becomes arbitrarily small, envy bounded by a single good becomes arbitrarily similar to exact envy freeness. Formally:

Proposition 5. *If $\max_{x \in 2^{\mathcal{C}}, j \in x} [u_i(x) - u_i(x \setminus \{j\})] < \delta$ for all $i \in \mathcal{S}$, and $\mathbf{x}^* = (x_i^*)_{i \in \mathcal{S}}$ satisfies envy bounded by a single good, then $[u_i(x_{i'}^*) - u_i(x_i^*)] < \delta$ for all $i, i' \in \mathcal{S}$.*

Dictatorships fail to bound envy in the sense of Definition 3. If there are two diamonds, Definition 3 requires that each agent gets one diamond, but in a dictatorship some agent gets both. This directly implies the following result, analogous to Proposition 4:

Proposition 6. *There is no combinatorial assignment mechanism that is strategyproof, ex-post efficient, and which always yields allocations that satisfy envy bounded by a single good.*

Theorem 3 shows that the Approximate CEEI mechanism always satisfies envy bounded by a single good.

4.3. Discussion: Dictatorships and Fairness. In combinatorial assignment problems, dictatorships violate the maximin share guarantee and envy bounded by a single good. The criteria formalize the common intuition that dictatorships are unfair in contexts like course allocation.

In contrast to combinatorial assignment, dictatorships are commonly observed in practice for *single-unit* assignment, e.g., for school choice and house-allocation problems (see

the references in footnote 5). But in single-unit assignment – e.g., two agents dividing one diamond and one rock – we have the following simple observation.

Remark 1. *In single-unit assignment problems, dictatorships satisfy the maximin share guarantee and envy bounded by a single good.*

Hence, the properties help us to make sense of the empirical patterns of dictatorship usage.²⁶

5. FAIRNESS THEOREMS FOR APPROXIMATE CEEI

In a traditional general equilibrium economy agents' utilities are continuous with respect to their budget. So we would expect an arbitrarily small amount of budget inequality to have an arbitrarily small effect on properties of agents' outcomes.

By contrast, in our economy agents' utilities are discontinuous with respect to their budget. Moreover, an agent's most preferred affordable bundle need not exhaust his budget. So it is not obvious how to think about the arbitrarily small, but strictly positive amount of budget inequality required in Theorem 1.

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 that for small enough β , an (α, β) -CEEI guarantees that envy is bounded by a single good.

5.1. 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 the main result of this section, which is that the Approximate CEEI of Theorem 1 guarantees an approximation to maximin shares that is based on adding one more agent to the economy.

Proposition 7. *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 $u_i(x_i^*) < u_i(x_i^{MS})$ for each $x_i^{MS} \in \mathbf{x}^{MS}$. By conditions (i) and (iii) of Definition 1 we have $\mathbf{p}^* \cdot x_i^{MS} > b_i^*$ for each $x_i^{MS} \in \mathbf{x}^{MS}$, and $\mathbf{p}^* \cdot x_i^* \leq b_i^*$ for each $x_i^* \in \mathbf{x}^*$, 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:

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

□

²⁶See also Abdulkadiroğlu and Sönmez (1998) for a normative argument in favor of the Random Serial Dictatorship for the case of single-unit demand.

The proof of Proposition 7 relies on two facts about $(0, 0)$ -CEEs: (i) $\beta = 0$ means that each agent has $\frac{1}{N}$ of the budget endowment; (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 (i). Issue (ii) is handled with the following approximation parameter and minor extension of Theorem 1.

Definition 4. *Fix an economy.* 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. *Fix an economy.* 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 $\mathbf{p}^* \in \mathcal{P}(\delta, \mathbf{b}^*)$.

The key to the proof of Lemma 1 is that the $(\frac{\sqrt{\sigma M}}{2}, \beta)$ -CEEI price vector \mathbf{p}^* guaranteed to exist by Theorem 1 is near a fixed point of the correspondence $F(\cdot)$ defined informally in (3.2). At each price near to the fixed point agents can afford their demands, and a convex combination of these demands is feasible. Hence, since \mathbf{p}^* is nearly a fixed point of $F(\cdot)$, agents can approximately afford the endowment at \mathbf{p}^* .

By choosing β, δ small enough we can ensure that each agent's budget is at least $\frac{1}{N+1}$ of the cost of the endowment at the Approximate CEEI price vector \mathbf{p}^* .

Theorem 2. *Fix an economy.* 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 $\beta < \frac{1-\delta N}{N(1+\delta)}$, then \mathbf{x}^* satisfies the $(N+1)$ -maximin share guarantee.

We might worry, especially in small markets, about the difference between exact maximin shares and $(N+1)$ -maximin shares. For instance, in the two diamonds and two rocks example described in the introduction, the exact maximin share guarantees each agent a diamond, whereas the $(N+1)$ -maximin share does not. Fortunately, we can often provide a slightly stronger guarantee than Theorem 2.

Proposition 8. *In any (α, β) -CEEI with $\alpha = 0$ and $\beta < \frac{1}{N-1}$, each agent is guaranteed the weaker of:*

- (1) *an outcome weakly better than her N^{th} favorite bundle in any $(N+1)$ -maximin split.*
- (2) *an outcome strictly better than her $(N+1)$ -maximin share (i.e., her $(N+1)^{\text{st}}$ favorite bundle in any $(N+1)$ -maximin split).*

Proposition 8 guarantees that each agent receives a diamond in the example described in the introduction. The approximation error is that the agent who gets the small diamond may get the ugly rock.²⁷

5.2. 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, and 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. *Fix an economy. If $[\mathbf{x}^*, \mathbf{b}^*, \mathbf{p}^*]$ is an (α, β) -CEEI with $\beta < \frac{1}{k-1}$, 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 $j_1, \dots, j_{k'}$. Then we have:

$$\begin{aligned} u_i(x_{i'}^* \setminus \{j_1\}) &> u_i(x_i^*) \\ &\vdots \\ u_i(x_{i'}^* \setminus \{j_{k'}\}) &> u_i(x_i^*) \end{aligned}$$

Condition (i) of Definition 1 indicates that i cannot afford any of these k' bundles formed by removing an object from $x_{i'}^*$:

$$\begin{aligned} \mathbf{p}^* \cdot (x_{i'}^* \setminus \{j_1\}) &> b_i^* \\ &\vdots \\ \mathbf{p}^* \cdot (x_{i'}^* \setminus \{j_{k'}\}) &> b_i^* \end{aligned}$$

Since $p_1^* + p_2^* + \dots + p_{k'}^* = \mathbf{p}^* \cdot x_{i'}^* \leq b_{i'}^*$ we can sum these inequalities to obtain

$$(k' - 1)b_{i'}^* \geq (k' - 1)(\mathbf{p}^* \cdot x_{i'}^*) > k'b_i^*$$

which implies that $\frac{b_{i'}^*}{b_i^*} \geq \frac{k'}{k-1}$. Since $k' \leq k$ we have $\frac{b_{i'}^*}{b_i^*} \geq \frac{k}{k-1}$. So if $\beta < \frac{1}{k-1}$ we have a contradiction. \square

Note that budget inequality plays slightly different roles in the two proofs. In Theorem 2, $\beta < \frac{1-\delta N}{N(1+\delta)}$ ensures that the poorest agent's budget is at least $\frac{1}{N+1}$ of the cost of the

²⁷For these examples it is easy to see that there exists a $(0, \beta)$ -CEEI for any $\beta > 0$. Label the big diamond as a , the small diamond as b , the pretty rock as c , and the ugly rock as d . Each agent's $(N+1)$ -Maximin Split is $\{\{a\}, \{b\}, \{c, d\}\}$, so part (i) of the guarantee is a bundle weakly better than $\{b\}$, and part (ii) of the guarantee is a bundle strictly better than $\{c, d\}$ (which happens to coincide with part (i); this is not the case generally). Market clearing then implies that each agent gets a bundle at least as good as $\{b, d\}$.

goods endowment. In Theorem 3, $\beta < \frac{1}{k-1}$ ensures that the poorest agent has at least $\frac{k-1}{k}$ of the richest agent's budget. Theorem 1 allows the market administrator to choose β arbitrarily small, so we can ensure that we simultaneously satisfy the requirements of both Fairness Theorems.

6. THE APPROXIMATE CEEI MECHANISM (A-CEEI)

Based on the efficiency and fairness results of Sections 3-5, and with an eye to incentives as will be discussed in Section 7, I propose the Approximate Competitive Equilibrium from Equal Incomes mechanism (A-CEEI). Informally, the mechanism is: (i) students report their preferences over permissible schedules of courses; (ii) if there are one or more exact CEEI's, choose one randomly; (iii) else, randomly assign students approximately equal budgets of the artificial currency; (iv) Theorem 1 now guarantees existence of an Approximate CEEI: choose one randomly, restricting attention to those with the least-possible market-clearing error. Formally:

Mechanism 1. *Fix an economy. The **Approximate CEEI mechanism (A-CEEI)** is the following procedure:*

- (1) *Each agent i reports her utility function \hat{u}_i , with $\hat{u}_i(x) = 0$ for any $x \notin \Psi_i$, and $\hat{u}_i(x') \neq \hat{u}_i(x'')$ for any $x', x'' \in \Psi_i$. Modify the original economy by replacing each u_i with \hat{u}_i .*
- (2) *Compute the set of (0,0)-CEEIs. If this set is non-empty, choose uniform randomly from this set and announce the resulting $[\mathbf{x}^*, \mathbf{b}^*, \mathbf{p}^*]$.*
- (3) *If the set of (0,0)-CEEI allocations is empty, choose a target budget vector \mathbf{b}' by drawing each b'_i uniform randomly from $[1, 1 + \beta]$, for some $\beta \ll \min(\frac{1}{N}, \frac{1}{k-1})$.*
- (4) *Then, for $\varepsilon \gtrsim 0$, $\delta < 1 - N\beta$, and $\alpha \leq \frac{\sqrt{\sigma M}}{2}$, compute the set of (α, β) -CEEI allocations with price vectors in $\mathcal{P}(\delta, \mathbf{b}')$ and budget vectors \mathbf{b}^* that satisfy $\|\mathbf{b}^* - \mathbf{b}'\|_\infty < \varepsilon$. This set is guaranteed to be non-empty by Theorem 1 and Lemma 1. Choose from this set as follows. First, restrict attention to those with the smallest α . Second, restrict attention to those with $\mathbf{b}^* = \mathbf{b}'$, unless this leaves an empty set. Third, choose uniform randomly out of the remaining set. Announce $[\mathbf{x}^*, \mathbf{b}^*, \mathbf{p}^*]$.*

In Step (1) agents report their preferences over schedules. In theory the only restriction on preferences is that utility from an impermissible schedule is zero and that otherwise preferences are strict. In practice, the number of possible schedules can be quite large, and so the market administrator must provide a way for students to express their preferences

concisely. Othman et al (2010) propose a language that tries to balance expressiveness with conciseness.²⁸

Step (2) seeks an exact CEEI, which is particularly attractive but need not exist. In case no CEEI exists, Step (3) chooses a set of approximately equal incomes. These are selected uniform randomly to ensure that the mechanism satisfies the procedural fairness requirement of symmetry. The budget inequality parameter β can be arbitrarily small, and in particular should be set smaller than the relevant bounds from the Fairness Theorems.

Step (4) computes the set of (α, β) -CEEIs that have market-clearing error below the Theorem 1 bound and then selects from this set. The first part of the tiebreaking rule minimizes market-clearing error to the extent possible; the second part respects the market administrator’s exact choice of budgets, if possible; the third part, as well as the analogous rule in Step (2), helps ensure that prices and budgets are exogenous to zero-measure agents’ reports.

A-CEEI clears the market with a small amount of error; see the discussion in Section 3.3. Two variants of A-CEEI without market-clearing error are described in Appendix E. Each variant involves greater budget inequality, introduced in slightly different ways.

An important, if imprecisely defined concern in practical market design is “transparency”. The computation of A-CEEI prices is opaque (cf. Section 9.2), but the prices can then be announced publicly. The assignment of random budgets can be public as well. Perhaps most importantly, each student is allocated her most-preferred affordable schedule at the announced prices; that is, they can verify that their allocations are correct.

7. INCENTIVE PROPERTIES OF THE APPROXIMATE CEEI MECHANISM

The dictatorship theorems suggest that the most attractive mechanisms for combinatorial assignment will not be strategyproof, i.e., dominant-strategy incentive compatible. Sections 7.1-7.2 show that A-CEEI is *strategyproof in the large (SPITL)*, which requires that truth-telling is a dominant strategy for the zero-measure agents economists traditionally think of as price takers. Section 7.3 discusses the relationship between incentives in the limit and incentives in finite markets.

7.1. Strategyproof in the Large (SPITL). A mechanism is said to be strategyproof in the large (SPITL) if it is exactly strategyproof in a certain continuum economy.

²⁸Students report their ordinal preferences over permissible schedules using a certain linear program. First, students report a single item value for each of the M courses. If their preferences are additive-separable but for scheduling constraints, these M numbers can be used to generate ordinal preferences over any Ψ_i . Next, students enhance their report by specifying that certain sets of courses are either complements or substitutes: e.g., courses A and B go well together so their total value is $v_A + v_B + v_{AB}$, with the complementarity term $v_{AB} > 0$; or, I want at most two courses in Finance, so any schedule with three or more Finance courses has utility of zero. In principle, students can express arbitrary ordinal preferences over permissible schedules, so the Othman et al (2010) language is fully expressive in the sense of Nisan (2006).

Definition 5. The *continuum replication* of economy $(\mathcal{S}, \mathcal{C}, (q_j)_{j=1}^M, (\Psi_i)_{i=1}^N, (u_i)_{i=1}^N)$, written, $(\mathcal{S}^\infty, \mathcal{C}, (q_j)_{j=1}^M, (\Psi_i)_{i \in \mathcal{S}^\infty}, (u_i)_{i \in \mathcal{S}^\infty})$ is constructed as follows:

- The set of agents is $\mathcal{S}^\infty = (0, N]$
- The set of objects and their capacities are left as is, except now we understand an object's capacity constraint as a Lebesgue measure.
- Agent $i \in \mathcal{S}^\infty$ in the continuum-replication economy has the same permissible-schedule set and utility function as agent $\lceil i \rceil$ in the original economy, where $\lceil \cdot \rceil$ is the ceiling operator. That is, agents numbered $(0, 1]$ in the continuum are identical to agent 1 in the original, agents numbered $(1, 2]$ in the continuum are identical to agent 2 in the original, etc.

In words, each atomic agent in the original economy is replaced with a continuum of measure one of agents identical to him. Note that the distribution of agent types and the set of object types are the same in the continuum replication as they are in the original finite economy.²⁹ This allows each type of agent's strategy set to be the same in the continuum as in the original, and means that any distribution of agents' reports that is feasible in the finite economy is also feasible in the continuum.

Definition 6. A mechanism is *strategyproof in the large (SPITL)* if it is exactly strategyproof in the continuum replication of any finite economy. Otherwise it is *manipulable in the large*.

While it typically is obvious how to define the continuum-economy version of finite mechanisms that use prices,³⁰ it can be subtle to define the continuum-economy version of matching and assignment algorithms. Recent contributions by Abdulkadiroğlu, Che and Yasuda (2008) and Che and Kojima (2010) provide a methodology for this task, as well as explicit continuum definitions of many important matching and assignment mechanisms.

7.2. Theorem 4: A-CEEI is SPITL.

Theorem 4. *The Approximate CEEI mechanism is strategyproof in the large.*

The proof of Theorem 4 is straightforward and contained in Appendix C. One interesting detail is that results in Mas-Colell (1977) and Yamazaki (1978) can be adapted to provide a much more direct proof of Theorem 1 in the continuum economy.

²⁹Definition 4 combines elements of the classic Debreu and Scarf (1963) and Aumann (1964) conceptions of a large market. As in Debreu and Scarf (1963) the number of types does not grow with market size, but as in Aumann (1964) we look at the continuum limit in which each agent is zero measure.

³⁰For instance, the statement of Mechanism 1 directly accommodates both finite economies and continuum economies, once the statement of Definition 1 is extended appropriately.

It is possible to generalize Step (4) of the Approximate CEEI mechanism in certain ways without jeopardizing SPITL. For instance, the market administrator might specify a penalty function in terms of α and β (or more complicated statistics of market-clearing error and budget inequality) and seek the best such (α, β) -CEEI. What matters is that the tiebreaking is based on aggregate features of the (α, β) -CEEI rather than individuals' reports and allocations, so that in a continuum economy the probability that any agent affects the market administrator's choice of prices or budgets is zero.

7.3. Discussion of SPITL. Definition 6 draws a conceptual distinction between two ways a mechanism can fail to be strategyproof (SP). The Approximate CEEI mechanism is not SP because atomic agents' reports might affect prices, and prices affect budget sets. But, agents for whom prices are exogenous – which is literally the case for zero-measure agents and may be a reasonable approximation in practice³¹ – have a dominant strategy of reporting truthfully. Equivalently, truthful reporting selects the most-preferred outcome from any realized budget set.

