Efficient Disposal Equilibria of Pseudomarkets

Andrew McLennan*
School of Economics
The University of Queensland
a.mclennan@economics.uq.edu.au

February 3, 2020

Abstract. For an economy with compact consumption and production sets and some goods that can be freely disposed, an *efficient disposal equilibrium* specifies prices, consumptions, and production plans such that: a) each agent maximizes utility among bundles costing no more than her income and minimizes expenditure among bundles providing the same utility; b) an unsated agent consumes a bundle that is at least as valuable as her income; c) each producer maximizes profits; d) the aggregate endowment plus aggregate production minus aggregate consumption is a nonnegative bundle of disposable goods; e) disposable goods that are not completely consumed have the minimal price of disposable goods. We prove an existence of equilibrium result that nests those of Hylland and Zeckhauser (1979a), Mas-Colell (1992), and Budish, Che, Kojima, and Milgrom (2013). It significantly improves the latter by increasing flexibility and relaxing assumptions that are not satisfied by applications such as course allocation. Open problems concerning generic finiteness of the set of equilibria and efficient algorithms for computing equilibria are described.

Key Words: Market design, general equilibrium, competitive equilibrium, pseudomarkets, assignment problems, course allocation, PPAD, generic finiteness, mechanism design.
JEL Classification Numbers: D47, D50, D82

*I have benefitted from helpful comments by Federico Echenique, Hülya Eraslan, Hervé Moulin, Tim Roughgarden, Shino Takayama, and Richard Zeckhauser.

1 Introduction

This paper establishes an existence of general competitive equilibrium result that subsumes the existence results provided by Hylland and Zeckhauser (1979a), Mas-Colell (1992), and Budish, Che, Kojima, and Milgrom (2013). It also subsumes the quite general classical results concerning existence of equilibrium that are special cases of Mas-Colell's result. A common theme of these papers is the use of so-called pseudomarkets to achieve efficient and equitable division of a collectively owned endowment, so we begin with an overview of this line of research.

It has long been recognized (Varian (1974)) that for a given bundle of goods owned collectively by a group of agents, a competitive allocation of the exchange economy in which the agents have equal incomes (or equivalently, each agent is endowed with her pro rata share of the common endowment) gives an allocation that is both efficient and envy free. When there are finitely many divisible goods and finitely many agents, existence of such an allocation is a special case of standard existence results. (Weller (1985) establishes existence for a cake cutting problem in which the collective endowment is a measurable space and each agent's utility function is an atomless measure.) The study of such allocations goes back at least to Eisenberg and Gale (1959) who noted that the equilibrium of a pari-mutuel betting system has this form when the bettors have equal stakes. Generalizing this work, Eisenberg (1961) showed that when the agents' utility functions are concave and homogeneous of degree 1, the unique equilibrium-from-equal-incomes allocation gives the Nash (1950) bargaining solution, which is to say that it maximizes the sum of the logarithms of the agents' utilities, and is thus the solution of a convex program.

It is natural to try to apply this general idea to other types of fair division problems, and explorations of this sort have led to novel settings and challenges. Bogomolnaia, Moulin, Sandomirskiy, and Yanovskaya (2016b) study the allocation of bads, finding that it is not the simple mirror image of the allocation of goods. Bogomolnaia, Moulin, Sandomirskiy, and Yanovskaya (2017, 2016a) provide generalizations of Eisenberg's characterization results for allocation problems in which there are both goods and bads, and commodities which are good for some agents and bad for others.

Hylland and Zeckhauser (1979a) (henceforth HZ) pioneered the application of competitive-equilibrium-from-equal-incomes to the allocation of indivisible objects. They consider a model in which each element of a set of agents is to be matched with exactly one element of a finite set of objects. Each object has an integral capacity, and the sum of the capacities is at least as large as the number of agents. (One of the objects may represent being unassigned.) There is a profile of vNM utility functions assigning a utility of each object for each agent. HZ propose endowing all agents with positive incomes, and then allowing the agents to trade in a market in which the goods are probabilities of being assigned to the various objects. They show that an equilibrium exists for any assignment of incomes. An easy generalization of the Birkhoff-von Neumann theorem implies that if a profile of probability distributions for the agents does not assign more probability to any object than its capacity, then it can be implemented by a probability distribution over pure assignments. As of July 2018, HZ has been cited over 460 times, and it is often described as a seminal touchstone of the market design literature, but there has been very little technical follow up, so their existence results have not been well understood in relation to some larger framework.

Budish, Che, Kojima, and Milgrom (2013) (henceforth BCKM) study a course allocation problem in which each of finitely many students must receive a package of seats in various courses. Each student's allocation is constrained by a system of linear inequalities, each of which requires that the number of courses in a certain set not lie above a quota (ceiling constraint) or not lie below a quota (floor constraint). They provide a deep analysis of the conditions that the various constraints must satisfy in order to insure that an assignment of probabilities satisfying them can be implemented as a probability distribution over pure assignments. They also prove existence of a competitive equilibrium for a market in which probabilities of receiving seats are traded. In the Online Appendices (pp. 23–24) they point out that such an equilibrium need not be efficient if the agents are not required to minimize expenditure among the bundles providing the equilibrium utility, and they show how to modify their argument to attain this condition.

Their existence result requires that each agent's consumption set is the portion of the positive orthant satisfying a system of ceiling constraints (that is, there are no floor constraints) and that the set of bundles of seats that the school can provide is a singleton. It can be used to prove the HZ existence result in the following manner. Instead of requiring agents to consume a probability distribution, allow them to consume anywhere in the portion of the positive orthant lying below the probability simplex, but rescale utilities so that each good has a positive utility for each agent. In equilibrium all agents consume probability distributions because otherwise some good is incompletely consumed and has price zero.

A number of recent papers (Le (2018), He, Miralles, Pycia, and Yan (2018), Echenique, Miralles, and Zhang (2018)) study models similar to the one in BCKM.

There is a strand of general equilibrium literature (Bergstrom (1976), Polemarchakis and Siconolfi (1993), Mas-Colell (1992)) (henceforth M-C) that considers possibly sated consumers, but without free disposal. M-C attains the most general result, which asserts the existence of an equilibrium in which the unspent income of sated consumers is redistributed to unsated consumers. (The results in M-C imply the existence claims in Bogomolnaia, Moulin, Sandomirskiy, and Yanovskaya (2017), which are proved in that paper in a more constructive manner.)

We provide an existence of equilibrium result that subsumes the results described above. In comparison with HZ and BCKM our result is more general in several ways such as more general consumption sets and nonlinear utilities, and because the consumers are endowed with commodity bundles in addition to artificial currency. (In Section 2 we explain that although our model requires the aggregate endowment to be the sum of individual commodity endowments, the case of pure currency endowment is a special case.) In comparison with M-C our result is more general because we allow some goods to be freely disposable.

These generalizations significantly enhance the applicability of the model. As Echenique, Miralles, and Zhang (2018) stress, commodity endowments can be used to guarantee certain types of entitlements and fairness. In the context of school choice, they can insure that a student has the option of going to a school in her walk zone, or the school a sibling is attending. In the course allocation problem, commodity endowments can insure that a student has access to at least one program of study that fulfills the requirements of her degree program. In most practical applications of matching the total capacity will exceed the number of agents, and since the sum of the agents' initial endowment must be society's endowment, free disposal is required in order to satisfy the constraint that each agent's total assignment probability is one. In the context of the BCKM result we are able to dispense with the assumption that there are no floor constraints, and that the school can offer only one bundle of seats. As we explain in more detail in Section 3, our result provides many dimensions of flexibility that seem natural, and potentially quite useful, in the course allocation application. Thus we provide a powerful unified framework that easily encompasses existing results and presents many new possibilities for market design.