In the Bidding Points Auction studied by Sönmez and Ünver (2010) and described in Section 8.3, students' reports also might affect prices. In the Harvard Business School Draft mechanism studied by Budish and Cantillon (2009) and described in Section 9, students' reports might affect the algorithm time at which different courses reach capacity, which is that mechanism's analogue of prices. But each of these mechanisms is not SP for a second, conceptually distinct reason. In the auction, even a student who faces exogenous prices should misreport her preferences; see the example in Section 8.3. In the draft, even a student who faces exogenous capacity times should misreport; see Section 3 of Budish and Cantillon (2009). That is, in these mechanisms, truthful reporting does *not* select an agent's most preferred outcome from any realized budget set. Such mechanisms are not only manipulable in finite markets, they are manipulable in the large.

Empirical evidence suggests that the second kind of manipulability is more important in practice. In addition to the two course-allocation mechanisms mentioned above, other

³¹Intuition and evidence from other contexts suggests that price-taking behavior more closely approximates optimal behavior in markets that are large or in which agents are uncertain about the realized distribution of other agents' preferences. Market size will tend to reduce an agent's ability to affect price. Uncertainty will tend to increase the downside risk of misreporting, i.e., the risk that the agent is assigned something other than her most-preferred affordable bundle at the realized prices. An additional thing to note is that the kinds of manipulations that do exist in small markets are reasonably non-obvious, by contrast with, e.g., bid shading in a double auction or quantity shading in Cournot competition. Here is an example of the kinds of manipulations I have found in computational exploration. Suppose there are two agents and four objects, and that i_1 's true preferences are $u_{i_1} : \{b, c\}, \{a, c\}, \{c, d\}, \{a, b\}, \dots$, while i_2 's true preferences are $u_{i_2} : \{a, b\}, \{a, c\}, \{a, d\}, \{b, c\}, \dots$. If both agents report truthfully, then the allocations in which i_1 gets $\{b, c\}$ or $\{c, d\}$, while i_2 gets $\{a, d\}$ or $\{a, b\}$, each can be supported as exact CEEIs. A-CEEI will randomize between the two. If i_1 instead misreports her preferences as $\hat{u}_{i_1} : \{b, c\}, \{a, c\}, \{a, b\}, \{c, d\}, \dots$ she can cause the allocation in which she gets $\{b, c\}$ while i_2 gets $\{a, d\}$ to be the unique exact CEEI. That is, by pretending to favor $\{a, b\}$ over $\{c, d\}$, she increases the likelihood that at the realized prices she can afford $\{b, c\}$. The downside risk to this manipulation is that if she is wrong about i_2 's preferences she might end up being allocated the bundle that she is pretending to like instead of some other bundle that she actually likes.

mechanisms that are manipulable in the large and that have been shown to have important incentive problems in practice include non-stable matching algorithms (cf. Roth, 2002), the Boston mechanism for school choice (Abdulkadiroğlu et al, 2005a) and discriminatory price multi-unit auctions (e.g., Friedman, 1991). Deferred acceptance algorithms (Roth and Peranson, 1999; Kojima and Pathak, 2009) and double auctions (Rustichini et al, 1994; Perry and Reny, 2006) are well-known examples of non-strategyproof mechanisms that are strategyproof in the large and thought to have attractive incentive properties in practice.³² To the best of my knowledge, there are no empirical examples of market designs that are strategyproof in the large but which 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 thought to be truthful in large finite markets.

8. COMPARISON TO ALTERNATE MECHANISMS

Impossibility theorems indicate that there is no perfect mechanism for combinatorial assignment. The Approximate CEEI mechanism (A-CEEI) developed in Sections 3-7 offers a compromise of competing design objectives. It is approximately ex-post efficient (Theorem 1, Proposition 2), guarantees approximate maximin shares (Theorem 2), bounds envy by a single good (Theorem 3), and is strategyproof in the large (Theorem 4). Additionally, it satisfies the procedural fairness requirement of symmetry.

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

[Insert Table 1 Here]

Every known 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 all except the dictatorship and the descending demand procedure 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 Section 9) A-CEEI is the more efficient mechanism, both ex-ante and ex-post.

³²For both deferred acceptance and double auctions, as well as for 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 non-constructive and highly discontinuous (cf. Section 3.4), as well as somewhat random (cf. Section 6). Developing a better understanding of A-CEEI's incentive properties away from the limit is a natural topic for future research.

The remainder of this Section compares A-CEEI to three specific alternatives in more detail. The first, Random Serial Dictatorship, is the mechanism suggested by the extant theory literature on combinatorial assignment. The second, a multi-unit generalization of Hylland and Zeckhauser’s (1979) probability-shares CEEI mechanism (HZ), is a natural candidate for an attractive solution to the combinatorial assignment problem, given the attractiveness of HZ for the single-unit case. The third, the Bidding Points Auction widely used for course allocation by business schools in the United States (Sönmez and Ünver, 2010), is easily confused with *exact* CEEI by the casual observer.

8.1. Comparison to Random Serial Dictatorship. Random Serial Dictatorship (RSD) can be interpreted as a competitive equilibrium mechanism. In the case of single-unit demand, unequal budgets play the same role as the serial order. An arbitrarily small amount of budget inequality suffices to implement the dictatorship, and RSD coincides with A-CEEI.

In the case of multi-unit demand, however, the magnitude of budget inequality matters: the analogue of being l^{th} in line in the dictatorship is to have a budget that dwarfs that of anyone later than l^{th} in line, and is dwarfed by that of anyone earlier than l^{th} in line. The following proposition, proved in Appendix D, summarizes the relationship.

Proposition 9. *In the case of single-unit demand, RSD coincides with A-CEEI. In general combinatorial assignment problems, RSD corresponds to a competitive equilibrium mechanism in which \mathbf{b}^* is a uniform random permutation of $\mathbf{b}^{\text{RSD}} = (1, (k+1), (k+1)^2, \dots, (k+1)^{N-1})$, with k the maximum number of objects per agent as defined in Theorem 1.*

Proposition 9 may help us to further understand why RSD is often observed in practice for single-unit assignment (e.g., school choice) but not for multi-unit assignment. See also Remark 1.

8.2. Comparison to the Multi-Unit Hylland Zeckhauser Mechanism. In a seminal paper, Hylland and Zeckhauser (1979) propose CEEI in “probability shares” as a solution to the single-unit assignment problem. Specifically, agents report von Neumann Morgenstern utility functions, and then the mechanism computes a set of prices such that, when each agent purchases the *lottery* over objects that maximizes their vNM utility subject to their (common) budget constraint, the market clears. The Birkhoff-von Neumann theorem then indicates that this lottery can be resolved into a deterministic allocation.

Note the contrast with A-CEEI, in which randomness enters via the budgets and agents purchase sure bundles of goods.

An HZ generalization is not possible in general due to limitations on what kinds of lotteries can be resolved into deterministic allocations. Suppose there are three agents, $\mathcal{S} = \{i_1, i_2, i_3\}$, and three objects, $\mathcal{C} = \{a, b, c\}$, each in unit supply. Agent i_1 derives positive utility only from the bundle $\{a, b\}$; agent i_2 only from the bundle $\{a, c\}$; and agent i_3 only from the bundle $\{b, c\}$. The lottery in which each agent obtains the bundle they like with probability one half clears the market in probability shares, and can be supported as a probability-shares CEEI. However, this lottery cannot be resolved: in the one-half probability event that agent i_1 gets $\{a, b\}$ it is not possible to give agent i_2 bundle $\{a, c\}$ or to give agent i_3 bundle $\{b, c\}$; and, similarly for the one-half probability events in which agents i_2 or i_3 get their desired bundles.

Pratt (2007) extends the HZ mechanism to the following special class of combinatorial assignment problems. First, each object j is in unit supply: $q_j = 1$ for all $j \in \mathcal{C}$. Second, each agent i is permitted to consume any bundle of objects: $\Psi_i = 2^{\mathcal{C}}$ for all $i \in \mathcal{S}$. Last, each agent i 's vNM utility function is additive-separable over objects. That is, for each $i \in \mathcal{S}$, there exists a vector of object values $v_i = (v_{i1}, \dots, v_{iM})$ such that, for any $x \in \Psi_i$, $u_i(x) = v_i \cdot x$. A generalization of the Birkhoff-von Neumann theorem given by Budish, Che, Kojima and Milgrom (2010) allows Pratt's first and second restrictions to be somewhat relaxed: objects can be in arbitrary integer supply, and permissible schedule sets can be any subset of $2^{\mathcal{C}}$ that satisfies a technical condition called *hierarchy*.

In this class of problems, the HZ mechanism has the same efficiency, procedural fairness and incentives properties as the original HZ: it is ex-ante efficient, symmetric and SPITL. It does not satisfy the outcome fairness criteria proposed herein, though in special cases it bounds envy by at most twice the value of a single good.

However, the class of combinatorial assignment problems for which these results apply may be restrictive in ways that matter for practice. First, the assumption that vNM preferences are additive-separable over objects is very strong, and creates the possibility that some students are allocated a lottery in which they sometimes receive zero desirable courses. For instance, in the diamonds and rocks example from the introduction, it generically will be the case that one student sometimes receives no diamonds, while the other student sometimes receives both.³³ Second, the restrictions on each u_i and Ψ_i together rule out many kinds of complementarities and substitutabilities (see Milgrom, 2009). Third, the restrictions on Ψ_i limit the kinds of schedule and curricular constraints that may be imposed by the market administrator.

³³Assume that the two agents' reports are such that they do not have identical marginal rates of substitution between the big diamond and the small diamond. Then whichever agent has the higher relative value for the big diamond will get the big diamond with probability greater than 0.5, but will get the small diamond with probability zero. The other agent always gets the small diamond, and sometimes gets both. The first agent may prefer the second agent's lottery (if he really dislikes getting zero diamonds), but the preference reporting language implicitly assumes otherwise.

A-CEEI avoids these three issues. First, it allocates sure bundles (schedules) rather than probability shares, and places ex-post guarantees on the attractiveness of those schedules (Theorems 2 and 3). Second, it allows students to express arbitrary ordinal preferences over permissible schedules. Third, it allows the market administrator to define students' permissible schedule sets without restriction.

8.3. Comparison to the Bidding Points Auction. Many prominent universities – including Berkeley, Chicago, Columbia, MIT, Michigan, NYU, Northwestern, Penn, Princeton, and Yale – use equal income artificial currency procedures to allocate courses to their students.³⁴ The procedures work 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.³⁵ Bids can be interpreted as reports of additive-separable utility functions, scaled to sum to the budgets. The q_j^{th} highest bid for course j is frequently described as the “clearing price” p_j for course j , and the allocation itself as a “market equilibrium” (e.g., Wharton, 2009).

To the casual observer, this Bidding Points Auction (BPA) looks like an equal-incomes competitive equilibrium procedure. A clue that something is funny is that we know exact CEEI need not always exist, yet the BPA yields a well-defined outcome for any set of preference reports.

It turns out the BPA makes two mistakes. First, it finds the wrong prices. Second, for a particular set of prices, it calculates the wrong demands.³⁶ Both mistakes share a common conceptual cause: the BPA treats fake money as if it were real money that enters the utility function. (Currency is fake if it has no value outside the problem at hand). That is, it treats a general-equilibrium problem as if it were an auction-theory problem.³⁷

³⁴Sönmez and Ünver (2010) describe the Bidding Points Auction used at the University of Michigan Ross School of Business, and close variants used at the Haas School of Business at UC Berkeley, Columbia Business School, Kellogg Graduate School of Management at Northwestern, Princeton University, and Yale School of Management. Adler et al (2008) and MIT (2008) describe similar procedures used at NYU Law and MIT Sloan, respectively. Graves et al (1993) describe the Primal-Dual Linear Programming procedure used until 2007 at the University of Chicago Graduate School of Business, which then adopted a Bidding Points Auction. See <https://ibid.chicagobooth.edu/registrar-student/Home.tap> for details. Information on the Bidding Points Double Auction procedure used at the Wharton School of the University of Pennsylvania can be found in Bartlett (2008), Guernsey (1999), Wharton (2007), and <http://technology.wharton.upenn.edu/auction/>.

³⁵More 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 (ii), a student whose bid for course j is amongst the q_j highest might not get it, while some other student who bids less does. Strategic issues aside, (ii) can lead to inefficient allocations. Sönmez and Ünver (2010) and Krishna and Ünver (2008) propose a mechanism that eliminates the inefficiencies that arise from this specific aspect of the Bidding Points Auction.

³⁶There are two other differences between the BPA and A-CEEI. First, A-CEEI has a more general preference reporting language. Second, the A-CEEI uses approximately equal incomes instead of exactly equal incomes.

³⁷Some of the institutions named above allow fake money to be carried over from one period (e.g., semester) to the next. If there were infinitely many such periods, then fake money would be like real money, because it always has a future use. Each of these institutions treats students in their final period as if they have quasi-linear preferences over courses and fake money. See the discussion following Proposition 11.

With fake money, student i 's correct demand, given a budget of b_i and a price vector of \mathbf{p} is

$$(8.1) \quad x_i^* = \arg \max_{x' \in 2^{\mathcal{C}}} (u_i(x') : \mathbf{p} \cdot x' \leq b_i).$$

Prices that clear the market according to (8.1) may not exist, and when they do a sophisticated fixed-point algorithm (or Walrasian auctioneer) is required to find them. A-CEEI finds prices that approximately clear the market according to (8.1) when students have approximately equal incomes; the algorithm that implements the mechanism is described in Section 9.2 and Othman et al (2010).

Suppose instead that students were bidding real money for courses, and that their preferences were quasi-linear over courses and real money. Then the correct demand for student i would be

$$(8.2) \quad x_i^* = \arg \max_{x' \in 2^{\mathcal{C}}} (u_i(x') - \mathbf{p} \cdot x').$$

This is exactly what students receive under the BPA: students receive every course for which their bid strictly exceeds price, and no courses for which their bid is strictly lower than price. Prices that clear the market according to (8.2) have the virtue of being simple to compute; no Walrasian auctioneer is necessary. But treating fake money as if it were real money creates incentives to misreport, even in large markets in which students regard prices as exogenous to their own report. And, it leads to highly unfair outcomes.

The incentives to misreport are easy to see. Suppose that there are three courses, $\mathcal{C} = \{a, b, c\}$, and that Alice's true additive-separable preferences, scaled to sum to her budget of 10000, are $u_{Alice} = (7000, 2000, 1000)$. Suppose that other students' bids are such that the q_j^{th} highest bids without Alice are $\mathbf{p}_{-Alice}^{BPM} = (8000, 3000, 1500)$. At these prices, if Alice bids truthfully, the BPA will allocate her *zero* courses. Her best reply is to misreport her preferences as $\hat{u}_{Alice} = (8001, 0, 1501)$ (or the equivalent). Note that \hat{u}_{Alice} misrepresents both cardinal and ordinal aspects of her preferences. Now, the BPA allocates her the bundle $\{a, c\}$; note that this is her correct competitive-equilibrium demand at $\mathbf{p}_{-Alice}^{BPM}$.

But what is undesirable about this? Alice simply tricked a demand function of the form (8.2) into behaving like one of the form (8.1). The answer is that it is impossible for all students to obtain their competitive-equilibrium demand simultaneously, even when an exact CEEI actually exists. In the example, Alice obtains a for 8001, but by doing so causes the unlucky student who bid 8000 of her points on a (say, Betty) no longer to get it. Betty now at best gets her correct demand given a budget of 2000. More generally, for the equilibrium price of course j to be p_j with strictly positive probability, there must be

at least one unlucky student who bids $p_j - 1$ with strictly positive probability. Otherwise, the students who get j for p_j would want to bid a little bit less.

Essentially, some students must be the marginal losers who support prices, and these students' outcomes may be highly unfair. The following results, proved in Appendix D, formalize that the BPA is not a backdoor method for implementing true competitive equilibria, and that the BPA can lead to highly unfair outcomes.

Proposition 10. *Suppose that (i) agents are able to describe their ordinal preferences over bundles using the message space provided by the BPA: that is, for each $i \in \mathcal{S}$, there exists a value vector $v_i = (v_{i1}, \dots, v_{iM}) \in \mathbb{R}_+^M$ such that $v_i \cdot x > v_i \cdot x' \iff u_i(x) > u_i(x')$; and (ii) an exact CEEI exists. Then*

- (1) *Truthful play of the BPA yields non-CEEI outcomes.*
- (2) *All complete information Nash equilibria of the BPA yield non-CEEI outcomes.*

Proposition 11. *Suppose that agents are able to describe their ordinal preferences over bundles using the message space provided by the BPA. Then*

- (1) *Truthful play of the BPA yields outcomes in which an agent's ex-post realized utility is zero.*
- (2) *All complete information Nash equilibria of the BPA yield outcomes in which an agent's ex-post realized utility is zero.*

Proposition 11 emerges 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 number of students allocated zero courses in the main round of bidding has been 53, 37, 64, and 17.³⁸ It is somewhat difficult to interpret these numbers, because in Chicago's BPA, budget that is unspent in one quarter carries over to future quarters. The cleanest evidence comes from focusing on students who get zero courses in their last term. Of the 17 students who got zero courses in Spring 2010, 5 were full-time MBA students in their last term. One student bears an uncanny resemblance to Alice: he bid 5466, 5000, 1500, and 1 for courses that had prices of 5741, 5104, 2023, and 721. Another case that is instructive is a student who bid 11354, 3, 3, 3, and 2 for courses that then had prices of 13266, 2023, 1502, 1300 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", he also did not get a small diamond or even any rocks.