We can briefly describe the main technical ideas. When consumers have compact consumption sets and preferences are only assumed to be continuous and convex, there are no restrictions on the equilibrium price vector. In particular, it is possible that all prices are zero, and indeed when it is possible for all agents to simultaneously consume bliss points, this is an equilibrium. However, the budget sets of consumers are not lower hemicontinuous at this price vector. A second difficulty is that when some goods can be freely disposed, the value of aggregate consumption may not be equal to the sum of the values of the aggregate endowment and the aggregate production. In particular, in order to maximize utility of net consumption a consumer may need to purchase a gross (prior to disposal) bundle that is less valuable than her income. Each of these problems invalidates standard methods of proof. To overcome both of them we introduce an additional artificial good that is always desirable. In an equilibrium of the expanded economy the price of this good is necessarily positive, even if the prices of the given goods are all zero. In the expanded economy consumers spend all income because any income that is not used to improve the bundle of given goods can be used to increase consumption of the artificial good. We take a sequence of expanded economies along which the aggregate endowment of the artificial good goes to zero, finding that the limiting (along a suitable subsequence) prices, consumptions, and productions have all desired properties.

This approach seems quite novel for the general equilibrium literature. Hart and Kuhn (1975) is in a sense similar insofar as the topological basis of their proof is unusual, but they study economies with unbounded consumption sets and no free disposal. M-C uses a construction originated by Bergstrom (1976) in which the space of prices is the unit ball and the budget constraint is relaxed as the price vector approaches the origin.

Other Related Literature

Recently it has been increasingly recognized that the general notion of using pseudomarkets to resolve allocation problems for which monetary transfers are infeasible or undesirable, is in many ways quite appealing. Most obviously, one expects Pareto optimal outcomes. If all agents have the same income, pseudomarket equilibria yield envy free outcomes and are fair in an ex ante sense, although a symmetric endowment may admit asymmetric equilibria, so certain kinds of ex post fairness are not guaranteed. HZ produce examples showing that their mechanism is not strategy proof, insofar as an agent may obtain a more favorable price by misreporting her vNM utility over objects, but one may expect that pseudomarkets become asymptotically strategy proof in the sense of Roberts and Postlewaite (1976) as the economy becomes large. (Cf. Theorem 2 of He, Miralles, Pycia, and Yan (2018).)

If cardinal utilities are part of the input to a pseudomarket, it may provide efficiency gains over mechanisms such as random priority and the probabilistic serial rule (Bogomolnaia and Moulin (2001)) that have ordinal preferences as their input. In fact Liu and Pycia (2008) show that all ordinal mechanisms satisfying certain conditions are asymptotically equivalent, and Pycia (2014) shows that the loss in comparison with efficient cardinal mechanisms may be arbitrarily large. (Featherstone and Niederle (2008), Miralles (2009), Abdulkadiroğlu, Che, and Yasuda (2011), Troyan (2012), Abdulkadiroğlu, Che, and Yasuda (2015), and Ashlagi and Shi (2016) discuss this issue in the context of school choice.)

Pratt and Zeckhauser (1990) is a concrete application of HZ, and He, Miralles, Pycia, and Yan (2018), Le (2018), and Echenique, Miralles, and Zhang (2018) (which is further described in Section 2) are recent papers proposing variations or extensions of the HZ mechanism. Pratt (2007) and Budish (2011) (see also Budish and Kessler (2016)) are other papers describing pseudomarkets, and of course a wide variety of mechanisms for allocating objects without monetary payments (e.g., Sönmez and Ünver (2010) and Budish and Cantillon (2012)) have market-like features. In particular, Immorlica, Lucier, Mollner, and Weyl (2017) study a "tricky tray" raffle mechanism that has incentives resembling those of the HZ mechanism.

Contents

The next section defines efficient disposal equilibrium and states Theorem 1, which is our main result. Section 3 explains in detail how Theorem 1 implies the various existence results mentioned above, and it explains the additional generality of our result beyond that of BCKM, and how this may be useful in the context of course allocation. Section 4 presents the proof of Theorem 1.

Almost all of the traditional concerns of mathematical economics are applicable to the EDE concept, so there are many new problems and possible conjectures. Section 5 highlights two of these, which already arise in the HZ model with each object having capacity one, the number of agents equal to the number of objects, and each agent endowed with an equal share of all objects. First, is the set of equilibria finite for generic utility profiles? Second, is the problem of computing an EDE a PPAD complete computational problem?

Appendix A provides a brief review of the theory of the vector field index and the Poincaré-Hopf theorem. Appendix B contains proofs omitted from the body of the paper.

2 Efficient Disposal Equilibrium

We work in a classical general equilibrium setting with ℓ commodities indexed by h, m consumers indexed by i, and n producers indexed by j. For nonnegative integers ℓ_c and ℓ_d such that $\ell_c + \ell_d = \ell$, the goods indexed by $1, \ldots, \ell_c$ are not disposable and the goods indexed by $\ell_c + 1, \ldots, \ell$ are freely disposable. The *disposal cone* is

$$C = \{ x \in \mathbb{R}^\ell_+ : x_1 = \cdots = x_{\ell_c} = 0 \}.$$

Consumer i has an *endowment* $\omega_i \in \mathbb{R}^\ell$, a nonempty compact convex *consumption set* $X_i \subset \mathbb{R}^\ell$, and a continuous *utility function* $u_i : X_i \to \mathbb{R}$ that is semi-strictly quasiconcave[1]: that is, for all $x_i^0, x_i^1 \in X_i$ and $t \in (0,1)$, if $u_i(x_i^1) > u_i(x_i^0)$, then $u_i\big((1-t)x_i^0 + tx_i^1\big) > u_i(x_i^0)$. If $x_i \in X_i$ and $u_i(x_i) = \max_{x_i' \in X_i} u_i(x_i')$, then we say that consumer i is *sated* at x_i, and that x_i is a *bliss point* for i, and otherwise we say that i is *unsated* at x_i.

Producer j has a compact convex *production set* $Y_j \subset \mathbb{R}^\ell$ that contains the origin. For each j and $p \in \mathbb{R}^\ell$ let maximal profits and the set of optimal production plans be

$$\pi_j(p) = \max_{y_j \in Y_j} \langle p, y_j \rangle \quad \text{and} \quad M_j(p) = \operatorname{argmax}_{y_j \in Y_j} \langle p, y_j \rangle.$$

There is an $m \times n$ matrix θ of ownership shares with $\theta_{ij} \geq 0$ for all i and j and $\sum_i \theta_{ij} = 1$ for all j. For each i and $p \in \mathbb{R}^\ell$, i's income is

$$\mu_i(p) = \langle p, \omega_i \rangle + \sum_j \theta_{ij} \pi_j(p).$$

Berge's theorem gives:

Lemma 1 *For each j, $\pi_j : \mathbb{R}^\ell \to \mathbb{R}_+$ is a continuous function and $M_j : \mathbb{R}^\ell \to Y_j$ is an upper hemicontinuous convex valued correspondence. Consequently each $\mu_i : \mathbb{R}^\ell \to \mathbb{R}$ is continuous.*

[1] A simple argument applying the intermediate value theorem shows that a continuous and semi-strictly quasiconcave u_i is quasiconcave: for all $x_i^0, x_i^1 \in X_i$ and $t \in [0,1]$, if $u_i(x_i^1) \geq u_i(x_i^0)$, then $u_i\big((1-t)x_i^0 + tx_i^1\big) \geq u_i(x_i^0)$.

Let $\omega = \sum_i \omega_i$, and let $X = \prod_i X_i$ and $Y = \prod_j Y_j$. A triple $(p, x, y) \in (\mathbb{R}^{\ell_c} \times \mathbb{R}^{\ell_d}_+) \times X \times Y$ is an *efficient disposal equilibrium* (EDE) if:

(a) For each i there is no $x'_i \in X_i$ such that either $\langle p, x'_i \rangle \leq \langle p, x_i \rangle$ and $u_i(x'_i) > u_i(x_i)$ or $\langle p, x'_i \rangle < \langle p, x_i \rangle$ and $u_i(x'_i) \geq u_i(x_i)$.

(b) For each i, if i is unsated at x_i, then $\langle p, x_i \rangle \geq \mu_i(p)$.

(c) For each j, $y_j \in M_j(p)$.

(d) $\omega + \sum_j y_j - \sum_i x_i \in C$.

(e) For each $h = \ell_c + 1, \ldots, \ell$, if $\sum_i x_{ih} < \omega_h + \sum_j y_{jh}$, then $p_h = 0$.

The less common elements of this definition are that consumers minimize expenditure, subject to attaining the equilibrium utility, and that the price of a good that is disposed of is not greater than the price of another disposable good. Of course for many general equilibrium models these are consequences of the model's assumptions and other equilibrium conditions.

In an EDE (p, x, y) the excess income of sated consumers is redistributed to unsated consumers. Let S (U) be the set of i that are sated (unsated) at x_i. For $\alpha \in \mathbb{R}^m_{++}$ we say that (p, x, y) is an α-EDE if there is a $\Pi \geq 0$ such that $\langle p, x_i \rangle - \mu_i(p) = \Pi \alpha_i$ for all $i \in U$. As in Le (2018) (the correct interpretation of the α in Echenique, Miralles, and Zhang (2018) is somewhat different) we may interpret α as a vector of endowments of currency, so that Π is the purchasing power of this currency relative to p. Combining (d) and (e) gives $\langle p, \sum_i x_i \rangle = \langle p, \omega + \sum_j y_j \rangle$, so $\sum_i \langle p, x_i \rangle = \langle p, \sum_i (\omega_i + \sum_j \theta_{ij} y_j) \rangle = \sum_i \mu_i(p)$. For each $i' \in U$, $\alpha_{i'} / \sum_{i \in U} \alpha_i$ is i''s share of the unspent income of sated consumers because

$$\Pi \cdot \sum_{i \in U} \alpha_i = \sum_{i \in U} \langle p, x_i \rangle - \mu_i(p) = \sum_{i \in S} \mu_i(p) - \langle p, x_i \rangle.$$

We now state the main result. Let $\mathbf{e} = (0, \ldots, 0, 1, \ldots, 1) \in \mathbb{R}^\ell$ be the vector such that $\mathbf{e}_h = 1$ if $\ell_c + 1 \leq h \leq \ell$ and $\mathbf{e}_h = 0$ otherwise. Let $V_0 = \{ x \in \mathbb{R}^\ell : \langle \mathbf{e}, x \rangle = 0 \}$ be the orthogonal complement of $\mathbf{e}$.

Theorem 1 *Suppose that:*

(a) For each i there is an $x^0_i \in X_i$ such that $X_i \subset x^0_i + V_0$, x^0_i is in the interior (relative to $x_i + V_0$) of X_i, and $\omega_i - x^0_i \in C$.

(b) For each j, $Y_j \subset V_0$.

Then for any $\alpha \in \mathbb{R}^m_{++}$ there is an α-EDE.

Several remarks are in order:

- The requirement that each consumption set is contained in a translate of V_0 arises naturally in HZ, where the consumption set is the set of probability distributions over a finite set of objects. In many settings it is natural to introduce an additional disposable good that may be thought of as the negation of the sum of the other disposable goods, in which case it is natural to have each production set contained in V_0 and each consumption set contained in a translate of V_0. As we will see in the next section, one can pass from the model of BCKM to our framework using this technique.

- Note that we do not assume that the endowments and consumption sets stand in any particular relationship to the positive orthant. The critical assumption is that each consumer can reach a point in the interior of the consumption set using only free disposal.

- The existence proofs of HZ and BCKM endow the consumers with positive incomes. A setting in which the expenditure of each unsated consumer i is proportional to α_i can be realized in our framework by setting $\omega_i = \alpha_i \omega$ and $\theta_{ij} = \alpha_i$ for all j. In this sense our framework is more general and flexible. (But note that when the commodities are bads, a larger α_i may be a burden rather than a blessing, and when the endowment includes both goods and bads, the sign of the value of the aggregate endowment may not be known a priori, or may be different in different equilibria.)

- HZ and BCKM are unable to work with commodity bundle endowments because the budget set fails to be lower hemicontinuous at the zero price vector. (See pp. 23-24 of BCKM's Online Appendices for a more complete description of the mathematical difficulties.) A simple example due to HZ (footnote 14, pp. 301–2) illustrates these issues. Let $\ell = 2$, $\omega = (1, 2)$, $m = 3$, and $X_i = \{ x_i \in \mathbb{R}^2_+ : x_{i1} + x_{i2} = 1 \}$ for all i, and suppose that each consumer's utility function is affine. (This is the problem of probabilistically allocating two objects with capacities 1 and 2.) Assume that consumers 1 and 2 prefer good 1 to good 2 and consumer 3 prefers good 2 to good 1.

As HZ point out, if each consumer is endowed with $(\frac{1}{3}, \frac{2}{3})$ and has no other source of income, then there is no equilibrium: if good 1 is more expensive than good 2, then consumers 1 and 2 consume their endowments while consumer 3 consumes $(0, 1)$, and otherwise good 1 is overdemanded.

- Following Gale and Mas-Colell (1975, 1979), M-C allows each consumer's income to be a continuous function of the price vector p, and also the production vector y, that is more general than the sum of the value of the consumer's endowment and her shares of the firms' profits. In addition, consumers' preferences may be affected by quite general externalities. These possibilities may be realistic or conceptually interesting, and in those papers pointing them out clarifies the foundations of the results without introducing burdensome complications. Such additional generality seems to be possible here as well, but for the sake of avoiding distractions we have not done this formally.

- Echenique, Miralles, and Zhang (2018) study a pseudomarket model in which, for some $\alpha \in (0, 1]$, each consumer's budget constraint is $\langle p, x_i \rangle \leq \alpha + (1 - \alpha)\langle p, \omega_i \rangle$. (There is no production.) Superficially this seems quite similar, but in fact it is a distinct model with its own features. To see this let $\ell = 2$, $m = 2$, $X_i = \{ x_i \in \mathbb{R}^2_+ : x_{i1} + x_{i2} = 1 \}$ for both i, $\omega_1 = (\frac{1}{3}, \frac{2}{3})$, and $\omega_2 = (\frac{2}{3}, \frac{1}{3})$, and suppose that both utility functions are affine, with both consumers preferring good 1 to good 2. For the model of this paper the only equilibrium allocation is the initial endowment, but as α goes from 0 to 1 the equilibrium allocation of their model goes from the initial endowment to equal division. Whereas the vector α of our model is used to redistribute excess income of sated consumers, the parameter α of their model traces a path from the endowment economy to equal incomes.