A second implication of Proposition 11, and more broadly 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. Amongst full-time MBA students graduating in

³⁸Students allocated zero courses in the main round can fill their schedules in subsequent rounds of bidding, typically with courses that had a price of zero in the main round, i.e., courses that are unpopular and in excess supply.

Spring 2010, the median student graduated with a budget of 6,601 unspent points, which is nearly a full quarters' budget (8,000 points). The 90th percentile student graduated with 17,547 unspent points, and the 99th with 26,675 unspent points.

By contrast to Proposition 10, truthful play of the A-CEEI yields an exact CEEI outcome whenever one exists. And, Theorems 2 and 3 indicate that the highly unfair outcomes that occur for Alice when she plays truthfully, or for Betty when Alice plays strategically, never occur under A-CEEI.

9. 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 (2009) data from course allocation at Harvard Business School (HBS), and Othman, Budish and Sandholm's (2010) computational procedure for A-CEEI. Sections 9.1-9.2 describe the data and the computational procedure. Section 9.3 reports results on market-clearing error. Section 9.4 reports results on outcome fairness. Section 9.5 performs an ex-ante welfare comparison versus HBS's own draft mechanism.

9.1. Data and Key Assumptions. I use the Budish and Cantillon (2009) data on course allocation at Harvard Business School (HBS) for the 2005-2006 academic year. The data consist of students' true and stated ordinal preferences over 50 Fall semester courses and 47 Spring semester courses, as well as these courses' capacities. Data on true preferences are generally difficult to obtain for manipulable mechanisms. Budish and Cantillon (2009) use a survey conducted by HBS a few days' prior to the run of its mechanism as a proxy for true preferences, and support this assumption using equilibrium analysis and a follow-on survey. There are 916 HBS students; 456 filled out HBS's survey. In the analysis I consider an economy with just the 456, adjusting course capacities proportionally. Robustness checks reported in Budish and Cantillon (2009) suggest that there are no systematic differences in strategic preferences between the 456 who filled out the survey and the 460 who did not.

A-CEEI utilizes students' ordinal preferences over bundles of courses, but the HBS data consist of ordinal preferences over individual courses. To convert preferences over courses into preferences over bundles, I follow Budish and Cantillon (2009) and assume that students compare bundles based on the "average rank" of the courses in each bundle. For instance, a student prefers the bundle consisting of her 2nd and 3rd favorite courses to that consisting of her 1st and 5th 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 mechanism³⁹ students are unable to express the intensity of their preference for individual courses beyond ordinal rank. Exploration of preferences more complex than average-rank suggests that the reported results are robust. In particular, the performance of the HBS draft mechanism deteriorates relative to A-CEEI when there are complementarities or intense preferences, making the welfare difference found in Section 9.5 more pronounced.

I also assume that students report their preferences truthfully under A-CEEI. While I have no formal way of assessing whether truthful reporting is an equilibrium in this environment,⁴⁰ Section 7.3 suggests that this assumption is reasonable.

9.2. Approximate CEEI Algorithm. Theorem 1 is non-constructive and so computing A-CEEI prices is non-trivial. There are two computational challenges. The first is that calculating demands is NP-Complete: the problem of solving for an agent’s demand at a particular price vector is related to the knapsack problem. The complexity of solving for an agent’s demand grows with the number of bundles he must consider, which itself grows exponentially with the maximum number of courses per bundle. The second is that even if excess demand were easy to compute, finding an approximate zero of excess demand is a challenging search problem.

Othman, Budish and Sandholm (2010) develop a computational procedure that overcomes these two challenges in life-size problems.⁴¹ Agents’ demands are calculated using an integer program solver, CPLEX. The search procedure takes a traditional tâtonnement search process – which Scarf (1960) showed can cycle even in economies with divisible goods and convex preferences – and enhances it using an artificial-intelligence method called Tabu Search.⁴²

³⁹The HBS draft mechanism works as follows. Students report their ordinal preferences over individual courses. A computer assigns each student a random priority number. Then, over a series of rounds, it chooses courses for the students one at a time based on their reported preferences. In rounds 1, 3, 5, . . . it proceeds through students in ascending order of the random priority numbers, whereas in rounds 2, 4, 6, . . . it proceeds in descending order. At each turn, the choosing student is given his most-preferred course that (i) he has not yet received; (ii) is not yet at capacity.

⁴⁰There are 50 courses per semester, and each student ranks about 15 courses per semester. So there are about $\frac{50!}{(50-15)!} \approx 2 \times 10^{12}$ 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, unlike e.g. in Roth and Peranson (1999).

⁴¹The most recent version of the algorithm 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 twenty minutes, and 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 eleven hours. The paper reports results for semester-sized problems so that results can be reported for a large number of trials.

⁴²There are two basic ideas to the enhancement. First, the algorithm considers not only a tâtonnement adjustment of the form $\mathbf{p}^{t+1} = \mathbf{p}^t + \mathbf{z}(\mathbf{p}^t)$ but also adjustments that raise or lower just a single price at a time. This set of potential adjustments is called the “neighborhood” of \mathbf{p}^t . Second, of this neighborhood, the algorithm travels to the price vector that has the lowest market-clearing error, except that it avoids prices that have an excess demand vector that has been encountered recently, as recorded on the “Tabu List”. That is, the algorithm often travels in a seemingly less attractive direction, in an attempt to avoid cycles. Russell and Norvig (2002; Chapter 4) provide an overview of Tabu Search. The algorithm stops when it has (i) found a price vector with market clearing error within the Theorem 1 bound; and (ii) gone 100 iterations without further improvement. The algorithm does not explicitly calculate the full set of $(\frac{\sqrt{\sigma M}}{2}, \beta)$ -CEEI price vectors.

9.3. Market-Clearing Error. Theorem 1 indicates that there exist Approximate CEEI prices that clear the HBS course-allocation market to within market-clearing error of $\frac{\sqrt{2kM}}{2}$, where k is the number of courses per student, and M is the number of courses. Here, $k = 5$ and $M = 50, 47$ for the Fall and Spring semesters, respectively.

Figure 1 reports the actual market clearing error over 100 runs of A-CEEI on the Fall and Spring semesters of course allocation at HBS. Each run corresponds to a different set of random budgets.

[Insert Figure 1: Market-Clearing Error]

The actual error is meaningfully smaller than the bound implied by Theorem 1. The maximum observed error in Euclidean Distance is $\sqrt{14}$ in the Fall and $\sqrt{15}$ in the Spring, as compared with the Theorem 1 bounds of $\sqrt{125}$ and $\sqrt{117.5}$, respectively. Part of the explanation is that only a subset of courses ever have a strictly positive price: 21 in the Fall, and 22 in the Winter. If we reformulated the problem as one of allocating only the potentially scarce courses (see Section 3.3) this would reduce the bounds to $\sqrt{52.5}$ and $\sqrt{55}$, respectively.

In terms of seats, the maximum (mean) observed error is 14 (6.04) seats in the Fall and 11 (5.50) seats in the Spring. Budish and Cantillon (2009) find that HBS's own mechanism yields substantially more inefficiency ex-post. They find Pareto-improving trades involving 15% of course seats, and 85% of students.

9.4. 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. This can be seen by examining the average rank distribution shown in Figure 2, which will be described in Section 9.5 below. The worst outcome any student receives is an average rank of around 9, while students' maximin shares have an average rank of around 18-20.⁴³

Theorem 3 indicates that A-CEEI bounds each student's envy by a single good. Table 2 describes the distribution of the amount of realized envy in the HBS economy over 100 runs of A-CEEI, as measured in ranks.

[Insert Table 2: Degree of Ex-Post Envy]

⁴³In the Fall, the HBS economy has 50 courses and 29% excess capacity. A divider in divide-and choose with average-rank preferences will use all of the seats in her roughly 39 most-preferred Fall courses, and will not use any of the seats in courses she ranks below ≈ 39 . By dividing these seats equally she can guarantee that her least preferred schedule in the division has an average rank of around 20. The Spring has slightly fewer courses and slightly more excess capacity, so the same calculation yields 18.5 instead of 20. In the data, students only rank at most 30 courses overall (≈ 15 per semester) so it is not possible to calculate each student's precise maximin share.

Around 99% of students have no envy, that is, they weakly prefer their own allocation to any other student’s allocation. For the few students who do envy, the degree of envy is small. The worst observed case is envy of 5 course ranks, e.g., a student who receives her 2nd-6th favorite courses while someone else receives the student’s 1st-5th favorite courses.

9.5. Ex-Ante Efficiency Comparison Versus the HBS Draft Mechanism. Figure 2 reports the distribution of average ranks over 100 runs of A-CEEI and the HBS draft mechanism.

[Insert Figure 2: Ex-Ante Efficiency Comparison]

In each semester, the distribution of average ranks under A-CEEI first-order stochastically dominates that under HBS. First-order stochastic dominance is an especially strong comparison relation: we do not need to make any further assumptions on how utility responds to average rank to reach a welfare comparison.⁴⁴

There are two equivalent ways to interpret the f.o.s.d. finding. First, in this environment, utilitarian school administrators should prefer A-CEEI to the HBS draft mechanism. Second, a student who knows the distribution of outcomes but does not yet know his own preferences – i.e., a student behind a veil of ignorance in the sense of Harsanyi (1953) or Rawls (1971) – should prefer A-CEEI to the HBS draft mechanism.

The magnitude of the improvement is economically meaningful. The mean average ranks under A-CEEI are 4.24 in the Fall and 4.44 in the Spring, versus 4.72 and 4.76 for HBS. Thus on average the quality of a student’s schedule improves by 0.40 ranks per course. Simulations suggest that the outperformance of A-CEEI versus HBS grows when either some students have more extreme preferences (e.g., an especially high value for their single favorite course) or there are substitutabilities or complementarities amongst courses.

That A-CEEI outperforms the HBS draft mechanism on HBS’s own data makes sense, because A-CEEI avoids what Budish and Cantillon (2009) identify as the central weakness of the HBS draft mechanism while incorporating what they identify as its most attractive feature.

The central weakness of the HBS draft mechanism is that it is simple to manipulate even in large markets; this misreporting harms efficiency both in theory and in the data. If students reported their preferences truthfully under the HBS draft mechanism then, at least under the assumption of average rank preferences, HBS would actually look a bit

⁴⁴We have assumed that students’ ordinal preferences over bundles are based on the average rank of the courses contained in each bundle. We have not made any additional assumption about how their cardinal utilities depend on average rank. For instance, if x' has a lower average rank for student i than x'' , then we have assumed $u_i(x') > u_i(x'')$ but have not made any assumptions on the magnitude of this difference.

better than A-CEEI. The mean average ranks under truthful play of HBS would be 4.09 in the Fall and 4.40 in the Spring, versus 4.24 and 4.44 for A-CEEI.⁴⁵

The most attractive feature of the HBS draft mechanism is that it distributes choosing times as equally as possible amongst the students. The fairness benefits of this design feature are obvious. Budish and Cantillon (2009) show that there are important ex-ante efficiency benefits as well, and that the distribution of average ranks under HBS actually second-order stochastically dominates that under the Random Serial Dictatorship, despite the fact that the HBS draft mechanism is harmfully manipulated.⁴⁶ The approximately equal incomes of A-CEEI can be interpreted as analogous to the approximately equal choosing times of the HBS draft mechanism. By contrast, RSD can be interpreted either as a draft mechanism with maximally unequal choosing times, or as a competitive equilibrium mechanism with maximally unequal incomes (Proposition 9).

10. CONCLUSION

Most of what is known about the combinatorial assignment problem are impossibility theorems which 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 (Theorem 1), 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 (Theorems 2 and 3), and gives exact maximin shares and is 99% envy free in the data. Ideally, a mechanism would be strategyproof. A-CEEI is strategyproof in the large (Theorem 4), 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 than the Theorem 1 worst-case bound, consisting of just a single seat in on average 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

⁴⁵Here is the basic intuition for why HBS does better than A-CEEI on measures of average rank under truthful play. Consider a student i whose, say, four favorite courses are unpopular, but whose fifth course, j , is very popular. Under A-CEEI she is very likely to get j , since she spends zero on her top-four favorite courses. Under truthful play of the HBS draft mechanism she is very likely not to get j ; instead, her seat will go to someone who ranks it more highly than fifth.

⁴⁶Their theoretical explanation for why RSD is unattractive ex-ante is simple: lucky students who get early serial numbers in RSD might use their *last* choices to take courses that would have been some unlucky students' *first* choices. Note that the unattractiveness of RSD is independent of risk aversion: even risk-neutral students regard a "win a little, lose a lot" lottery as unattractive.

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 where budget constraints are non-zero but often bind, or where 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 $(\mathcal{S}, \mathcal{C}, (q_j)_{j=1}^M, (\Psi_i)_{i=1}^N, (u_i)_{i=1}^N)$, and fix $\beta > 0, \varepsilon > 0$. Let $\mathbf{b}' = (b'_1, \dots, b'_N)$ be any vector of budgets that satisfies $\max_i(b'_i) \leq 1 + \beta$ and $\min_i(b'_i) = 1$. In particular, \mathbf{b}' can be the target budgets specified by the market administrator.

Let $\bar{b} = 1 + \beta + \varepsilon$. Define an M -dimensional price space by $\mathcal{P} = [0, \bar{b}]^M$. For much of the proof we will work with an enlargement of this space $\tilde{\mathcal{P}} = [-1, \bar{b} + 1]^M$ in order to handle a boundary issue (described informally in 4.3.3).

Define a truncation function $t : \tilde{\mathcal{P}} \rightarrow \mathcal{P}$ that takes any price vector in $\tilde{\mathcal{P}}$ and truncates all prices to be within $[0, \bar{b}]$. Formally, for $\tilde{\mathbf{p}} \in \tilde{\mathcal{P}}$, $t(\tilde{\mathbf{p}}) = (\min[\bar{b}, \max(0, \tilde{p}_1)], \dots, \min[\bar{b}, \max(0, \tilde{p}_M)])$.

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 $\tilde{\mathbf{p}} \cdot x - \tau_{ix}$ (that is, a positive τ_{ix} decreases the price of x to i).

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

$$(A.1) \quad d_i(\tilde{\mathbf{p}}; b_i, \tau_i) = \arg \max_{x' \in 2^{\mathcal{C}}} (u_i(x') : \tilde{\mathbf{p}} \cdot x' \leq b_i + \tau_{ix'})$$

Let $\boldsymbol{\tau} \equiv (\tau_i)_{i \in \mathcal{S}}$. Excess demand $\mathbf{z}(\cdot)$ is defined by

$$(A.2) \quad \mathbf{z}(\tilde{\mathbf{p}}; \mathbf{b}, \boldsymbol{\tau}) = \left(\sum_{i=1}^N d_i(\tilde{\mathbf{p}}; b_i, \tau_i) \right) - \mathbf{q}.$$

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. Usually we are interested in how $d_i(\cdot)$ and $\mathbf{z}(\cdot)$ move with price.

Since each agent i consumes either 0 or 1 of each object j , it is without loss of generality to assume $q_j \in \{1, \dots, N\}$ and so $-N \leq z_j \leq N - 1$ for all $j \in \mathcal{C}$.

For each agent $i \in \mathcal{S}$ and schedule $x \in 2^{\mathcal{C}}$, define the budget-constraint-hyperplane $H(i, x)$ by $H(i, x) \equiv \{\tilde{\mathbf{p}} \in \tilde{\mathcal{P}} : \tilde{\mathbf{p}} \cdot x = b_i + \tau_{ix}\}$. Each budget-constraint hyperplane is of dimension $M - 1$.

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 $u_i(x) > u_i(x')$ (taxes favor more-preferred bundles)

- (iii) $\max_{i,x}(b'_i + \tau_{ix}^*) \leq \max_i(b'_i)$, $\min_{i,x}(b'_i + \tau_{ix}^*) \geq \min_i(b'_i)$; (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.

It will be important for obtaining the approximation bound that no two hyperplanes coincide (iv), and that no more than M of the hyperplanes intersect at any particular price vector (v). We could ensure the former by perturbing just the budgets, but to ensure (v) we need to perturb each budget-constraint hyperplane separately.

At the end of the proof, if agent i is actually assigned bundle x' in the Approximate CEEI we will adjust i 's budget to $b_i^* = b'_{i'} + \tau_{i'x'}^*$. Then we will set all of the taxes to zero. Property (ii) will ensure that x' is i 's most-preferred choice at a budget of b_i^* . Property (i) will ensure that $|b_i^* - b_i| < \varepsilon$, and property (iii) will ensure that \mathbf{b}^* preserves the inequality bound β .

Existence of a set of taxes $(\tau_{ix}^*)_{i \in \mathcal{S}, x \in 2^c}$ satisfying (i)-(v) follows from the fact that the number of agents, number of permissible schedules per agent, and the number of budget-constraint hyperplanes are all finite. Specifically, consider any set of taxes that satisfies (i) - (iv): these taxes define a perturbation of the set of b-c-h's. No two of the perturbed b-c-h's are homogeneous, i.e., have the same constant on the RHS, due to (iv). Generically, no more than M of a finite set of inhomogeneous hyperplanes intersect at a single point of an M -dimensional space. So there exist sets of taxes that satisfy (i)-(v).

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.

We define a tâtonnement price-adjustment function on the enlarged price space $\tilde{\mathcal{P}}$. Let $\gamma \in (0, \frac{1}{N})$ be a small positive constant. Given budgets \mathbf{b} and taxes $\boldsymbol{\tau}$ define $f : \tilde{\mathcal{P}} \rightarrow \tilde{\mathcal{P}}$ by:

$$(A.3) \quad f(\tilde{\mathbf{p}}) = t(\tilde{\mathbf{p}}) + \gamma \mathbf{z}(t(\tilde{\mathbf{p}}); \mathbf{b}, \boldsymbol{\tau})$$

The reason we impose $\gamma < \frac{1}{N}$ is to ensure 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 $i \in \mathcal{S}$. To see this, 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 $z_j(\mathbf{p}^*; \mathbf{b}', \boldsymbol{\tau}) \leq 0 - q_j < 0$ which contradicts $\tilde{p}_j^* \geq \bar{b}$ being part of a fixed point since $f_j(\tilde{\mathbf{p}}^*) = \bar{b} + \gamma z_j(\mathbf{p}^*; \mathbf{b}', \boldsymbol{\tau}) < \bar{b}$. Second, note that $p_j^* \in (0, \bar{b})$ implies $z_j(\mathbf{p}^*; \mathbf{b}', \boldsymbol{\tau}) = 0$. 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 $\boldsymbol{\tau}$, and $z(\mathbf{p}^*; \mathbf{b}', \boldsymbol{\tau}) = z(\mathbf{p}^*; \mathbf{b}^*, \mathbf{0})$. So $z_j(\mathbf{p}^*; \mathbf{b}^*, \mathbf{0}) \leq 0$ with $z_j(\mathbf{p}^*; \mathbf{b}^*, \mathbf{0}) < 0 \Rightarrow p_j^* = 0$, as required for competitive equilibrium.

Of course, f is not continuous, so there is no guarantee that such a fixed point will exist.

Step 3. Define an upper hemicontinuous set-valued correspondence F which is a “convexification” of f , and which 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:

$$(A.4) \quad F(\mathbf{p}) = \text{co}\{\mathbf{y} : \exists \text{ a sequence } \mathbf{p}^w \rightarrow \mathbf{p}, \mathbf{p}^w \neq \mathbf{p} \text{ such that } f(\mathbf{p}^w) \rightarrow \mathbf{y}\}$$

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 (A.4) is upper hemicontinuous, and hence has a fixed point (the other conditions for Kakutani's fixed point theorem – F is non-empty; $\tilde{\mathcal{P}}$ 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 b-c-h, 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 (A.4) 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 the analysis in Step 2 confirms that \mathbf{p}^* is an exact competitive equilibrium price vector. For $\mathbf{p}^* \neq \tilde{\mathbf{p}}^*$, that is, for cases where the fixed point lies in $\tilde{\mathcal{P}} \setminus \mathcal{P}$, we need the following simple lemma.

Lemma 2. For any $\tilde{\mathbf{p}} \in \tilde{\mathcal{P}} \setminus \mathcal{P}$: (i) $f(\tilde{\mathbf{p}}) = f(t(\tilde{\mathbf{p}}))$; (ii) $F(\tilde{\mathbf{p}}) \subseteq F(t(\tilde{\mathbf{p}}))$.

Proof. (i) Follows immediately from (A.3). (ii) Consider a \mathbf{y} for which there exists a sequence $\tilde{\mathbf{p}}^w \rightarrow \tilde{\mathbf{p}}, \tilde{\mathbf{p}}^w \neq \tilde{\mathbf{p}}$ such that $f(\tilde{\mathbf{p}}^w) \rightarrow \mathbf{y}$. Now consider the sequence $t(\tilde{\mathbf{p}}^w)$. By the continuity of $t(\cdot)$, this sequence converges to $t(\tilde{\mathbf{p}})$, and, using 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.⁴⁷ \square

Combining Lemma 22 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 the analysis of Step 2 gives that \mathbf{p}^* is an exact competitive equilibrium price vector.

Step 5. Suppose \mathbf{p}^* is on $L \geq 1$ budget-constraint hyperplanes. By Step 1 we know $L \leq M$. 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}^ϕ is close enough to \mathbf{p}^* that there is a path from \mathbf{p}^ϕ to \mathbf{p}^* that does not cross any budget-constraint hyperplane (until the moment it reaches \mathbf{p}^*).
- (ii) each \mathbf{p}^ϕ is on the “affordable” side of the l^{th} 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 l^{th} b-c-h, i_l , can weakly afford the bundle named on the l^{th} b-c-h, 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 $\phi = (\phi_1, \dots, \phi_L) \in \Phi$ an L -dimensional vector of 0’s and 1’s. The convex polytope $\pi(\phi) := \bigcap_{l=1}^L H_l^{\phi_l}$ 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 : $\mathcal{H} = \{H(i, x)_{i \in \mathcal{S}, x_i \in 2^c}\}$. 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}^*

⁴⁷Note that the reverse inclusion $F(t(\tilde{\mathbf{p}})) \subseteq F(\tilde{\mathbf{p}})$ need not be true. If $\tilde{p}_j < 0$ then every price $\tilde{\mathbf{p}}'$ near to $\tilde{\mathbf{p}}$ has $t_j(\tilde{\mathbf{p}}') = 0$, whereas some prices \mathbf{p}'' near to $t(\tilde{\mathbf{p}})$ will have $p_j'' > 0$. So it is possible for there to be sequences of prices that converge to $t(\tilde{\mathbf{p}})$ at which some agent’s choice set is different from that along any sequence converging to $\tilde{\mathbf{p}}$. For instance, ignoring taxes, suppose $b_i = 1, \tilde{p}_j = 1$ and $\tilde{p}_{j'} = -1$. There exist sequences converging to $t(\tilde{\mathbf{p}})$ along which i can afford $\{j\}$ but not $\{j, j'\}$. There exist no such sequences converging to $\tilde{\mathbf{p}}$ at which i can afford $\{j\}$ but not $\{j, j'\}$ at the truncated prices, because the truncated price of j' is always zero.

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 the a $\{\mathbf{p}^\phi\}_{\phi \in \Phi}$ satisfying the requirements above: each \mathbf{p}^ϕ is an arbitrary element of $\pi(\phi) \cap B_\delta(\mathbf{p}^*)$.⁴⁸

Step 6: For any $\mathbf{y} \in F(\mathbf{p}^*)$ there exist non-negative weights $\{\lambda^\phi\}_{\phi \in \Phi}$ with $\sum_{\phi \in \Phi} \lambda^\phi = 1$ such that $\mathbf{p}^* + \sum_{\phi \in \Phi} \lambda^\phi \gamma \mathbf{z}(\mathbf{p}^\phi) = \mathbf{y}$.

The idea of this step is: for any price \mathbf{p}' close enough to \mathbf{p}^* , excess demand at \mathbf{p}' is determined entirely by the combination of b-c-h half spaces to which it belongs.

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 , 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}^ϕ .

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

$$(A.5) \quad f(\mathbf{p}^{w,\phi}) \rightarrow \mathbf{p}^* + \gamma \mathbf{z}(\mathbf{p}^\phi)$$

Note too that any sequence $\mathbf{p}^w \rightarrow \mathbf{p}^*$ for which $f(\mathbf{p}^w)$ converges must converge to $\mathbf{p}^* + \mathbf{z}(\mathbf{p}^{\phi'})$ for some $\phi' \in \Phi$. This follows because $\bigcup_{\phi \in \Phi} \pi(\phi) \cap B_\delta(\mathbf{p}^*) = \tilde{\mathcal{P}} \cap B_\delta(\mathbf{p}^*)$.

Combining these facts, if $\mathbf{y} \in F(\mathbf{p}^*)$ then

$$(A.6) \quad \exists \{\lambda^\phi\}_{\phi \in \Phi} \text{ with } \sum_{\phi \in \Phi} \lambda^\phi = 1 \text{ and } \lambda^\phi \geq 0, \text{ all } \phi \in \Phi \text{ s.t.:}$$

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

Step 7. 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, $\boldsymbol{\zeta}$, lies in the convex hull of $\{\mathbf{z}(\mathbf{p}^\phi)\}_{\phi \in \Phi}$.

⁴⁸It is possible, if \mathbf{p}^* is on the boundary of P , that $\pi(\phi) \cap P = \emptyset$ for some combinations ϕ . (For instance, it is impossible to be below $x + y = 1$ and above $y = 1$ while $x, y \geq 0$). In that case \mathbf{p}^ϕ might have some components which are strictly negative. We have defined demand and excess demand to be well defined for such prices, but note that at the final step of the proof we will ensure that all prices are weakly positive.

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 $\tilde{\mathbf{p}}^* \in F(\mathbf{p}^*)$, and we can apply Step 6 to obtain:

$$(A.7) \quad \exists \{\lambda^\phi\}_{\phi \in \Phi} \text{ with } \sum_{\phi \in \Phi} \lambda^\phi = 1 \text{ and } \lambda^\phi \geq 0, \text{ all } \phi \in \Phi \text{ s.t.:$$

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

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

$$(A.8) \quad \sum_{\phi \in \Phi} \lambda^\phi \mathbf{z}(\mathbf{p}^\phi) = \frac{\tilde{\mathbf{p}}^* - \mathbf{p}^*}{\gamma}$$

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 $\sum_{\phi \in \Phi} \lambda^\phi \mathbf{z}(\mathbf{p}^\phi) \leq 0$ with $\sum_{\phi \in \Phi} \lambda^\phi z_j(\mathbf{p}^\phi) < 0 \Rightarrow p_j^* = 0$, as required for market clearing. That is, a convex combination of excess demands for prices near \mathbf{p}^* exactly clears the market at prices \mathbf{p}^* .

Consider the set of excess demands $\{\mathbf{z}(\mathbf{p}^\phi)\}_{\phi \in \Phi}$ and define $\boldsymbol{\zeta} \equiv \sum_{\phi \in \Phi} \lambda^\phi \mathbf{z}(\mathbf{p}^\phi)$. The vector $\boldsymbol{\zeta}$ is a ‘‘perfect market clearing’’ ideal at prices \mathbf{p}^* since $\boldsymbol{\zeta} \leq 0$ with $\zeta_j < 0 \Rightarrow p_j^* = 0$. Note that $\boldsymbol{\zeta} \in \text{co}\{\mathbf{z}(\mathbf{p}^\phi)\}_{\phi \in \Phi}$, where co denotes the convex hull.

Step 8. *The set of excess demands $\{\mathbf{z}(\mathbf{p}^\phi)\}_{\phi \in \Phi}$ has a special geometric structure. In particular, if the L hyperplanes correspond to L distinct agents then $\{\mathbf{z}(\mathbf{p}^\phi)\}_{\phi \in \Phi}$ are the vertices of a zonotope.*

The L intersecting budget-constraint hyperplanes name $L' \leq L$ distinct agents. Renumber the agents in \mathcal{S} so that agents $\{1, \dots, L'\}$ are the ones named on some b-c-h intersecting at \mathbf{p}^* . Denote by w_i the number of intersection b-c-h's that name agent $i \in \{1, \dots, L'\}$, and let $x_i^1, \dots, x_i^{w_i}$ denote the bundles pertaining to i 's b-c-h's, numbered so that $u_i(x_i^1) > \dots > u_i(x_i^{w_i})$. Note that $\sum_{i=1}^{L'} w_i = L$.

The following argument illustrates that agent $i \in \{1, \dots, L'\}$ purchases at most $w_i + 1$ distinct bundles at prices near to \mathbf{p}^* . In the halfspace $H^0(i, x_i^1)$ 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 $H(i, x_i^2), \dots, H(i, x_i^{w_i})$ 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 $H(i, x_i^3), \dots, H(i, x_i^{w_i})$ 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 $d_i^2, \dots, d_i^{w_i}$.

The process ends when we have crossed to the unaffordable side of all w_i of i 's budget-constraint hyperplanes, and so cannot afford any of $x_i^1, \dots, x_i^{w_i}$.

The demand of any agents other than the L' named on b-c-h's is unchanging near \mathbf{p}^* . Call the total demand of such agents $d_{S \setminus \{1, \dots, L'\}}(\mathbf{p}^*) = \sum_{i=L'+1}^N d_i(\mathbf{p}^*; b_i, \tau_i^*)$, and let $z_{S \setminus \{1, \dots, L'\}}(\mathbf{p}^*) = d_{S \setminus \{1, \dots, L'\}}(\mathbf{p}^*) - \mathbf{q}$.

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}^* :

$$(A.9) \quad \{\mathbf{z}(\mathbf{p}^\phi)\}_{\phi \in \Phi} = \left\{ z_{S \setminus \{1, \dots, L'\}}(\mathbf{p}^*) + \sum_{i=1}^{L'} \sum_{f=0}^{w_i} a_i^f d_i^f \right\}$$

subject to

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

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

At any price vector near to \mathbf{p}^* , each agent $i = 1, \dots, L'$ demands exactly one of their $w_i + 1$ demand bundles. Over the set $\Phi = \{0, 1\}^L$ every combination of the L' agents' demands is possible. Informally, it is possible to “walk through price space” near to \mathbf{p}^* in such a way that we cross just a single budget-constraint hyperplane (and hence change just a single agent's demand) at a time. This would not be possible if agents had identical budgets, because then their hyperplanes would coincide. (Also, we would not be able to guarantee that at most M intersect.)

Step 7 tells us that there exists a market-clearing excess demand vector in the convex hull of (A.9). This convex hull can be written as

$$(A.10) \quad \{\mathbf{z}(\mathbf{p}^\phi)\}_{\phi \in \Phi} = \left\{ z_{S \setminus \{1, \dots, L'\}}(\mathbf{p}^*) + \sum_{i=1}^{L'} \sum_{f=0}^{w_i} a_i^f d_i^f \right\}$$

subject to

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

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

The set (A.9) has a particularly interesting structure in case $L' = L$ (and so $w_i = 1$ for $i = 1, \dots, L'$). Define a vector $v_i = d_i^1 - d_i^0$. The vector v_i describes how i 's demand changes as we raise price from \mathbf{p}^* in a way that makes d_i^0 unaffordable. Observe that total excess demand at \mathbf{p}^* satisfies $\mathbf{z}(\mathbf{p}^*) = z_{-L'}(\mathbf{p}^*) + \sum_{i=1}^{L'} d_i^0$. The set $\{\mathbf{z}(\mathbf{p}^\phi)\}_{\phi \in \Phi}$ can

be rewritten as

$$(A.11) \quad \{\mathbf{z}(\mathbf{p}^\phi)\}_{\phi \in \Phi} = \{\mathbf{z}(\mathbf{p}^*) + \sum_{i=1}^L a_i v_i\}$$

subject to

$$a_i \in \{0, 1\} \text{ for all } i = 1, \dots, L$$

The set (A.11) gives the vertices of a geometrical object called a zonotope. The zonotope itself is the convex hull of (A.11). A zonotope is the Minkowski sum of a set of generating vectors; here, the generating vectors are v_1, \dots, v_L . If the vectors are linearly independent then the zonotope is a parallelotope, the multi-dimensional generalization of a parallelogram. (See Ziegler, 1995)

Step 9. *There exists a vertex of the geometric structure from Step 8, (A.9), that is within $\frac{\sqrt{\sigma M}}{2}$ distance of the perfect market clearing excess demand vector, ζ , found in Step 7. That is, for some $\mathbf{z}(\mathbf{p}^{\phi'}) \in \{\mathbf{z}(\mathbf{p}^\phi)\}_{\phi \in \Phi}$, $\|\mathbf{z}(\mathbf{p}^{\phi'}) - \zeta\|_2 \leq \frac{\sqrt{\sigma M}}{2}$.*

We are interested in bounding the distance between an element of (A.9) and an element of its convex hull (A.10), which we know contains ζ .⁴⁹

Fix an arbitrary interior point of (A.10). That is, fix a set of $a_i^f \in [0, 1]$ that satisfy the constraint $\sum_{f=0}^{w_i} a_i^f = 1$ for all $i = 1, \dots, L'$. For each $i = 1, \dots, L'$ define a random vector $\Theta_i = (\Theta_i^0, \dots, \Theta_i^{w_i})$ where the support of each Θ_i^f is $\{0, 1\}$, $E(\Theta_i^f) = a_i^f$ for all $f = 1, \dots, w_i$, and in any realization θ_i , $\sum_{f=0}^{w_i} \theta_i^f = 1$. Define the random matrix $\Theta = (\Theta_1, \dots, \Theta_{L'})$, and suppose that the Θ_i 's are independent. Let:

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

Linearity of expectations yields

$$(A.12) \quad \rho^2 = \sum_{i=1}^{L'} \mathbb{E}_{\Theta_i} \left(\left\| \sum_{f=0}^{w_i} [(a_i^f - \theta_i^f)] d_i^f \right\|_2 \right)^2$$

$$+ \sum_{j \neq i} \sum_{f=0}^{w_i} \sum_{g=0}^{w_j} \mathbb{E}_{\Theta_i^f, \Theta_j^g} [(a_i^f - \theta_i^f)(a_j^g - \theta_j^g)] (d_i^f \cdot d_j^g)$$

⁴⁹The 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 for this Step would be to use the Shapley Folkman theorem (Starr, 1969).