- The usual argument proves the first fundamental theorem for EDE's: in a Pareto improving allocation each consumer would necessarily be spending at least as much, and some consumer would be spending more, but since incompletely consumed goods have price zero, the value of the equilibrium allocation is equal to the value of the initial endowment plus aggregate profits, which were already maximal. (The connection between expenditure minimization and efficiency has been noted by many authors, including Eisenberg (1961), Bogomolnaia, Moulin, Sandomirskiy, and Yanovskaya

(2017), the technical appendices of HZ (Hylland and Zeckhauser (1979b), henceforth HZA), M-C, BCKM (p. 24 of the Online Appendices), Le (2018), and Echenique, Miralles, and Zhang (2018).)

3 Comparison with Other Existence Results

In this section we explain, in some detail, how Theorem 1 implies the equilibrium existence results of M-C, HZ, and BCKM, emphasizing how the additional flexibility we provide is applicable to the course allocation application. The next section proves Theorem 1. These two sections can be read in either order.

There are several papers in the general equilibrium literature in which consumers have compact consumption sets $X_i \subset \mathbb{R}^\ell$, each ω_i is a point in $\mathbb{R}^\ell$ that is typically contained in the interior of X_i, and the production sets Y_j are compact subsets of $\mathbb{R}^\ell$ containing the origin. This is the special case of our framework in which there are no disposable goods. Bergstrom (1976) attains a competitive equilibrium existence result by assuming that at any feasible allocation all consumers are locally nonsatiated. Note that this result subsumes various classical existence results for consumers with monotonic preferences who are necessarily unsated at equilibrium if the consumption sets are sufficiently large compact subsets of $\mathbb{R}^\ell_+$. Polemarchakis and Siconolfi (1993) prove existence of a weak (insofar as consumers are not allowed to consume bundles costing less than their incomes) competitive equilibrium, and give conditions somewhat more general than Bergstrom's under which a competitive equilibrium exists. M-C weakens the notion of competitive equilibrium by allowing the excess income of sated consumers to be redistributed to the other consumers. (His definition of slack corresponds to setting $\alpha = (1, \dots, 1)$, but he notes (p. 204) that his framework allows much more general forms of redistribution.) His notion of *strong Walrasian equilibrium* requires expenditure minimization in addition to utility maximization, as per (a) of the definition of an EDE. Roughly, his Theorem 3 corresponds to the special case of our Theorem 1 given by $\ell_c = \ell$, so that no goods are disposable and each ω_i is an element of the interior of X_i. As we have already noted, his framework incorporates more general income functions and consumption externalities that can, to at least some extent, be included in our framework.

In the HZ model the commodities are the probabilities of being assigned to each of the

ℓ objects, so each X_i is the unit simplex $\{ x \in \mathbb{R}^\ell_+ : \sum_h x_h = 1 \}$. There is no production. For each h and i, u_{ih} is consumer i's utility if assigned to object h, and $u_i(x_i) = \sum_h u_{ih} x_{ih}$ is consumer i's expected utility. There is a collectively owned endowment ω of capacities where each ω_h is a positive integer, and $\sum_h \omega_h \geq m$, so that feasible allocations exist. HZ propose endowing the consumers with positive amounts $B_1, \ldots, B_m$ of some artificial currency, and in the discussion in Appendix D of HZA they demonstrate existence of an EDE in which the expenditures of the unsated consumers are proportional to their incomes. This result follows from Theorem 1 if we let $\omega_i = (B_i / \sum_{i'} B_{i'})\omega$ and $\alpha = (B_1, \ldots, B_m)$.

In the setup of Theorem 6 of BCKM each consumer's consumption set $\hat{X}_i$ is a compact subset of $\mathbb{R}^{\ell-1}_+$ given by a finite number of inequalities of the form $\sum_h c_h \hat{x}_{ih} \leq C$ where $c_1, \ldots, c_{\ell-1} \in \{0, 1\}$ and C is a positive integer. (That is, there are ceiling constraints but no floor constraints.) Note that the origin is an element of $\hat{X}_i$. There is a collectively owned endowment $\hat{\omega} \in \mathbb{R}^{\ell-1}_{++}$ of capacities, and there is no production. As in the HZ model they assume linear utility. The statement of Theorem 6 asserts that there is a competitive equilibrium from equal wealths in which goods that are not fully demanded have price zero, and on p. 24 of their Online Appendix they explain how to prove existence of an equilibrium in which sated consumers minimize expenditure.

To understand this strengthened version of their result as a consequence of Theorem 1 we introduce an additional good ℓ, thought of in the context of course allocation as "not taking any course," whose consumption gives no utility. Let $\ell_c = 0$ and $\ell_d = \ell$, so all goods are disposable. For each i let $b_i = \max_{\hat{x}_i \in \hat{X}_i} \sum_h \hat{x}_{ih}$, and let $X_i = \{ (\hat{x}_{i1}, \ldots, \hat{x}_{i,\ell-1}, b_i - \sum_h \hat{x}_{ih}) : \hat{x}_i \in \hat{X}_i \}$. Let B be a positive number greater than $\sum_i b_i$, and for each i let $\omega_i = (\hat{\omega}_1/m, \ldots, \hat{\omega}_{\ell-1}/m, B)$. Let $V_0 = \{ x \in \mathbb{R}^\ell : \sum_h x_h = 0 \}$. Since the origin is an element of X_i and $B > b_i$, an x_i^0 in the interior of X_i near $(b_i, 0, \ldots, 0)$ satisfies $x_i^0 \ll \omega_i$, so the hypotheses of Theorem 1 are satisfied, and applying it with $k = \ell$ and $\alpha = (1, \ldots, 1)$ gives an EDE of the economy we have constructed. Since the supply of good 0 exceeds any possible demand, its price is zero, and the projection of this EDE onto the space of the given economy is easily seen to be an EDE for it.

As mentioned previously, Echenique, Miralles, and Zhang (2018) prove existence of an α-slack equilibrium where $\alpha \in (0, 1]$ is the parameter of a homotopy between an initial endowment and equal endowments of currency ($\alpha = 1$). This does not seem to follow directly from Theorem 1. However, it seems likely that the argument presented in Section

4 can be adapted to give a rather heavy handed proof. The key point is that for each value of α and each price vector, each consumer can afford a point in her consumption set, leading to a well defined excess demand correspondence.

In the HZ framework with an equal number of consumers and objects, Le (2018) considers commodity endowments that gives each consumer a positive probability of each object, and shows that if each consumer has a unique favorite object, then there is an equilibrium without any redistribution of income. This does not seem to be an easy consequence of Theorem 1: in Le's setting our result only proves the existence of an equilibrium with redistribution. It also seems unlikely that our methods could be adapted to prove the existence result of He, Miralles, Pycia, and Yan (2018), which uses a different sort of machinery to handle priorities.

Our framework allows the BCKM result to be generalized in several ways. First of all, by allowing a nontrivial production set, we can incorporate some flexibility concerning the package of seats offered by the school.

The pseudomarket mechanisms currently used to allocate course seats at some business schools generally provide each student with a budget of some artificial currency. This will not be guaranteed to be feasible unless each student is sure to be able to afford some feasible package of courses. In fact the use of such mechanisms currently seems to be restricted to settings in which the students are quite homogeneous, e.g., MBA students at the beginning of the program.

We can endow students with commodities in addition to artificial currency. In more complex settings it is quite unnatural for all students to have the same endowment, since they differ in seniority, program, and the courses that they have already taken. Our requirement that each student's endowment is greater than or equal to some element of her consumption set (instead of the origin as in BCKM) is quite natural, corresponding to each student's entitlement to some schedule of courses that fulfills the requirements of her program. Consumption sets that satisfy both floor and ceiling constraints are possible.

Our framework has additional flexibility that is potentially useful in course allocation. The share of excess purchasing power of sated consumers going to an unsated consumer is a parameter that can be chosen. Since a high share prioritizes the consumer's interests when many consumers are achieving satiation, it may (for example) be desirable to give recipients of academic merit scholarships higher shares, thereby making it less likely that one of them will be unsated in this circumstance. The ownership shares can also be chosen with various objectives in mind, and their significance may be altered by subtracting a vector from the production set while adding the same vector to the aggregate endowment.

From the point of view of course allocation, perhaps the most unfortunate assumption of Theorem 1 is that each X_i has a nonempty interior in $x_i^0 + V_0$, since it requires that any student is potentially capable of taking any course. In computational practice one may attempt to overcome this by passing from the actual consumption sets, which do not satisfy this condition, to larger consumption sets with nonempty interiors that have the given consumption sets as faces, but imposing a severe disutility of moving away from the given consumption sets.

There are a number of other problems with the application of pseudomarkets to course allocation that can be foreseen. While the various dimensions of flexibility described above are in one sense potentially useful, they lack the sorts of symmetry that make current business school course allocation mechanisms "obviously" fair. Perhaps they will come to be accepted because they produce good results that are not obviously unfair, but guidance concerning how best to tune their parameters, how to choose among multiple equilibria, and how to convince people that such choices are not unjust, is likely to come from practical experience rather than a priori reasoning.

Another potential problem is computational tractability. The proof of existence given here exploits a fixed point theorem. While algorithms for computing fixed points have achieved considerable practical success, there are theoretical results (e.g., Hirsch, Papadimitriou, and Vavasis (1989)) that suggest that such algorithms cannot provide any guarantee of rapid convergence. The limits of practical computation are unclear, and evolving.

4 The Proof of Theorem 1

We work in the setting laid out in Section 2. Let V denote a linear subspace of $\mathbb{R}^\ell$ that does not contain $\mathbf{e}$. In the analysis prior to the proof of Theorem 1 this space will contain net trades prior to disposal, and it will also be the space of possible price vectors. (The reader should be warned that although our continuity and hemicontinuity results are expressed in terms of the given economy, in the proof of Theorem 1 they will be applied to an expanded economy with $V = \mathbb{R} \times V_0$.) From this point to the end of the proof of Proposition 3 we assume that each consumption set X_i is contained in a translate of V, and that V contains all of the production sets.

For the time being we consider a fixed consumer i. Let $A_i^V = (\omega_i + V) \cap (X_i + C)$ be the set of points in $\omega_i + V$ that are feasible predisposal consumptions for i, in the sense that free disposal can reach a point in X_i. Of course A_i^V is closed, the minimum value of each component is given by X_i and the maximal value is not greater than the maximal value in X_i plus $\langle \omega_i - x_i^0, \mathbf{e} \rangle$, so A_i^V is compact.

Let U_i^V be the set of $p \in V \setminus \{0\}$ such that there is a point w_i in the interior (relative to $\omega_i + V$) of A_i^V such that $\langle p, w_i \rangle = \mu_i(p)$ and $w_i' \mapsto \langle p, w_i' \rangle$ is not constant on A_i^V. It is easy to show that U_i^V is an open subset of $V \setminus \{0\}$.

Define correspondences $B_i^V : U_i^V \to A_i^V$, $F_i^V : U_i^V \to X_i$, $G_i^V : U_i^V \to X_i$, and $D_i^V : U_i^V \to A_i^V$ by setting

$$B_i^V(p) = \{\, w_i \in A_i^V : \langle p, w_i \rangle = \mu_i(p) \,\}, \quad F_i^V(p) = X_i \cap (B_i^V(p) - C),$$

$$G_i^V(p) = \operatorname{argmax}_{x_i \in F_i^V(p)} u_i(x_i), \quad D_i^V(p) = \{\, w_i \in B_i^V(p) : G_i^V(p) \cap (w_i - C) \neq \emptyset \,\}.$$

In words, $B_i^V(p)$ is i's budget set in $\omega_i + V$ when she is compelled to spend all her income, $F_i^V(p)$ is the set of points in X_i that can be reached from points in $B_i^V(p)$ by disposal, $G_i^V(p)$ is the set of optimal points in $F_i^V(p)$, and $D_i^V(p)$ is the set of points in $\omega_i + V$ that can pass to points in $G_i^V(p)$ by disposal.

Our handling of the continuity properties of consumption depends on a technical trick, which is based on the following geometric fact. Recall that a *polyhedron* in $\mathbb{R}^\ell$ is an intersection of finitely many closed half spaces, and a *polytope* is a bounded polyhedron.

Proposition 1 *If P_1 and P_2 are polyhedra in $\mathbb{R}^\ell$, $Q = \{\, q \in \mathbb{R}^\ell = (P_1 + q) \cap P_2 \neq \emptyset \,\}$, and $I : Q \to \mathbb{R}^\ell$ is the correspondence $I(q) = (P_1 + q) \cap P_2$, then I is continuous.*

Appendix B contains the proof of this and Proposition 2 below. (The arguments have a prosaic point-set theoretic character.) The main point is:

Proposition 2 *If X_i is a polytope, then D_i^V is an upper hemicontinuous convex valued correspondence.*

In the course of proving this Appendix B shows that if X_i is a polytope, then B_i^V, F_i^V, and G_i^V are convex valued correspondences, B_i^V and F_i^V are continuous, and G_i^V is upper hemicontinuous.

Let $U^V = \bigcap_i U_i^V$. The excess demand set for $p \in U^V$ is the Minkowski sum

$$Z^V(p) = -\omega - \sum_j M_j(p) + \sum_i D_i^V(p).$$

We have $D_i^V(p) \subset \omega_i + V$ for all i and $M_j(p) \subset Y_j \subset V$ for all j, so $Z^V(p) \subset V$. Fix a good h^* such that V is not contained in $\{ p \in \mathbb{R}^\ell : p_{h^*} = 0 \}$. For $\varepsilon \in (0,1)$ let

$$S_\varepsilon = \{ p \in V : \|p\| = 1 \text{ and } p_{h^*} \geq \varepsilon \}.$$

As the intersection of a sphere with a half space, S_ε is a smooth manifold with boundary. The *tangent space* of S_ε at p is $T_pS_\varepsilon = \{ z \in V : \langle p, z \rangle = 0 \}$. As the proof below explains more formally, Walras' law is tantamount to $Z^V(p) \in T_pS_\varepsilon$.

Proposition 3 *If each X_i is a polytope, $\varepsilon > 0$, $S_\varepsilon \subset U^V$, and $z_{h^*} \geq 0$ for all $p \in S_\varepsilon$ such that $p_{h^*} = \varepsilon$ and all $z \in Z^V(p)$, then there is a $p \in S_\varepsilon$ such that $0 \in Z(p)$.*

Appendix A explains the notion of a vector field correspondence, the Poincaré-Hopf theorem, and related concepts that are employed in the following argument.