Independence yields

$$(A.13) \quad \begin{aligned} & \mathbb{E}_{\Theta_i^f, \Theta_j^g} [(a_i^f - \theta_i^f)(a_j^g - \theta_j^g)] \\ &= \mathbb{E}_{\Theta_i^f} [a_i^f - \theta_i^f] \mathbb{E}_{\Theta_j^g} [a_j^g - \theta_j^g] = 0 \end{aligned}$$

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

Lemma 3. For each $i = 1, \dots, L'$, $\mathbb{E}_{\Theta_i} \left(\left\| \sum_{f=0}^{w_i} [(a_i^f - \theta_i^f)] d_i^f \right\|_2 \right)^2 \leq \frac{\sigma w_i}{4}$

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 $\sigma = \min(2k, M)$ non-zero 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 $\bar{d}_i = \sum_{f=0}^{w_i} a_i^f d_i^f$. In words, \bar{d}_i is i 's average demand as used in the convex combination. Now rewrite

$$(A.14) \quad \begin{aligned} & \mathbb{E}_{\Theta_i} \left(\left\| \sum_{f=0}^{w_i} [(a_i^f - \theta_i^f)] d_i^f \right\|_2 \right)^2 \\ &= \sum_{f=0}^{w_i} a_i^f \left(\left\| \bar{d}_i - d_i^f \right\|_2 \right)^2 \end{aligned}$$

If $w_i = 1$ then (A.14) 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 $\frac{\sigma}{4}$ which is equal to the bound. If $w_i = 2$ then (A.14) 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 $\frac{\sigma}{3}$ which is strictly lower than the bound of $\frac{\sigma}{2}$. If $w_i = 3$ then (A.14) 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 $\frac{3\sigma}{8}$ which is strictly lower than the bound of $\frac{3\sigma}{4}$. 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 $\leq \sigma$, while the RHS of the bound $\frac{\sigma w_i}{4} \geq \sigma$. \square

Combining Lemma 3, (A.12), and (A.13) yields

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

This means that there must exist at least one realization of Θ 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 (A.10), in particular for weights \tilde{a}_i^f such that $z_{\mathcal{S} \setminus \{1, \dots, L'\}}(\mathbf{p}^*) + \sum_{i=1}^{L'} \sum_{f=0}^{w_i} \tilde{a}_i^f d_i^f = \zeta$. Call this realization $\tilde{\theta}$. This realization points us to an element of $\{\mathbf{z}(\mathbf{p}^\phi)\}_{\phi \in \Phi}$, namely $z_{\mathcal{S} \setminus \{1, \dots, L'\}}(\mathbf{p}^*) + \sum_{i=1}^{L'} \sum_{f=0}^{w_i} \tilde{\theta}_i^f d_i^f$, that is within $\frac{\sqrt{\sigma M}}{2}$ Euclidean distance of the perfect market clearing ideal point, ζ .

Step 10. *Use the vertex found in Step 9 to produce prices, budgets and an allocation that satisfy the statement of Theorem 1.*

In Step 9 we showed that the excess demand vector $\mathbf{z}(\mathbf{p}^{\phi'})$ approximately clears the market at prices \mathbf{p}^* . 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, and 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 9.

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 $b_i^* = b'_i + \tau_{ix_i^*}^*$. Observe that 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. The label ϕ' tells us which side $\mathbf{p}^{\phi'}$ is on of each of i 's hyperplanes $H(i, x^1), \dots, H(i, x^{w_i})$. If $\mathbf{p}^{\phi'} \in H^1(i, x^f)$, i.e., x^f is unaffordable for i at $\mathbf{p}^{\phi'}$, then 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. At the price vector \mathbf{p}^* and initial taxes $\boldsymbol{\tau}^*$ agent i could exactly afford bundle x^f , i.e., \mathbf{p}^* was on the budget-constraint hyperplane $H(i, x^f)$. This tiny perturbation ensures that at taxes $\boldsymbol{\tau}^{**}$ he can no longer afford x^f .⁵⁰ 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 $\mathbf{p}^{\phi'}$, i.e., $d_i(\mathbf{p}^*, b'_i, \tau_i^{**}) = d_i(\mathbf{p}^{\phi'}, b'_i, \tau_i^*)$.

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 requirement (ii) of Step 1 still obtains, we ensure that x_i^* remains optimal for i at prices \mathbf{p}^* and a budget of b_i^* (i.e., Condition (i) of Definition 1). 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 9. So budgets of \mathbf{b}^* , prices of \mathbf{p}^* , and the allocation \mathbf{x}^* satisfy all of the requirements of Theorem 1. \square

⁵⁰In Step 5 we indicated that if \mathbf{p}^* is on the boundary of \mathcal{P} then it is possible that some of the combinations of half spaces $\phi \in \Phi$ are entirely disjoint from \mathcal{P} . By perturbing budgets rather than prices and keeping prices at \mathbf{p}^* we avoid the worry that we end up with an illegal price vector.

APPENDIX B. PROOF OF THEOREM 2 AND RELATED RESULTS

Proof of Lemma 1. In Step 7 of the proof of Theorem 1 we showed that $\sum_{\phi \in \Phi} \lambda^\phi \mathbf{p}^* \cdot \mathbf{z}(\mathbf{p}^\phi) = 0$ for a set of prices $(\mathbf{p}^\phi)_{\phi \in \Phi}$ arbitrarily close to \mathbf{p}^* . Recall that $\varepsilon > 0$ is both the maximum τ_{ix} and the maximum discrepancy between each b'_i and b_i^* . At each price \mathbf{p}^ϕ , 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 $\sum_{\phi \in \Phi} \lambda^\phi \mathbf{p}^\phi \cdot \left(\sum_{i=1}^N d_i(\mathbf{p}^\phi, b'_i, \tau_i^*) \right) \leq \sum_{i=1}^N b_i^* + 2N\varepsilon$. Let $\varepsilon_2 > 0$ denote the maximum distance in the L_1 norm between \mathbf{p}^* and any of the \mathbf{p}^ϕ 's. Note that 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}^ϕ we have

$$\begin{aligned} & \sum_{\phi \in \Phi} \lambda^\phi \mathbf{p}^\phi \cdot \left(\sum_{i=1}^N d_i(\mathbf{p}^\phi, b'_i, \tau_i^*) \right) \\ & \geq \sum_{\phi \in \Phi} \lambda^\phi \mathbf{p}^* \cdot \left(\sum_{i=1}^N d_i(\mathbf{p}^\phi, b'_i, \tau_i^*) \right) - N\varepsilon_2 \\ & = \sum_{\phi \in \Phi} \lambda^\phi \mathbf{p}^* \cdot (\mathbf{z}(\mathbf{p}^\phi) + \mathbf{q}) - N\varepsilon_2 \\ & = \sum_{\phi \in \Phi} \lambda^\phi \mathbf{p}^* \cdot \mathbf{q} - N\varepsilon_2 \\ & = \mathbf{p}^* \cdot \mathbf{q} - N\varepsilon_2 \end{aligned}$$

So $\sum_{i=1}^N b_i^* + 2N\varepsilon \geq \mathbf{p}^* \cdot \mathbf{q} - N\varepsilon_2$. In the proof of Theorem 1 we are free to choose any $\varepsilon, \varepsilon_2 > 0$, after we have already fixed \mathbf{b}' . Choose them sufficiently small such that $N(2\varepsilon + \varepsilon_2) < \delta \sum_{i=1}^N (b'_i - \varepsilon)$. Since $b'_i - \varepsilon < b_i^*$ for all i , $\mathbf{p}^* \cdot \mathbf{q} \leq \sum_{i=1}^N b_i^* (1 + \delta)$, as required. \square

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

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

So set β sufficiently small that $N(1 + \beta)(1 + \delta) > (N + 1)$ is a contradiction, i.e., set $\beta < \frac{1 - N\delta}{N(1 + \delta)}$. \square

Proof of Proposition 8. Suppose $\alpha = 0$. Then condition (ii) of Definition 1 implies that $N(1 + \beta) \geq \mathbf{p}^* \cdot \mathbf{q}$. Let $\mathbf{x}^{MS'} = (x_1^{MS'}, \dots, x_N^{MS'}, x_{N+1}^{MS'})$ denote an $(N + 1)$ -maximin split for agent i , ordered such that $u_i(x_l^{MS'}) \geq u_i(x_{l+1}^{MS'})$ for $l = 1, \dots, N$. We know from Theorem 2 that we can at least guarantee $u_i(x_i^*) \geq u_i(x_{N+1}^{MS'})$. Suppose that i gets exactly his $N + 1$ -maximin share, i.e., $u_i(x_i^*) = u_i(x_{N+1}^{MS'})$. If $u_i(x_N^{MS'}) = u_i(x_{N+1}^{MS'})$ then Condition (1) of Proposition 8 is satisfied and we are done. Suppose $u_i(x_N^{MS'}) > u_i(x_{N+1}^{MS'})$. Condition (i) of Definition 1 implies

$$(B.1) \quad b_i^* < \mathbf{p}^* \cdot x_i^{MS'} \text{ for } l = 1, \dots, N$$

Since $\alpha = 0$, condition (ii) of Definition 1 implies that the other $N - 1$ agents can collectively afford the endowment but for $x_{N+1}^{MS'}$, i.e.,

$$(B.2) \quad \sum_{l=1}^N \mathbf{p}^* \cdot x_l^{MS'} \leq \sum_{l \neq i} b_l^*$$

The β inequality bound implies

$$(B.3) \quad \sum_{l \neq i} b_l^* \leq (N - 1)(1 + \beta)b_i^*$$

Combining (B.1), (B.2) and (B.3) gives $Nb_i^* \leq (N - 1)(1 + \beta)b_i^*$. So if $\beta < \frac{1}{N-1}$ we have a contradiction, and so $u_i(x_i^*) > u_i(x_{N+1}^{MS'})$, as required for Condition (2) of Proposition 8. \square

APPENDIX C. PROOF OF THEOREM 4

Before presenting the proof of Theorem 4, we first need to extend Definition 1 to accommodate continuum economies.

Definition 7 (Extension of Definition 1). *Fix an economy $(\mathcal{S}, \mathcal{C}, (q_j)_{j=1}^M, (\Psi_i)_{i=1}^N, (u_i)_{i=1}^N)$ and consider its continuum replication $(\mathcal{S}^\infty, \mathcal{C}, (q_j)_{j=1}^M, (\Psi_i)_{i \in \mathcal{S}^\infty}, (u_i)_{i \in \mathcal{S}^\infty})$. The allocation $\mathbf{x}^* = (x_i^*)_{i \in \mathcal{S}^\infty}$, budgets $\mathbf{b}^* = (b_i^*)_{i \in \mathcal{S}^\infty}$ and prices $\mathbf{p}^* = (p_1^*, \dots, p_M^*)$ constitute an **(α, β) -Approximate Competitive Equilibrium from Equal Incomes** $((\alpha, \beta)$ -CEEI) if:*

- (i) $x_i^* = \arg \max_{x' \in \Psi_i} [u_i(x') : \mathbf{p}^* \cdot x' \leq b_i^*]$ for all $i \in \mathcal{S}^\infty$
- (ii) $\| (z'_1(\mathbf{p}^*), \dots, z'_M(\mathbf{p}^*)) \| \leq \alpha$ where
 - (a) $z'_j(\mathbf{p}^*) = \int_{\mathcal{S}^\infty} x_{ij}^* di - q_j$ if $p_j^* > 0$
 - (b) $z'_j(\mathbf{p}^*) = \max(\int_{\mathcal{S}^\infty} x_{ij}^* di - q_j, 0)$ if $p_j^* = 0$
- (iii) $\inf_i (b_i^*) = 1 \leq \sup_i b_i^* \leq 1 + \beta$

Conditions (i) and (iii) are identical to those in the original Definition 1. Condition (ii) differs slightly, in that we now calculate $z'_j(\cdot)$ as a measure.

In finite economies, Theorem 1 indicates that as the number of agents and number of copies of each object grow large, market-clearing error goes to zero as a fraction of the endowment. Formally as $N, (q_j)_{j=1}^M \rightarrow \infty$, we have $\frac{\alpha}{\|q_1, \dots, q_M\|_2} \rightarrow 0$. In the continuum economy, the formal analogue of Theorem 1 is:

Lemma 4. *Fix an economy $(\mathcal{S}, \mathcal{C}, (q_j)_{j=1}^M, (\Psi_i)_{i=1}^N, (u_i)_{i=1}^N)$ and consider its continuum replication $(\mathcal{S}^\infty, \mathcal{C}, (q_j)_{j=1}^M, (\Psi_i)_{i \in \mathcal{S}^\infty}, (u_i)_{i \in \mathcal{S}^\infty})$. For any $\beta > 0$, there exists a $(0, \beta)$ -CEEI. Moreover, for any budget distribution $(b'_i)_{i \in \mathcal{S}^\infty}$ that satisfies $\sup_i(b'_i) \leq 1 + \beta$, $\inf_i(b'_i) = 1$, and is atomless, there exists a $(0, \beta)$ -CEEI with budgets of $(b_i^*)_{i \in \mathcal{S}^\infty}$ such that $b_i^* = b'_i$ for all $i \in \mathcal{S}^\infty$.*

Observe that Lemma 4 is slightly stronger than Theorem 1 in the sense that it is no longer necessary to be able to perturb budgets by $\varepsilon > 0$. All that is required is that the specified budget distribution is atomless.

Lemma 4 is essentially implied by results of Mas-Colell (1977) and Yamazaki (1978), who study exchange economies that satisfy analogues of the atomless budget distribution assumption in Lemma 4. For completeness I provide a self-contained proof.

Proof of Lemma 4. Fix a continuum replication economy $(\mathcal{S}^\infty, \mathcal{C}, (q_j)_{j=1}^M, (\Psi_i)_{i \in \mathcal{S}^\infty}, (u_i)_{i \in \mathcal{S}^\infty})$ and budget distribution $(b'_i)_{i \in \mathcal{S}^\infty}$ satisfying the requirements of the lemma statement. Define a price space $\mathcal{P} = [0, \bar{b}]^M$ with $\bar{b} \equiv 1 + \beta$. For $\mathbf{p} \in \mathcal{P}$, define individual and excess demand by

$$d_i(\mathbf{p}; b_i) = \arg \max_{x' \in 2^{\mathcal{C}}} (u_i(x') : \mathbf{p} \cdot x' \leq b_i)$$

$$z(\mathbf{p}; (b_i)_{i \in \mathcal{S}^\infty}) = \int_{\mathcal{S}^\infty} d_i(\mathbf{p}; b_i) di - \mathbf{q}$$

Step 1. *Excess demand is continuous in price.*

For $i \in \mathcal{S}^\infty$ and $x \in 2^{\mathcal{C}}$, let $H(i, x) = \{\mathbf{p} \in \mathcal{P} : \mathbf{p} \cdot x = b_i\}$ describe the hyperplane in price space along which i can exactly afford x . Since $2^{\mathcal{C}}$ is finite, each agent is associated with a finite number of such budget-constraint hyperplanes. As described in the proof of Theorem 1, each agent's demand is continuous in price except possibly on one of these finitely many b-c-h's.

Since the distribution of budgets $(b'_i)_{i \in \mathcal{S}^\infty}$ is atomless, at any $\mathbf{p} \in \mathcal{P}$, for any bundle $x \in 2^{\mathcal{C}}$ the measure of the set $\{i \in \mathcal{S}^\infty : \mathbf{p} \cdot x = b'_i\}$ is zero. Since $2^{\mathcal{C}}$ is finite, the measure of the set $\{i \in \mathcal{S}^\infty : \exists x \in 2^{\mathcal{C}} : \mathbf{p} \cdot x = b'_i\}$ is also zero. Hence, at any point in price space,

the measure of the set of agents whose demand is possibly discontinuous at this price is zero. The magnitude of a discontinuity in any one agent's demand is bounded above (by $\sqrt{\sigma}$). Hence $z(\mathbf{p}; (b'_i)_{i \in \mathcal{S}^\infty})$ is continuous in \mathbf{p} .