Proof. Lemma 1 and Proposition 2 imply that Z^V is an upper hemicontinuous convex valued correspondence. For $p \in S_\varepsilon$ and $z \in Z^V$ there are $w_i \in D_i^V(p)$ and $y_j \in M_j(p)$ such that $z = -\omega - \sum_j y_j + \sum_i w_i$, and

$$\langle p, z \rangle = -\langle p, \omega \rangle - \sum_j \langle p, y_j \rangle + \sum_i \langle p, w_i \rangle = 0$$

because $\langle p, w_i \rangle = \langle p, \omega_i \rangle + \sum_j \theta_{ij} \pi_j(p)$ and $\pi_j(p) = \langle p, y_j \rangle$. Thus Z^V is a vector field correspondence. By assumption it is not outward pointing. Therefore the Poincaré-Hopf theorem implies that the index of Z^V is $(-1)^{\dim V - 2}$ times the Euler characteristic of S_ε. Since S_ε is contractible, its Euler characteristic is 1. Thus the index is not zero, so $0 \in Z^V(p)$ for some p. ■

Proof of Theorem 1. We first expand the economy by adding a new good 0 that is not disposable, but which will be assumed to be universally desirable. For $\tilde{z} \in \mathbb{R}^{1+\ell}$, z denotes the projection that discards the first (good 0) component, so $\tilde{p} = (\tilde{p}_0, p)$, $\tilde{w}_i = (\tilde{w}_{i0}, w_i)$, $\tilde{x}_i = (\tilde{x}_{i0}, x_i)$, $\tilde{y}_j = (\tilde{y}_{j0}, y_j)$, etc. For $i = 1, \ldots, m$ let $\tilde{X}_i = [0, \tau_i] \times X_i$ where $\tau_i > 0$ is a number that is sufficiently large in a sense that will be specified later, and let $\tilde{u}_i : \tilde{X}_i \to \mathbb{R}$

be the function $\tilde{u}_i(\tilde{x}_i) = \tilde{x}_{i0} + u_i(x_i)$. For $j = 1, \ldots, n$ let $\tilde{Y}_j = \{0\} \times Y_j$, and define $\tilde{\pi}_j$ and $\tilde{M}_j$ in relation to $\tilde{Y}_j$ as we defined π_j and M_j in relation to Y_j.

We now define a sequence of economies. For each i let $\overline{X}_i = \mathrm{argmax}_{x_i \in X_i} u_i(x_i)$ be the set of bliss points. Since u_i is continuous and semi-strictly quasiconcave, $\overline{X}_i$ is nonempty, compact, and convex. Let $\{X_i^r\}$ be a sequence of polytopes contained in X_i such that $X_i^r \to X_i$ and $X_i^r \cap \overline{X}_i \to \overline{X}_i$ in the Hausdorff metric. After eliminating some of the initial terms of the sequence, x_i^0 is an interior point of X_i^r for all r. Let $\{\tilde{\omega}_0^r\}$ be a sequence in $(0,1)$ that converges to 0. Consumer i's endowment in the r^{th} economy is $\tilde{\omega}_i^r = (\alpha_i \tilde{\omega}_0^r, \omega_i)$. For each i and r define $\tilde{\mu}_i^r$ in relation to the production sets $\tilde{Y}_j$, the ownership shares θ_{ij}, and $\tilde{\omega}_i^r$, as we defined μ_i.

Let $\tilde{V} = \mathbb{R} \times V_0$. Define $\tilde{A}_i^{\tilde{V},r}$ and $\tilde{U}_i^{\tilde{V},r}$ in relation to $\tilde{V}$ and the r^{th} economy as we defined A_i^V and U_i^V above. Let $\tilde{U}^{\tilde{V},r} = \bigcap_i \tilde{U}_i^{\tilde{V},r}$. For each i and r define $\tilde{F}_i^{\tilde{V},r}$, $\tilde{G}_i^{\tilde{V},r}$, and $\tilde{D}_i^{\tilde{V},r}$ in relation to the r^{th} economy as we defined F_i^V, G_i^V, and D_i^V above. For $\tilde{p} \in \tilde{U}^{\tilde{V},r}$ let $\tilde{Z}^{\tilde{V},r}(\tilde{p}) = -\tilde{\omega}^r - \tilde{M}(\tilde{p}) + \sum_i \tilde{D}_i^{\tilde{V},r}(\tilde{p})$. For some $\varepsilon > 0$ let $\tilde{S}_\varepsilon = \{\, \tilde{p} \in \tilde{V} : \|\tilde{p}\| = 1 \text{ and } \tilde{p}_0 \geq \varepsilon \,\}$.

We seek conditions on ε and the τ_i that imply that the hypotheses of Proposition 3 hold. In particular, we want every $\tilde{p} \in \tilde{S}_\varepsilon$ to be in $\tilde{U}^{\tilde{V},r}$ because $0 < \alpha_i \tilde{\omega}_0^r + \sum_j \theta_{ij} \tilde{\pi}_j(\tilde{p})/\tilde{p}_0 < \tau_i$, so that $(\alpha_i \tilde{\omega}_0^r + \sum_j \theta_{ij} \tilde{\pi}_j(\tilde{p})/\tilde{p}_0, \omega_i)$ is in the interior of $\tilde{A}_i^{\tilde{V},r}$. It is clear that, once ε has been fixed, we can choose τ_i large enough that this will necessarily be the case. We also need total demand for good 0 to exceed supply at all $\tilde{p} \in \tilde{S}_\varepsilon$ such that $\tilde{p}_0 = \varepsilon$. Here we sketch the main ideas, omitting details. Consider a $\tilde{p} = (\varepsilon, p)$ in the boundary of S_ε. For each i there is a ball centered at x_i^0 that is contained in every X_i^r. It is possible to generate income to spend on good 0 by moving final consumption away from x_i^0 in the direction $-p$, and since $\|p\| = \sqrt{1 - \varepsilon^2}$ is close to 1, the amount of income that can be generated in this way is approximately equal to the radius of this ball. If ε is sufficiently small, the utility resulting from spending half of this income on good 0 exceeds $\max_{x_i, x_i' \in X_i} u_i(x_i) - u_i(x_i')$. In this case any bundle that spends less than half this amount on good zero is dominated by a bundle that spends this amount on good 0. Therefore, for large r, total demand for good 0 exceeds $\tilde{\omega}_0^r$.

Thus, if ε and the τ_i are chosen suitably, Proposition 3 implies that for each r there is a $\tilde{p}^r \in \tilde{S}_\varepsilon$ such that $0 \in \tilde{Z}^{\tilde{V},r}(\tilde{p}^r)$. For each r choose $\tilde{w}_i^r \in \tilde{D}_i^{\tilde{V},r}(\tilde{p}^r)$ for all i and $\tilde{y}_j^r \in \tilde{M}_j(\tilde{p}^r)$ for all j such that $\sum_i \tilde{w}_i^r = \tilde{\omega}^r + \sum_j \tilde{y}_j^r$. For each r and i let $\tilde{x}_i^r$ maximize $\tilde{u}_i(\tilde{x}_i^r)$ subject to $\tilde{x}_i^r \in \tilde{X}_i \cap (\tilde{w}_i^r - C)$. After passing to a subsequence we may assume that $w^r \to w$, $x^r \to x$, and $y^r \to y$.