Step 2. Define a tâtonnement price-adjustment function f on \mathcal{P} . Show that it has a Brouwer fixed point $\mathbf{p}^* = f(\mathbf{p}^*)$.

Define $f : \mathcal{P} \rightarrow \mathcal{P}$ by:

$$(C.1) \quad f_j(\mathbf{p}) = \min [0, \max (\bar{b}, p_j + z_j(\mathbf{p}; (b'_i)_{i \in \mathcal{S}}))]$$

for $j = 1, \dots, M$. In words, f adjusts each object's price in the direction of its excess demand at \mathbf{p} , truncating above and below to ensure the image of f lies within \mathcal{P} . Since $z(\cdot)$ is continuous in \mathbf{p} , so is $f(\cdot)$. Since \mathcal{P} is compact and convex we can apply Brouwer's fixed point theorem: there exists $\mathbf{p}^* = f(\mathbf{p}^*)$.

Step 3. Show that prices \mathbf{p}^* , budgets $(b'_i)_{i \in \mathcal{S}^\infty}$ and their associated allocation $(x_i^*)_{i \in \mathcal{S}^\infty}$ constitute a $(0, \beta)$ -CEEI.

We need to show that $z_j(\mathbf{p}^*; (b'_i)_{i \in \mathcal{S}^\infty}) \leq 0$ with equality whenever $p_j^* > 0$, for all $j = 1, \dots, M$. There are three cases. If $p_j^* = 0$, then $\mathbf{p}^* = f(\mathbf{p}^*)$ and (C.1) together imply $z_j(\mathbf{p}^*; (b'_i)_{i \in \mathcal{S}}) \leq 0$, as required. If $p_j^* \in (0, \bar{b})$, then $\mathbf{p}^* = f(\mathbf{p}^*)$ and (C.1) together imply $z_j(\mathbf{p}^*; (b'_i)_{i \in \mathcal{S}}) = 0$, as required. Finally, if $p_j^* = \bar{b}$, then $\mathbf{p}^* = f(\mathbf{p}^*)$ is impossible, because the measure of the set of agents who can afford j is zero and its quantity is strictly positive, and so $z_j(\mathbf{p}^*; (b'_i)_{i \in \mathcal{S}}) < 0$. So the price vector \mathbf{p}^* , budget distribution of $b_i^* = b'_i$ for all $i \in \mathcal{S}$, and allocation $x_i^* = d_i(\mathbf{p}^*; b'_i)$ for all $i \in \mathcal{S}$, satisfy Definition 7's requirements for a $(0, \beta)$ -CEEI. □

Proof of Theorem 4. Fix an economy $(\mathcal{S}, \mathcal{C}, (q_j)_{j=1}^M, (\Psi_i)_{i=1}^N, (u_i)_{i=1}^N)$ and consider its continuum replication $(\mathcal{S}^\infty, \mathcal{C}, (q_j)_{j=1}^M, (\Psi_i)_{i \in \mathcal{S}^\infty}, (u_i)_{i \in \mathcal{S}^\infty})$. Consider an arbitrary agent $i \in \mathcal{S}^\infty$, and fix arbitrary reports by all agents other than i , denoted $(\hat{u}_{i'})_{i' \in \mathcal{S}^\infty \setminus \{i\}}$.

Since i is zero measure, his report cannot affect the measure of excess demand at any particular price vector, and so he cannot affect whether any particular price vector is part of a $(0, 0)$ -CEEI, or whether or not the mechanism converges at Step (2) of A-CEEI. Suppose in fact the set of $(0, 0)$ -CEEIs is non-empty given $(\hat{u}_{i'})_{i' \in \mathcal{S}^\infty \setminus \{i\}}$. Since preferences are strict, randomly selecting a $(0, 0)$ -CEEI allocation is isomorphic to randomly selecting a $(0, 0)$ -CEEI price vector. At any such price vector, condition (i) of Definition 7 indicates i maximizes his utility by reporting truthfully.

Suppose the set of $(0,0)$ -CEEIs is empty given $(\hat{u}_{i'})_{i' \in \mathcal{S}^\infty \setminus \{i\}}$. In Step (3) of A-CEEI, i 's target budget b'_i is chosen randomly. Given Lemma 4 and the tie-breaking rule in Step (4), A-CEEI will restrict attention to the set of $(0,\beta)$ -CEEIs with $b_i^* = b'_i$ for all $i \in \mathcal{S}^\infty$. Since he is zero measure, i cannot affect which price vectors are part of this set. As above, since preferences are strict and budgets are fixed, randomly selecting a $(0,\beta)$ -CEEI allocation is isomorphic to randomly selecting a $(0,\beta)$ -CEEI price vector. At any such price vector, given his exogenous budget b_i^* , condition (i) of Definition 7 indicates i maximizes his utility by reporting truthfully. \square

APPENDIX D. OTHER OMITTED PROOFS

Proof of Proposition 1. For the case $M = 4$, consider the following example:

Example 1. *There are 4 objects, $\mathcal{C} = \{a, b, c, d\}$, each with capacity 2. There are 4 agents, $\mathcal{S} = \{i_1, i_2, i_3, i_4\}$, whose preferences are $u_{i_1} : \{a, b, c\}, \{d\}, \dots$, $u_{i_2} : \{a, b, d\}, \{c\}, \dots$, $u_{i_3} : \{a, c, d\}, \{b\}, \dots$, and $u_{i_4} : \{b, c, d\}, \{a\}, \dots$.*

Fix $\beta \gtrsim 0$, 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 (A.4) is \mathbf{p}^* as given by $p_a^* = \left(\frac{1+\beta_1+\beta_2+\beta_3-2\beta_4}{3}\right)$, $p_b^* = \left(\frac{1+\beta_1+\beta_2-2\beta_3+\beta_4}{3}\right)$, $p_c^* = \left(\frac{1+\beta_1-2\beta_2+\beta_3+\beta_4}{3}\right)$, $p_d^* = \left(\frac{1-2\beta_1+\beta_2+\beta_3+\beta_4}{3}\right)$. 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 $\frac{\sqrt{\sigma M}}{2} = 2$ from the perfect market clearing demand of $\mathbf{q} = (2, 2, 2, 2)$. To see why, consider the matrix that is formed by stacking the four agents' change-in-demand vectors at \mathbf{p}^* (see (A.11)):*

$$(D.1) \quad \begin{pmatrix} -1 & -1 & -1 & +1 \\ -1 & -1 & +1 & -1 \\ -1 & +1 & -1 & -1 \\ +1 & -1 & -1 & -1 \end{pmatrix}$$

This is an example of a Hadamard matrix: all of its entries are ± 1 and its rows are mutually orthogonal (Wallis et al, 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, $\sigma = M$, so $\frac{\sqrt{\sigma M}}{2} = 2$). Since \mathbf{q} is the hypercube's center, as here, each vertex is distance $\frac{\sqrt{\sigma M}}{2}$ from perfect market clearing.*

The Hadamard matrix (D.1) 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 4.⁵¹ We can use these regular Hadamard matrices to construct examples that are exactly analogous to Example 1 for $M = 16, 64, 256, \dots$ \square

Proof of Proposition 2. Suppose \mathbf{x}' Pareto improves upon \mathbf{x}^* in economy $(\mathcal{S}, \mathcal{C}, (q_j^*)_{j=1}^M, (\Psi_i)_{i=1}^N, (u_i)_{i=1}^N)$. 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 non-negative and \mathbf{x}^* allocates all units of positive-priced goods. \square

Proof of Proposition 3. Suppose there exists a feasible allocation \mathbf{x}' such that $u_i(x'_j) \geq u_i(\frac{\mathbf{q}}{n})$ for all $x'_j \in \mathbf{x}'$, with at least one strict. Now consider an economy with the same endowment but in which all N agents have i 's preferences. In this economy, the allocation \mathbf{x}' Pareto improves upon the allocation in which all agents receive $\frac{\mathbf{q}}{N}$, which is a contradiction when preferences are convex and monotonic (see e.g. Moulin 1990). Hence for any feasible allocation \mathbf{x}' , $u_i(\frac{\mathbf{q}}{n}) \geq \min(u_i(x'_1), \dots, u_i(x'_N))$. This in turn implies that $(\frac{\mathbf{q}}{n}, \dots, \frac{\mathbf{q}}{n})$ is a maximin split, and that $\frac{\mathbf{q}}{N}$ is a maximin share. If preferences are strictly convex then $(\frac{\mathbf{q}}{n}, \dots, \frac{\mathbf{q}}{n})$ is the unique maximin split, and so $\frac{\mathbf{q}}{N}$ is the unique maximin share. \square

Proof of Proposition 5. Follows immediately from Definition 3. \square

Proof of Proposition 9. (Single-unit case) A (0,0)-CEEI exists if and only if all agents have a different favorite object. In this case, RSD and the Approximate CEEI mechanism clearly coincide. Suppose a (0,0)-CEEI does not exist. Then the strict budget order selected in Step (3) of A-CEEI plays the same role as the serial order selected in RSD. Specifically, the RSD allocation can be supported as a $(0, \beta)$ -CEEI by the price vector in which an object's price is equal to the budget of the last agent to obtain a copy of it, or zero if it does not reach capacity. Any other allocation has strictly-positive market-clearing error at any price vector: if the agent with the l^{th} highest budget obtains an allocation at some price vector \mathbf{p}' that is strictly better than he receives under the serial dictatorship then one of the first l objects selected in the serial dictatorship must be over-allocated at \mathbf{p}' , and vice versa. So the RSD allocation is selected in Step (5) of A-CEEI.

(General case) Run a serial dictatorship with the same serial order as the randomly selected budget order. If an object reaches capacity at the l^{th} agent's turn, set its price equal to $\frac{1}{k}(k+1)^{N-l}$. If an object never reaches capacity, set its price equal to zero. At

⁵¹Here is Neil Sloane's proof. Let A be the matrix defined in (D.1). The tensor product of two Hadamard matrices is itself a Hadamard matrix, and the tensor product preserves the "same number of +1s per row" property. So $A \otimes A$ is a 16-Dimensional Hadamard matrix with the same number of +1s per row, $A \otimes (A \otimes A)$ is a 64-Dimensional example, etc. QED. It has been conjectured that there exist regular Hadamard matrices of order $(2n)^2$ for any integer n . Useful references are <http://www.research.att.com/~njas/hadamard/> and <http://www.research.att.com/~njas/sequences/A016742>.

this price vector, each agent consumes her most preferred bundle that consists of objects available at her turn in the serial dictatorship, and market-clearing error is zero. \square

Proof of Proposition 10. We prove both parts using the following example. There are two agents, $\mathcal{S} = \{i_1, i_2\}$, and three objects $\mathcal{C} = \{a, b, c\}$, each in unit supply. Agents' additive-separable preferences over objects are given by the following table.

	a	b	c
i_1	600	350	50
i_2	400	300	300

The allocation $(\{a\}, \{b, c\})$ can be supported as a Competitive Equilibrium from Equal Incomes. Since this is the only allocation that is Pareto efficient and envy free it must be the unique CEEI.

Part (1). If agents report truthfully, we reach the allocation $(\{a, b\}, \{c\})$ which is not a CEEI (i_2 envies i_1).

Part (2). Suppose there is a complete information Nash equilibrium of the BPA that yields the allocation $(\{a\}, \{b, c\})$ with probability one. We will show that it is a contradiction for i_1 's equilibrium best response to yield him the allocation $\{a\}$ with probability one.

Let y denote the common budget, and let \underline{v}_j and \bar{v}_j denote the lowest and highest bid i_2 ever submits for item j in the supposed equilibrium. Since i_2 's strategy always obtains b and c , we know he must always bid a strictly positive amount for each item; since bids are in integer amounts, $\underline{v}_b \geq 1$ and $\underline{v}_c \geq 1$. From the fact that bids always sum to y we have that $\bar{v}_a \leq y - 2$, and that $\min(\underline{v}_b, \underline{v}_c) \leq \lfloor \frac{y - \bar{v}_a}{2} \rfloor$.

Consider the following reply by i_1 . Bid $\bar{v}_a + 1$ for a . If $\underline{v}_b \leq \underline{v}_c$ then bid $y - (\bar{v}_a + 1)$ for b and zero for c . Otherwise if $\underline{v}_b > \underline{v}_c$ then bid $y - (\bar{v}_a + 1)$ for c and zero for b . This reply always obtains a , and obtains whichever of b or c he bids a positive amount for with strictly positive probability.⁵² Since this reply is possible, it cannot be that his best reply yields the allocation $\{a\}$ with probability one. A contradiction. \square

Proof of Proposition 11. We prove both parts using examples in which there are three agents, $\mathcal{S} = \{i_1, i_2, i_3\}$ three objects, $\mathcal{C} = \{a, b, c\}$, each in unit supply, $\Psi_i = 2^{\mathcal{C}}$ for all $i \in \mathcal{S}$, and agents' vNM utilities are additive separable over objects.

Part (1). Consider the following preferences:

⁵²If agents are allowed to submit real-valued bids, we have to allow for i_2 's bid for a to mix all the way up to y (all we know is that there is no atom at y). A slightly more involved argument goes through, in which i_1 obtains a with probability $1 - \varepsilon$ and mixes his bids for b and c to obtain expected value that more than compensates for the ε chance of losing a .

	a	b	c
i_1	334	333	333
i_2	600	400	0
i_3	0	600	400

If agents report truthfully, i_1 gets zero objects.

Part (2). Consider the following preferences:

	a	b	c
i_1	600	200	200
i_2	600	200	200
i_3	0	500	500

One complete information Nash equilibrium is for i_1 and i_2 each to bid $(1000, 0, 0)$ and for i_3 to bid $(0, 500, 500)$. In this equilibrium, whichever of i_1 or i_2 loses the coin toss for object a ends up with zero objects ex-post.

Suppose there exists some other Nash equilibrium in which each agent always gets exactly one object. If i_1 and i_2 each bid $(1000, 0, 0)$ with positive probability then there will be outcomes in which one of them gets zero objects (at best they tie i_3 for either b or c). Suppose i_1 never bids $(1000, 0, 0)$. Then i_2 can obtain a with probability one, and since we are supposing that each agent always gets exactly one object, it must be that i_2 always obtains exactly $\{a\}$. So i_2 always obtains exactly $\{b\}$ or $\{c\}$; given his utilities, he prefers to deviate and bid $(1000, 0, 0)$, a contradiction.

□

APPENDIX E. TWO VERSIONS OF A-CEEI WITH EXACT MARKET CLEARING

A-CEEI clears the market with a small amount of error. In the context of course allocation, a small amount of error might not be very costly in practice, for the envelope-theorem and secondary-market reasons discussed in Section 3.3. In other contexts, such as assigning pilots to planes, market-clearing error may be extremely costly. This appendix describes two variants of A-CEEI for such contexts.

The first variant, *A-CEEI with a Pareto-Improving Secondary Market*, has two components. First, we run A-CEEI almost exactly as above, but with the stipulation that excess demand is not allowed at the allocations in Step 4, only excess supply. It is an open question whether the no-excess-demand restriction necessitates a larger worst-case error bound than $\frac{\sqrt{\sigma M}}{2}$. It also may be necessary to consider price vectors that lie outside of $\mathcal{P}(\delta, \mathbf{b}')$, i.e., at which the goods endowment costs more than a small amount more than the budget endowment. In this case the initial stage will not guarantee $N + 1$ -Maximin Shares as in Theorem 2, but will provide a slightly weaker guarantee.

Second, we iteratively execute Pareto-improving trades – in particular, utilizing the excess capacity held by the market administrator – until we reach a Pareto efficient allocation. The Pareto-improvement stage could be implemented by a computer algorithm using the preference reports from Step 1 of A-CEEI, or it can be interpreted as a metaphor for an add-drop procedure similar to those in current use.

The benefit of the Pareto-improvement stage is that it yields exact ex-post efficiency. There are two costs. First, it may introduce additional envy. Second, the mechanism no longer is strategyproof in the large: executing a utility-improving trade for an agent corresponds to increasing that agent’s budget in a way that depends directly on his report.

The second variant is the *Competitive Equilibrium from Equal-as-Possible Incomes*. This mechanism simply seeks an (α, β) -CEEI with $\alpha = 0$ and β as small as possible. Specifically, if there is no exact CEEI, choose some large finite sequence of budget vectors $\{\mathbf{b}^t\}_{t=1}^T$ for which budget inequality is increasing in t . In particular, let $\mathbf{b}^T = (1, (k+1), (k+1)^2, \dots, (k+1)^{N-1})$, with k defined as in Theorem 1.⁵³ Then, for periods $t = 1, \dots, T$, search in a random order over the $N!$ permutations of \mathbf{b}^t for an exact competitive equilibrium price vector. As soon as one is found, stop.

Proposition 9 shows that this procedure will certainly converge at period T . Toy problems and some limited experimentation on the HBS data suggest that it typically will converge much sooner, though there are also toy examples in which dictatorship levels of budget inequality are necessary to achieve exact market clearing (e.g., the economy described in footnote 24).