For the usual reasons, multiplying all prices by a positive scalar does not change the consumers' budget sets or optimal consumptions. In addition, since each consumer i is constrained to choose a predisposal consumption in $\mathbb{R} \times (\{\omega_i\} + V_0)$, and production sets are contained in $\{0\} \times V_0$, adding a scalar multiple of $(0, \mathbf{e})$ to the price vector also has no effect. In more detail, suppose that $\hat{p}^r = \tilde{p}^r + \beta^r(0, \mathbf{e})$ for some $\beta^r \in \mathbb{R}$. For each j, $\tilde{Y}_j = \{0\} \times Y_j$, so $\tilde{M}_j(\hat{p}^r) = \tilde{M}_j(\tilde{p}^r)$ and $\tilde{\pi}_j(\hat{p}^r) = \tilde{\pi}_j(\tilde{p}^r)$. It follows that $\tilde{B}_i(\hat{p}^r) = \tilde{B}_i(\tilde{p}^r)$, $\tilde{F}_i(\hat{p}^r) = \tilde{F}_i(\tilde{p}^r)$, $\tilde{G}_i(\hat{p}^r) = \tilde{G}_i(\tilde{p}^r)$, and $\tilde{D}_i(\hat{p}^r) = \tilde{D}_i(\tilde{p}^r)$. We modify $\tilde{p}^r$ by adding the scalar multiple of $(0, \mathbf{e})$ that makes the minimum component of $(p^r_{\ell_c+1}, \ldots, p^r_\ell)$ zero. If, for the new $\tilde{p}^r$, $p^r = 0$, then each i is sated at x^r_i in X^r_i, and if this happens for arbitrarily large r, then each x_i is a bliss point in X_i, so $(0, x, y)$ is an EDE, and in fact it is (vacuously) an α-EDE for any α.

Therefore we may assume that for the new $\tilde{p}^r$, $p^r \neq 0$ for all r. We now multiply $\tilde{p}^r$ by the positive scalar that makes the maximum absolute value of a component of p^r one. After passing to a subsequence we may assume that p^r converges to a p such that $\min_{\ell_c+1 \le h \le \ell} p_h = 0$ and the maximum absolute value of any component is one. We will show that (p, x, y) is an α-EDE.

We begin by showing that (a) holds. First suppose that for some i there is an $x'_i \in X_i$ such that $\langle p, x'_i \rangle < \langle p, x_i \rangle$ and $u_i(x'_i) > u_i(x_i)$. For each r let $w'_i = x'_i + (w_i - x_i)$. Then $w'_i \in A^V_i$, and for large r we have $u_i(x'_i) > u_i(x^r_i)$ and $\langle p^r, w'_i \rangle \le \langle p^r, w^r_i \rangle$, which contradicts utility maximization for $\tilde{x}^r_i$, so this is not possible.

Next suppose that for some i there is a $x'_i \in X_i$ such that $\langle p, x'_i \rangle < \langle p, x_i \rangle$ and $u_i(x'_i) = u_i(x_i)$. If i is unsated at x_i, then moving x'_i in the direction of a bliss point yields $\langle p, x'_i \rangle < \langle p, x_i \rangle$ and $u_i(x'_i) > u_i(x_i)$, which we ruled out above. On the other hand if i is sated at x_i, then x_i and x'_i are bliss points for i. For large r we have $\langle p^r, x'_i \rangle < \langle p^r, x^r_i \rangle$, so we can improve on x^r_i by purchasing x'_i instead and spending more on good 0.

Finally suppose that for some i there is a $x'_i \in X_i$ such that $\langle p, x'_i \rangle = \langle p, x_i \rangle$ and $u_i(x'_i) > u_i(x_i)$. If there was any $x''_i \in X_i$ with $\langle p, x''_i \rangle < \langle p, x'_i \rangle$, then we could move x'_i in the direction of x''_i and achieve $\langle p, x'_i \rangle < \langle p, x_i \rangle$ and $u_i(x'_i) > u_i(x_i)$, which we know is impossible. Therefore $\langle p, x_i \rangle \le \langle p, x''_i \rangle$ for all $x''_i \in X_i$. In the expanded economies it is always possible to consume $(0, x^0_i)$, so $\langle p^r, x^0_i \rangle \le \langle p^r, x^r_i \rangle + \tilde{p}^r_0 \tilde{x}^r_{i0}$ for all r, and we are assuming that i is unsated at x_i, so $\tilde{p}^r_0 \tilde{x}^r_{i0} \to 0$ and thus $\langle p, x^0_i \rangle \le \langle p, x_i \rangle$. Since x^0_i is in the interior of X_i, this implies that $\langle p, x_i \rangle$ is constant on X_i, and that p is a scalar multiple of $\mathbf{e}$, which we have been assuming is not the case. This completes the verification of (a).

Since $\tilde{y}_j^r = (0, y_j^r) \in \tilde{M}_j(\tilde{p}^r)$ for all j and r, $y_j^r \in M_j(p^r)$, and Lemma 1 gives $y_j \in M_j(p)$, which is (c) of the definition of an EDE. Of course

$$\omega + \sum_j y_j - \sum_i x_i = \lim_r \left(\omega + \sum_j y_j^r - \sum_i x_i^r\right) \in C$$

because C is closed, so (d) holds. If, for some $h = \ell_c + 1, \ldots, \ell$, $\sum_i x_{ih} < \omega_h + \sum_j y_{jh}$, then $\sum_i x_{ih}^r < \omega_h + \sum_j y_{jh}^r$ for large r, and for some i we have $x_{ih}^r < w_{ih}^r$. This is inconsistent with utility maximization if there is an $h' = \ell_c + 1, \ldots, \ell$ such that $p_{h'}^r < p_h^r$ because i could modify w_i^r by purchasing more h' and less h, thereby obtaining the same element of X_i and a greater consumption of good 0. Thus $p_h^r = 0$ for large r, so $p_h = 0$, and (e) holds.

Since $x_i \leq w_i$ for all i and $p_h = 0$ if $\sum_i x_{ih} < \sum_i w_{ih}$, we have $p_h = 0$ whenever $x_{ih} < w_{ih}$, so $\langle p, x_i \rangle = \langle p, w_i \rangle$. For all i and r we have $\langle p^r, w_i^r \rangle + \tilde{p}_0^r \tilde{w}_{i0}^r = \tilde{\mu}_i^r(\tilde{p}^r) = \alpha_i \tilde{p}_0^r \tilde{\omega}_0^r + \mu_i(p_i^r)$. Let S (U) is the set of i that are sated (unsated) at x_i. The total benefit of consuming good 0 is bounded by $\tilde{\omega}_0^r$, so it goes to zero. If $i \in U$, then the amount spent on good 0 cannot result in a reduction, in the limit, of utility from other goods, so $\tilde{p}_0^r \tilde{w}_{i0}^r \to 0$. Therefore

$$\langle p, x_i \rangle - \mu_i(p_i) = \langle p, w_i \rangle - \mu_i(p_i) = \alpha_i \Pi$$

for $i \in U$, where $\Pi = \lim_r \tilde{p}_0^r \tilde{\omega}_0^r$. Since $\Pi \geq 0$, (b) holds, so (p, x, y) is an α-EDE. ■

5 Challenges

For the most part the traditional concerns of general equilibrium theory are meaningful and conceptually pertinent in relation to pseudomarkets, perhaps taking on a somewhat different flavor insofar as we expect these to be designed mechanisms rather than phenomena occurring "in nature." Thus one can easily produce a host of original and meaningful problems for further research. For example, extending the result to various infinite dimensional commodity spaces is one obvious direction of possible generalization.