The benefits of this procedure are that it is Pareto efficient and SPITL. The primary disadvantage is that it is not possible to provide fairness guarantees. A second disadvantage is computability. As stated, it requires $O(TN!)$ searches for a $(0, \beta)$ -CEEI.

⁵³For instance, one could use the sequence $\left\{ \mathbf{b}^t = \left(1, (k+1)^{\frac{t}{T}}, (k+1)^{\frac{2t}{T}}, \dots, (k+1)^{\frac{(N-1)t}{T}} \right) \right\}_{t=1}^T$ for a large integer T . Budget inequality can be defined as the ratio of the maximum to minimum budgets, or it can be a more complicated statistic such as the Gini coefficient.

REFERENCES

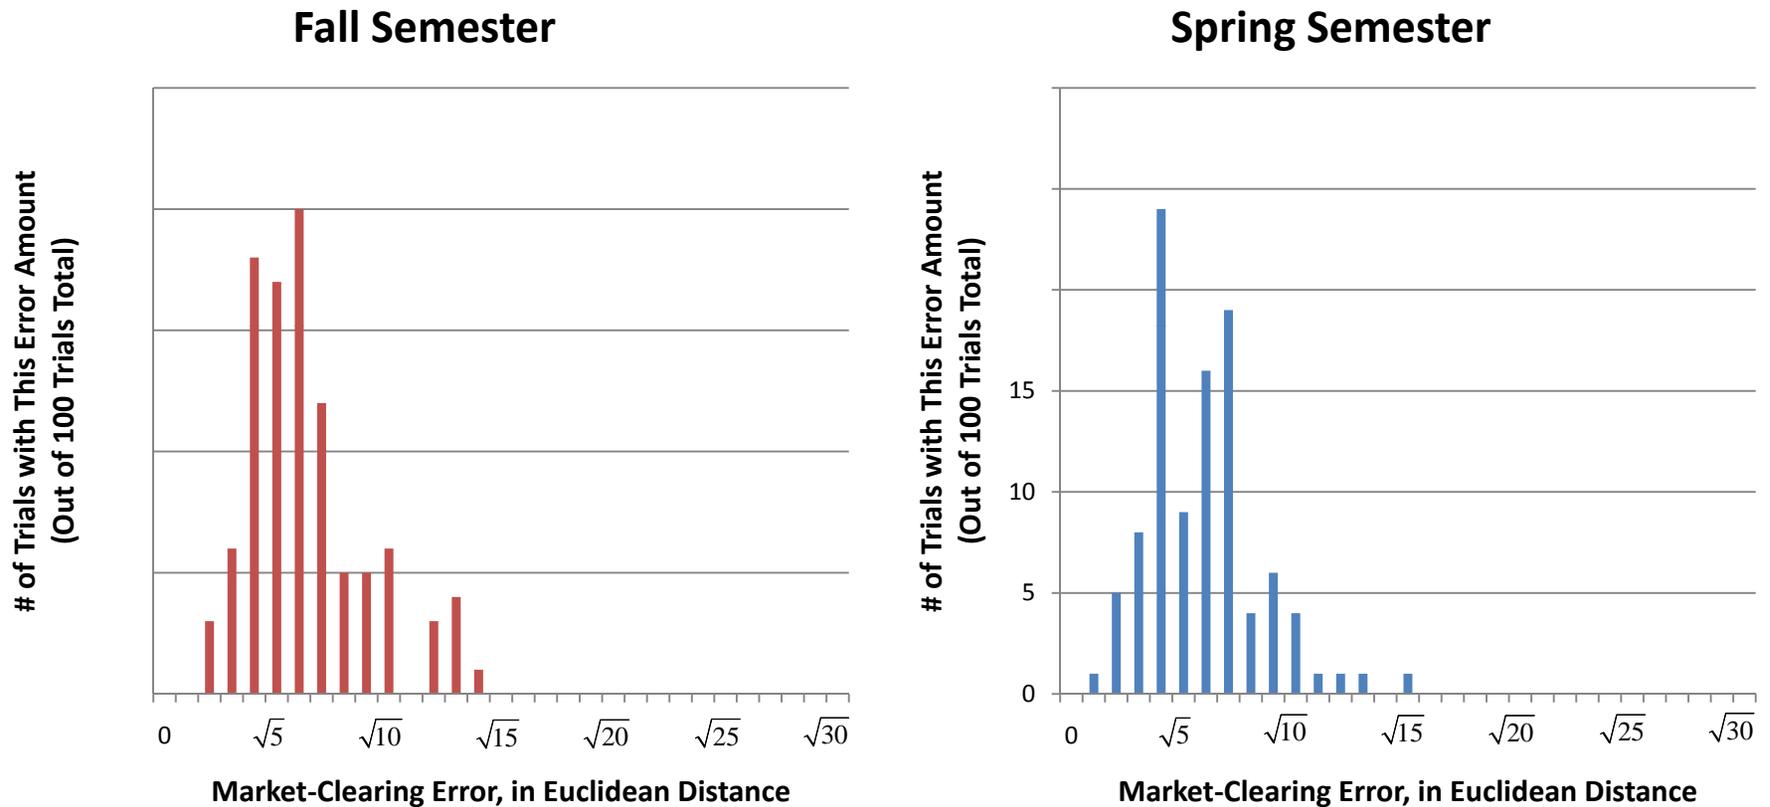
- [1] Abdulkadiroğlu, Atila, Yeon-Koo Che and Yosuke Yasuda (2008). "Expanding 'Choice' in School Choice." Mimeo.
- [2] Abdulkadiroğlu, Atila, Parag Pathak, and Alvin E. Roth (2005a). "The Boston Public School Match" *American Economic Review, Papers and Proceedings*, 95(2), 368-71.
- [3] Abdulkadiroğlu, Atila, Parag Pathak, Alvin E. Roth, and Tayfun Sönmez (2005b). "The New York City High School Match" *American Economic Review, Papers and Proceedings*, 95(2), 364-67.
- [4] Abdulkadiroğlu, Atila, Parag Pathak, and Alvin E. Roth (2009). "Strategyproofness versus Efficiency in Matching with Indifferences: Redesigning the NYC High School Match." *American Economic Review*, 99(5), 1954-78.
- [5] Abdulkadiroğlu, Atila and Tayfun Sönmez (1998). "Random Serial Dictatorship and the Core from Random Endowments in House Allocation Problems." *Econometrica*, 66(3), 689-702.
- [6] Abdulkadiroğlu, Atila and Tayfun Sönmez (1999). "House Allocation with Existing Tenants." *Journal of Economic Theory*, 88(2), 233-260.
- [7] Abdulkadiroğlu, Atila and Tayfun Sönmez (2003). "School Choice: A Mechanism Design Approach." *American Economic Review*, 93(3), 729-47.
- [8] Adler, Barry, Thomas Delaney, Michelle Kirkland, Liam Murphy (2008). "Memorandum re: Changes to the Registration System." April 25. Accessed at http://www.law.nyu.edu/ecm_dlv3/groups/public/@nyu_law_website__academics/documents/web_copytext/ecm_pro_061262.pdf
- [9] Albergotti, Reed (2010). "Why the NFL Draft Drives Economists Crazy." *Wall Street Journal*, April 22.
- [10] Alkan, Ahmet, Gabrielle Demange and David Gale (1991). "Fair Allocation of Indivisible Goods and Criteria of Justice." *Econometrica*, 59(4), 1023-1039.
- [11] Alon, Noga and Joel Spencer (2000). *The Probabilistic Method*. Wiley. Second Edition.
- [12] Anderson, Robert M., M. A., Khan and S. Rashid (1982). "Approximate Equilibria with Bounds Independent of Preferences." *Review of Economic Studies*, 44, 473-5.
- [13] Arnsperger, Christian (1994). "Envy-Freeness and Distributive Justice." *Journal of Economic Surveys*, 8(2), 155-86.
- [14] Arrow, Kenneth and F. H. Hahn (1971). *General Competitive Analysis*. Holden Day.
- [15] Aumann, Robert J., 1964. "Existence of Competitive Equilibria in Markets with a Continuum of Traders." *Econometrica*, 32: 39-50.
- [16] Bartlett, Thomas (2008). "Class Warfare: When Getting in is the Hardest Part." *Chronicle of Higher Education* 54(23), Feb 15, Page A1.
- [17] Bezakova, Ivona and Varsha Dani (2005). "Allocating Indivisible Goods." *ACM SIGecom Exchanges*, 5(3), 11-18.
- [18] Bogomolnaia, Anna and Hervé Moulin (2001). "A New Solution to the Random Assignment Problem." *Journal of Economic Theory*, 100, 295-328.
- [19] Brainard, William C. and Herbert E. Scarf (2005). "How to Compute Equilibrium Prices in 1891." *The American Journal of Economics and Sociology*, 64(1), 57-83.
- [20] Brams, Steven J and D. Marc Kilgour (2001). "Competitive Fair Division." *The Journal of Political Economy*, 109(2), 418-443.
- [21] Brams, Steven J and Philip Straffin (1979). "Prisoners' Dilemma and Professional Sports Drafts." *American Mathematical Monthly*, 86, 80-88.
- [22] Brams, Steven J and Alan D. Taylor (1996). *Fair Division: From Cake Cutting to Dispute Resolution*. Cambridge University Press
- [23] Brams, Steven J and Alan D. Taylor (1999). "The Win-Win Solution: Guaranteeing Fair Shares to Everybody." W.W. Norton and Company.
- [24] Budish, Eric and Estelle Cantillon (2009). "The Multi-Unit Assignment Problem: Theory and Evidence from Course Allocation at Harvard." Mimeo.
- [25] Budish, Eric, Yeon-Koo Che, Fuhito Kojima, and Paul Milgrom (2010). "Implementing Random Assignments: A Generalization of the Birkhoff-von Neumann Theorem." Mimeo.
- [26] Che, Yeon-Koo and Fuhito Kojima (2010). "Asymptotic Equivalence of Probabilistic Serial and Random Priority Mechanisms." *Econometrica*, forthcoming.
- [27] Chen, Yan and Tayfun Sönmez (2002). "Improving Efficiency of On-Campus Housing: An Experimental Study." *American Economic Review*, 92(5), 1669-86.

- [28] Cramton, Peter, Yoav Shoham and Richard Steinberg eds. (2006). *Combinatorial Auctions*. MIT Press.
- [29] Crawford, Vincent P. (1977). "A Game of Fair Division." *The Review of Economic Studies*, 44(2), 235-47.
- [30] Cromme, Ludwig J. and Immo Diener (1991). "Fixed Point Theorems for Discontinuous Mapping." *Mathematical Programming*, 51, 257-67.
- [31] Debreu, George and Herbert Scarf (1963). "A Limit Theorem on the Core of an Economy." *Econometrica*.
- [32] Dierker, Egbert (1971). "Equilibrium Analysis of Exchange Economies with Indivisible Commodities." *Econometrica*, 39(6), 997-1008.
- [33] Dubins, L.E. and E.H. Spanier (1961). "How to Cut a Cake Fairly." *American Mathematical Monthly*, 68, 1-17.
- [34] Dworkin, Ronald (1981). "What is Equality? Part 2: Equality of Resources." *Philosophy and Public Affairs*, 10(4), 283-345.
- [35] Dworkin, Ronald (2000). *Sovereign Virtue: The Theory and Practice of Equality*. Harvard University Press.
- [36] Ehlers, Lars and Bettina Klaus (2003). "Coalitional Strategy-proof and Resource-Monotonic Solutions for Multiple Assignment Problems." *Social Choice and Welfare*, 21: 265-80.
- [37] Foley, Duncan (1967). "Resource Allocation and the Public Sector." *Yale Economic Essays*, 7, 45-98.
- [38] Friedman, Milton (1991). "How to Sell Government Securities." *Wall Street Journal*, August 28.
- [39] Gale, David and Lloyd Shapley (1962). "College Admissions and the Stability of Marriage." *American Mathematical Monthly*, 69(1), 9-15.
- [40] Geller, William (1986). "An Improved Bound for Approximate Equilibria." *Review of Economic Studies*, LIII, 307-8.
- [41] Graves, Robert L., Linus Schrage and Jayaram Sankaran (1993). "An Auction Method for Course Registration." *Interfaces*, 23(5), 81-92.
- [42] Guernsey, Lisa (1999). "Business School Puts Courses in Hands of an On-Line Market." *New York Times*. Sept 9.
- [43] Harsanyi, John C. (1953). "Cardinal Utility in Welfare Economics and in the Theory of Risk-taking." *Journal of Political Economy*, 61(5), 434-435.
- [44] Hatfield, John William (2009). "Strategy-proof, Efficient, and Nonbossy Quota Allocations." *Social Choice and Welfare*, 33(3), 505-15.
- [45] Herreiner, Dorothea and Clemens Puppe (2002). "A Simple Procedure for Finding Equitable Allocations of Indivisible Goods." *Social Choice and Welfare*, 19, 415-30.
- [46] Hylland, Aanund and Richard Zeckhauser (1979). "The Efficient Allocation of Individuals to Positions." *Journal of Political Economy*, 87(2), 293-314.
- [47] Kaplow, Louis and Steven Shavell (2001). "Any Non-Welfarist Method of Policy Assessment Violates the Pareto Principle." *Journal of Political Economy*, 109(2), 281-286.
- [48] Kaplow, Louis and Steven Shavell (2007). "Moral Rules, the Moral Sentiments, and Behavior: Toward a Theory of an Optimal Moral System." *Journal of Political Economy*, 115(3), 494-514.
- [49] Klaus, Bettina and Eiichi Miyagawa (2001). "Strategy-proofness, Solidarity, and Consistency for Multiple Assignment Problems." *International Journal of Game Theory*, 30: 421-435.
- [50] Konishi, Hideo, Thomas Quint and Jun Wako (2001). "On the Shapley-Scarf economy: the case of multiple types of indivisible goods." *Journal of Mathematical Economics*, 35: 1-15.
- [51] Kojima, Fuhito (2009). "Random Assignment of Multiple Indivisible Objects." *Mathematical Social Sciences*, 57, 134-42.
- [52] Kojima, Fuhito and Parag Pathak (2009). "Incentives and Stability in Large Two-Sided Matching Markets." *American Economic Review*, 99, 608-27.
- [53] Konishi, Hideo, Thomas Quint, and Jun Wako (2001). "On the Shapley-Scarf Economy: the Case of Multiple Types of Indivisible Goods." *Journal of Mathematical Economics*, 35(1), 1-15.
- [54] Krishna, Aradhna and Utku Ünver (2008). "Improving the Efficiency of Course Bidding at Business Schools: Field and Laboratory Studies." *Management Science*, 27(March/April), 262-82.
- [55] Lehrer, Jim (2008). "College Students Squeezed by Rising Costs, Less Aid: Straining the System." *NewsHour*, aired December 9, 2008. Transcript accessed at http://www.pbs.org/newshour/bb/education/july-dec08/collegecosts_12-09.html.
- [56] Levitt, Steven (2008a). "What Do I Have in Common with Hannah Montana?" *Freakonomics Blog on the New York Times*. January 8.

- [57] Levitt, Steven (2008b). "The Coase Theorem Rules at NYU Law." Freakonomics Blog on the *New York Times*. August 11.
- [58] Lipton, Richard, Evangelos Markakis, Elchanan Mossel, and Amin Saberi (2004). "On Approximately Fair Allocations of Indivisible Goods." ACM EC'04.
- [59] Mas-Colell, Andreu (1977). "Indivisible Commodities and General Equilibrium Theory." *Journal of Economic Theory*, 16, 443-56.
- [60] Maskin, Eric S. (1987). "On the Fair Allocation of Indivisible Goods." In Arrow and the Foundations of the Theory of Economic Policy, ed. George R. Feiwel. NYU Press.
- [61] MIT (2008). "Sloan Course Bidding Information." Accessed at <https://sloanbid.mit.edu>.
- [62] McKesson (2008). Website description of eShift. Accessed October 2008 at http://www.mckesson.com/en_us/McKesson.com/For%2BHealthcare%2BProviders/Hospitals/Workforce%2BManagement%2BSolutions/eShift.html
- [63] Milgrom, Paul (2000). "Putting Auction Theory to Work: The Simultaneous Ascending Auction." *Journal of Political Economy*, 108(2), 245-72.
- [64] Milgrom, Paul (2004). Putting Auction Theory to Work. Cambridge University Press.
- [65] Milgrom, Paul (2009). "Assignment Messages and Exchanges." *American Economic Journal: Microeconomics*, 1(2), 95-113.
- [66] Moulin, Hervé (1990). "Uniform Externalities: Two Axioms for Fair Allocation." *Journal of Public Economics*, 43, 305-26.
- [67] Moulin, Hervé (1991). "Welfare Bounds in the Fair Division Problem." *Journal of Economic Theory*, 54, 321-37.
- [68] Moulin, Hervé (1995). Cooperative Microeconomics. Prentice Hall.
- [69] Mongell, Susan J. and Alvin E. Roth (1986). "A Note on Job Matching with Budget Constraints." *Economics Letters*, 21, 135-8.
- [70] Neil, Martha (2008). "NYU Students Seek Coveted Law School Classes, Will Pay Cash." *ABA Journal*. Jul 28.
- [71] Nisan, Noam (2006). "Bidding Languages for Combinatorial Auctions." In Cramton et al, 2006.
- [72] Othman, Abe, Eric Budish and Tuomas Sandholm (2010). "Finding Approximate Competitive Equilibria: Efficient and Fair Course Allocation." *Proc. of 9th Int. Conf. on Autonomous Agents and Multiagent Systems (AAMAS 2010)*.
- [73] Pápai, Szilvia (2001). "Strategyproof and Nonbossy Multiple Assignments." *Journal of Public Economic Theory*, 3(3): 257-71.
- [74] Parkes, David (2006). "Iterative Combinatorial Auctions." In Peter Cramton et al, Combinatorial Auctions. MIT Press.
- [75] Pathak, Parag and Tayfun Sönmez (2008). "Comparing Mechanisms by their Vulnerability to Manipulation." Mimeo.
- [76] Perry, Motty and Philip J. Reny (2006). "Toward a Strategic Foundation for Rational Expectations Equilibrium." *Econometrica*, 74, 1231-69.
- [77] Pratt, John W. (2007). "Fair (and Not So Fair) Division." *Journal of Risk and Uncertainty*, 35(3), 203-36.
- [78] Prendergast, Canice and Lars Stole (1999). "Restricting the Means of Exchange within Organizations." *European Economic Review*, 43, 1007-19.
- [79] Rawls, John (1971). A Theory of Justice. Revised Edition. Harvard University Press.
- [80] Roberts, John and Andrew Postlewaite (1976). "Incentives for Price-Taking Behavior in Large Exchange Economies." *Econometrica*, 44(1), 115-27.
- [81] Roth, Alvin E. (1984). "The Evolution of the Labor Market for Medical Interns and Residents: A Case Study in Game Theory." *Journal of Political Economy*, 92, 991-1016.
- [82] Roth, Alvin E. (2002). "The Economist as Engineer: Game Theory, Experimentation, and Computation as Tools for Design Economics." *Econometrica*, 70(4): 1341-1378.
- [83] Roth, Alvin E. (2007). "Repugnance as a Constraint on Markets." *Journal of Economic Perspectives*, 21(3), 37-58.
- [84] Roth, Alvin E. and Elliot Peranson (1999). "The Redesign of the Matching Market for American Physicians: Some Engineering Aspects of Economic Design." *American Economic Review*, 89, 748-80.
- [85] Roth, Alvin E. and Marilda Sotomayor (1990). Two-Sided Matching. Cambridge University Press.
- [86] Russell, Stuart and Peter Norvig (2002). Artificial Intelligence: A Modern Approach. Prentice Hall, 2nd Edition.
- [87] Rustichini, Aldo, Mark A. Satterthwaite and Steven R. Williams (1994). "Convergence to Efficiency in a Simple Market with Incomplete Information." *Econometrica*, 62(5), 1041-63.

- [88] Scarf, Herbert. (1960). "Some examples of global instability of competitive equilibria." *International Economic Review*, 1, 157-172.
- [89] Sen, Amartya (1979). "Equality of What?" Tanner Lecture on Human Values. Reprinted in *Choice, Welfare and Measurement*, 1982. Harvard University Press.
- [90] Sen, Amartya (2009). *The Idea of Justice*. Harvard University Press.
- [91] Shapley, Lloyd and Herbert Scarf (1974). "On Cores and Indivisibility." *Journal of Mathematical Economics*, 1(1), 23-37.
- [92] Shulman, Robin (2008). "Contention over Proposal for N.Y. Airports: Transportation Dept. Wants to Auction Off Takeoff, Landing Slots to Ease Delays." *Washington Post*, Sept 1, page A04.
- [93] Sönmez, Tayfun (1999). "Strategy-Proofness and Essentially Single-Valued Cores." *Econometrica*, 67(3): 677-89.
- [94] Sönmez, Tayfun and Utku Ünver (2010). "Course Bidding at Business Schools." *International Economic Review*, 51(1), 99-123.
- [95] Starr, Ross M. (1969). "Quasi-Equilibria in Markets with Non-Convex Preferences." *Econometrica*, 37(1), 25-38.
- [96] Steinhaus, Hugo (1948). "The Problem of Fair Division." *Econometrica*. 16(1), 101-4.
- [97] Svensson, Lars (1983). "Large Indivisibles: An Analysis with Respect to Price Equilibrium and Fairness." *Econometrica*, 51(4), 939-54.
- [98] Thomson, William and Hal R. Varian (1985). "Theories of Justice Based on Symmetry." In Hurwicz, Schmeidler and Sonnenschein (eds), *Social Goals and Social Organization: Essays in Memory of Elisha Pazner*. Cambridge University Press.
- [99] Varian, Hal (1974). "Equity, Envy and Efficiency." *Journal of Economic Theory*, 29(2), 217-244.
- [100] Vazirani, Vijay V. (2007). "Combinatorial Algorithms for Market Equilibria." In Nisan, Roughgarden, Tardos and Vazirani (eds), *Algorithmic Game Theory*. Cambridge University Press.
- [101] Vickrey, William (1961). "Counterspeculation, Auctions, and Competitive Sealed Tenders." *Journal of Finance*, 16, 8-37.
- [102] Wallis, W. D.; Street, Anne Penfold; Wallis, Jennifer Seberry (1972). *Combinatorics: Room Squares, Sum-free Sets, Hadamard Matrices*. Lecture Notes in Mathematics, Vol. 292. Springer-Verlag, Berlin-New York.
- [103] Wharton (2009). "The Course Registration Auction for MBA Electives." Version 16.0, Graduate Division, MBA Program Office, The Wharton School, University of Pennsylvania.
- [104] Wulf, William A. (1993). "The Collaboratory Opportunity." *Science*, 261(13), 854-5.
- [105] Yamazaki, Akira (1978). "An Equilibrium Existence Theorem without Convexity Assumptions." *Econometrica*, 46(3), 541-55.
- [106] Ziegler, Gunter M. (1995). *Lectures on Polytopes*. Springer-Verlag.
- [107] Zhou, Lin (1990). "On a Conjecture by Gale about One-Sided Matching Problems." *Journal of Economic Theory*, 52(1), 123-135.

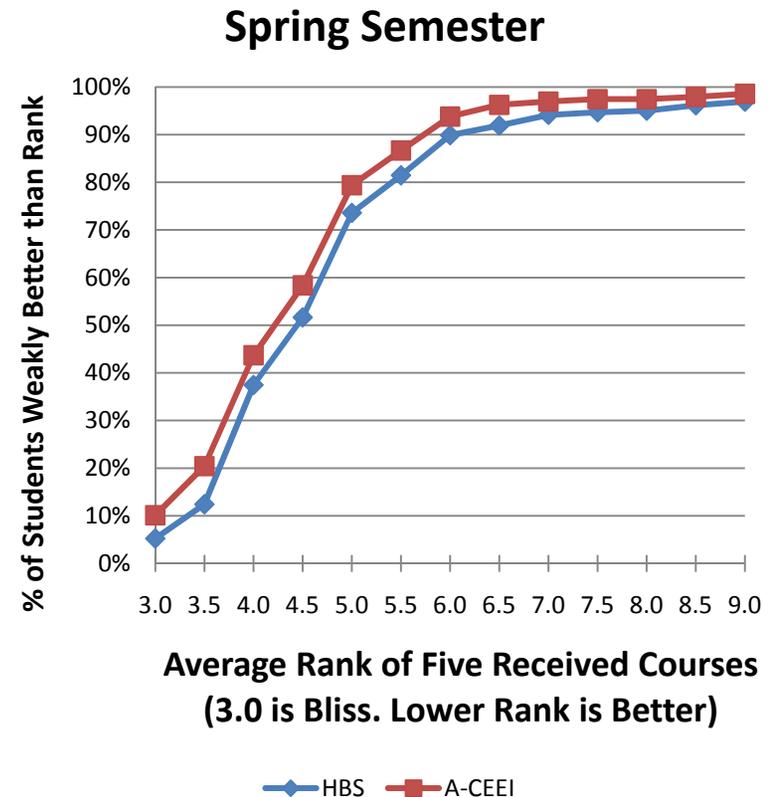
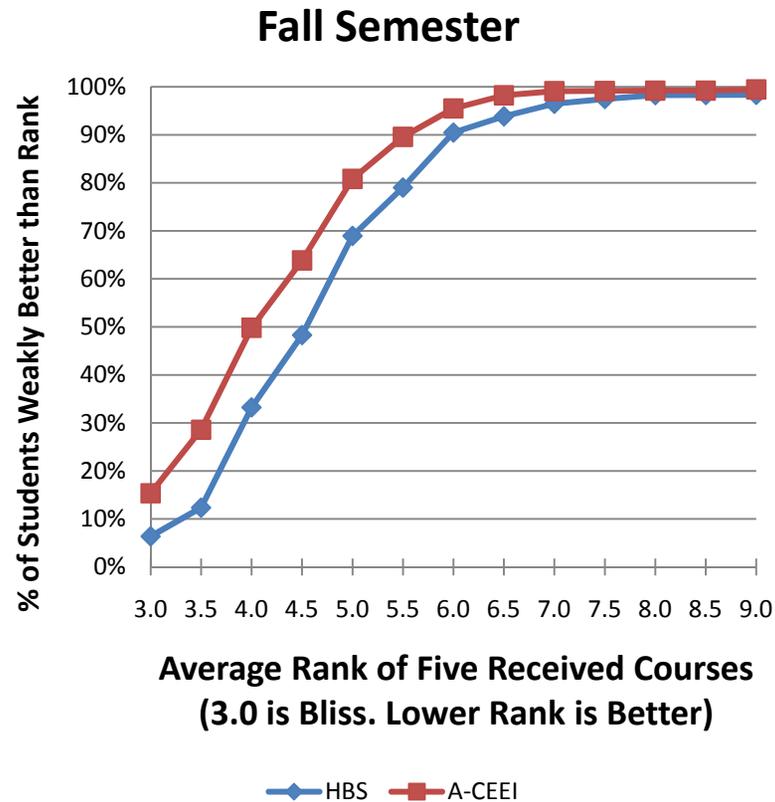
Figure 1: Market-Clearing Error



Description: The Othman, Budish and Sandholm (2010) Approximate CEEI algorithm is run 100 times for each semester of the Harvard Business School course allocation data (456 students, ~50 courses, 5 courses per student). Each run uses randomly generated budgets. This table reports the distribution of the amount of market-clearing error per trial, measured in Euclidean Distance. Both excess demand and excess supply count as error (except that courses priced at zero are allowed to be in excess supply without counting as error). See Section 9 for further description.

Figure 2: Ex-Ante Efficiency Comparison

Approximate CEEI Mechanism vs. HBS Draft Mechanism



Description: The Othman, Budish and Sandholm (2010) Approximate CEEI algorithm is run 100 times for each semester of the Harvard Business School course allocation data (456 students, ~50 courses, 5 courses per student). Each run uses randomly generated budgets. For each random budget ordering I also run the HBS Draft mechanism, using the random budget order as the draft order. The HBS Draft mechanism is run using students' actual strategic reports under that mechanism. The Approximate CEEI algorithm is run using students' truthful preferences. This table reports the cumulative distribution of outcomes, as measured by average rank, over the $456 \times 100 = 45,600$ student-trial pairs. Average rank is calculated based on the student's true preferences. For instance, a student who receives her 1,2,3,4 and 5th favorite courses has an average rank of $(1+2+3+4+5)/5 = 3$. See Section 9 for further description.

Table 1: Comparison of Alternative Mechanisms

Mechanism	Efficiency (Truthful Play)	Outcome Fairness (Truthful Play)	Procedural Fairness	Incentives	Preference Language
Approximate CEEI Mechanism (A-CEEI)	Pareto efficient w/r/t allocated goods Allocation error is small for practice and goes to zero in the limit If exact CEEI exists: Pareto efficient	N+1 – Maximin Share Guaranteed Envy Bounded by a Single Good If exact CEEI exists: Maximin Share Guaranteed and Envy Free	Symmetric	Strategyproof in the Large	Ordinal over Schedules
A-CEEI with a Pareto-Improving Secondary Market (Appendix E)	Pareto efficient	A bit weaker than N+1 – Maximin Share Guarantee, because prices in the initial allocation may be outside of $P(\delta, b')$. Initial allocation is Envy Bounded by a Single Good. The Pareto-improvement stage may exacerbate envy.	Symmetric	Manipulable in the Large	Ordinal over Schedules
Competitive Equilibrium from Equal-as-Possible Incomes (Appendix E)	Pareto efficient	Worst Case: coincides with dictatorship	Symmetric	Strategyproof in the Large	Ordinal over Schedules
Random Serial Dictatorship (Sec 8.1)	Pareto efficient	Worst Case: Get k worst Objects	Symmetric	Strategyproof	Ordinal over Schedules
Sequential or Serial Dictatorship (Refs in fn 6)	Pareto efficient	Worst Case: Get k worst Objects	Not Symmetric	Strategyproof	Ordinal over Schedules
Multi-unit generalization of Hylland Zeckhauser Mechanism (Sec 8.2)	If vNM preferences are described by assignment messages, Pareto efficient	If preferences are additive separable, envy bounded by the value of two goods Worst Case: Get Zero Objects	Symmetric	If vNM preferences are described by assignment messages, Strategyproof in the Large	Assignment messages
Bidding Points Auction (Sec 8.3)	If preferences are additive-separable, Pareto efficient but for quota issues described in Sonmez and Unver (2010)	Worst Case: Get Zero Objects	Symmetric	Manipulable in the Large	Cardinal over Items

Table 1: Comparison of Alternative Mechanisms (cont.)

Mechanism	Efficiency (Truthful Play)	Outcome Fairness (Truthful Play)	Procedural Fairness	Incentives	Preference Language
Sonmez-Unver (2010) Enhancement to Bidding Points Auction	If preferences are additive- separable, Pareto efficient	Worst Case: Get Zero Objects	Symmetric	Bidding Phase: Manipulable in the Large Allocation Phase: Strategyproof in the Large	Bidding Phase: Cardinal over Items Allocation Phase: Ordinal over Items
HBS Draft Mechanism (Sec 9)	If preferences are responsive, Pareto efficient with respect to the reported information (i.e., Pareto Possible)	If preferences are responsive and $k=2$, Maximin Share Guaranteed If preferences are responsive, Envy Bounded by a Single Good	Symmetric	Manipulable in the Large	Ordinal over Items
Utilitarian Solution (cf. Sen 1979)	Pareto Efficient	Worst Case: get zero objects (if a depressive and all other agents are hedonists)	Symmetric	Manipulable in the Large	Cardinal over Schedules
Bezakova and Dani (2005) Maximin Utility Algorithm (cf. Rawls 1971)	If preferences are additive- separable, ideal fractional allocation is Pareto efficient. Realized integer allocation is close to the fractional ideal.	Worst Case: Get approximately zero objects (if a hedonist and all other agents are depressives)	Symmetric	Manipulable in the Large	Cardinal over items
Brams and Taylor (1996) Adjusted Winner	If preferences are additive- separable, Pareto efficient	Worst Case: Get Zero Objects	Symmetric	Manipulable in the Large	Cardinal over Items
Herreiner and Puppe (2002) Descending Demand Procedure	Pareto efficient	Does not satisfy Maximin Share Guarantee or Envy Bounded by a Single Object	Symmetric	Manipulable in the Large	Ordinal over Schedules
Lipton et al (2004) Fair Allocation Mechanism	Algorithm ignores efficiency	If preferences are additive separable, Envy Bounded by a Single Good	Symmetric	Manipulable in the Large	Cardinal over items
University of Chicago Primal-Dual Linear Programming Mechanism (Graves et al 1993)	Pareto efficient when preference- reporting limits don't bind	Worst Case: Get Zero Objects	Symmetric	Manipulable in the Large	Cardinal over a Limited Number of Schedules

Table 2: Degree of Ex-Post Envy

Amount of Envy, in Ranks	% of Students with this Amount of Envy	
	Fall Semester	Winter Semester
No Envy	98.759%	99.522%
1	0.860%	0.274%
2	0.211%	0.171%
3	0.151%	0.013%
4	0.009%	0.020%
5	0.011%	0.000%
>5	0.000%	0.000%

Description: The Othman, Budish and Sandholm (2010) Approximate CEEI algorithm is run 100 times for each semester of the Harvard Business School course allocation data (456 students, ~50 courses, 5 courses per student). Each run uses randomly generated budgets. This table reports the distribution of the amount of envy students experience, measured in ranks. Here is an example calculation: a student who receives the bundle consisting of her 2,3,4,6 and 7th favorite courses, while some other student with a larger budget receives the bundle consisting of her 1,2,3,4 and 5th favorite courses, experiences envy in ranks of $(2+3+4+6+7) - (1+2+3+4+5) = 7$. See Section 9 for further description.