In the remainder we briefly describe two, possibly quite challenging, open problems that have a fundamental character. Both can be stated in a particularly simple version of the HZ model. (It will be evident that many variations are possible.) Assume that there are the same number of consumers and objects, and each object has unit capacity. Thus $m = \ell$, and for each i we have $X_i = \Delta$ where $\Delta = \{ x \in \mathbb{R}^\ell_+ : \sum_h x_h = 1 \}$ is the unit simplex. The (expected) utilities are linear: $u_i(x_i) = \sum_h u_{ih} x_{ih}$. For each i let the endowment ω_i be the barycenter $(1/\ell, \ldots, 1/\ell)$ of Δ; equivalently, the consumers are endowed with equal incomes. Let $\mathbf{e} = (1, \ldots, 1) \in \mathbb{R}^\ell$, and let $V = \{ p \in \mathbb{R}^\ell : \langle p, \mathbf{e} \rangle = 0 \}$ be its orthogonal complement.

Experience with general equilibrium theory (Debreu (1970)) and strategic form (Harsanyi (1973)) and extensive form (Kreps and Wilson (1982)) game theory, leads one to conjecture that for generic utilities there are finitely many EDE. We analyze this, arriving at a sharper characterization.

Generically each consumer has a unique favorite object, in which case any EDE decomposes as an assignment of favorites to some consumers and an EDE for the remaining economy in which all consumers are unsated and consequently spend all of their income. Thus it suffices to show that for generic u there are finitely many EDE in which each consumer is unsated and consequently consumes a bundle with the same value as the endowment.

The aggregate endowment is $\nabla = \sum_h \delta_h$ where each δ_h is the element of Δ whose h coordinate is 1 and whose other coordinates are 0. As before let $X = \prod_i X_i = \Delta^\ell$. The space of possible equilibrium price-allocation pairs is

$$B = \{ (p, x) \in V \times X : \langle p, x_i \rangle = 0 \text{ for all } i \text{ and } \sum_i x_i = \nabla \}.$$

For $p \in V$ let $A(p) = \{ z \in \Delta : \langle p, z \rangle \le 0 \}$ be the common budget set. The graph of the equilibrium correspondence is

$$E = \{ (u, (p, x)) \in B \times (\mathbb{R}^\ell)^\ell : x_i \in \operatorname{argmax}_{x_i' \in A(p)} u_i(x_i') \text{ for all } i \}.$$

Let $\pi : E \to (\mathbb{R}^\ell)^\ell$ be the projection $(u, (p, x)) \mapsto u$. We would like to show that for a generic set $G \subset (\mathbb{R}^\ell)^\ell$, $\pi^{-1}(u)$ is finite for all $u \in G$, so that $u \in G$ have finitely many EDE in which each consumer's consumption has value 0.

We partition E according to which consumers consume positive quantities of which goods. For $\emptyset \neq \sigma \subset \{1, \ldots, \ell\}$ let the interior of the face of Δ spanned by σ be

$$\Delta_\sigma = \{ x \in \Delta : \text{for each } h,\ x_h > 0 \text{ if and only if } h \in \sigma \}.$$

For a profile $\Sigma = (\sigma_1, \ldots, \sigma_\ell)$ in which each σ_i is a nonempty subset of $\{1, \ldots, \ell\}$, let

$$B_\Sigma = \{ (p, x) \in B : x_i \in \Delta_{\sigma_i} \text{ for all } i \} \quad \text{and} \quad E_\Sigma = \{ (u, (p, x)) \in E : (p, x) \in B_\sigma \}.$$

For each $(p, x) \in B_\Sigma$, the dimension of the set of $u_i \in \mathbb{R}^\ell$ such that $x_i \in \text{argmax}_{p \in B(p)} u_i(x)$ is a closed set whose interior is an open subset of an $(\min\{\ell, \ell - |\sigma_i| + 2\})$-dimensional linear subspace. Thus the set of $u \in (\mathbb{R}^\ell)^\ell$ such that $(u, (p, x)) \in E$ is a $(\sum_i \min\{\ell, \ell - |\sigma_i| + 2\})$-dimensional linear subspace.

A *semi-algebraic set* S in a Euclidean space is a finite union of sets, each of which is defined by a finite conjunction of polynomial equations and inequalities. A result first proved by Whitney (1957) (cf. Bochnak, Coste, and Roy (1998) and Benedetti and Risler (1990)) implies that such an S has a finite partition $M_1, \ldots, M_K$ whose cells are C^∞ manifolds of various dimensions. The *dimension* of S is the maximal dimension of the M_j. (This is the same for all partitions). It is easy to see that B_Σ is a semi-algebraic set. Let $M_1, \ldots, M_K$ be a partition of B_Σ into C^∞ manifolds. If the dimension of B_Σ is not greater than $\sum_i \max\{|\sigma_i| - 2, 0\}$, then the dimension of $\{\, (u, (p, x)) \in E : (p, x) \in M_j \,\}$ is not greater than ℓ^2, and application of Sard's theorem to the restriction of π to this set shows that generic u have finitely many equilibria in M_k. Conversely, if the dimension of B_Σ is greater than $\sum_i \max\{|\sigma_i| - 2, 0\}$, then there is an open subset of $(\mathbb{R}^\ell)^\ell$ whose elements each have a continuum of equilibria.

We now discuss some issues related to computation. For many applications of pseudo-markets, decentralized trading processes will not yield equilibria of the desired sort because they do not incentivize sated consumers to minimize expenditure. Thus implementation must be done by computing an equilibrium using preferences submitted by the consumers, so the complexity of this computation is a concern. (As many computer scientists have pointed out, market equilibration cannot perform intractible computations, so complexity concerns also give rise to potential doubts concerning the realism of general equilibrium as a model of decentralized trading in complex economies.)

The Debreu-Mantel-Sonnenschein theorem implies that in its general form, the problem of computing an approximate competitive equilibrium is equivalent to the problem of computing an approximate fixed point in the setting of Brouwer's fixed point theorem, and various algorithms (e.g., Ch. 3 of McLennan (2018)) for this problem are known, and may be used to compute EDE's of the expanded economies arising in the proof of Theorem 1. Thus it is possible to compute at least an approximation of an EDE, but this is far from ideal. A representation of an EDE as a fixed point would allow more direct computational methods to be employed. I do not know of such a representation, nor do I know of any way to show that no such representation exists.

Computation of market equilibrium under restrictive hypotheses has been studied extensively in the computer science literature. (Codenotti, Pemmaraju, and Varadarajan (2004) surveys less recent literature.) Of particular interest to us are Fisher markets with linear utilities. In a Fisher market there are finitely many divisible goods in fixed supply and finitely many consumers, each of whom has a utility function and a positive endowment of money. An equilibrium is a vector of nonnegative prices and an allocation such that each consumer is maximizing utility subject to her budget constraint. Gale (1960) showed that when utilities are linear, the set of equilibria is the set of solutions of a convex program, and Eisenberg (1961) generalized this result to homogeneous utility functions. Devanur, Papadimitriou, Saberi, and Vazirani (2008) provide a polynomial time algorithm for finding an equilibrium when utilities are linear. Jain and Vazarani (2010) define a class of Eisenberg-Gale markets, and provide polynomial time algorithms for several other markets in this class.

The difference between the linear Fisher market and our HZ market is that in the HZ market each consumer has the additional constraint that the sum of the quantities purchased (i.e., total probability) must be one. McLennan and Takayama (2018) present examples with multiple isolated equilibria, so the problem is not equivalent to a convex program, and they present calculations suggesting that even in the simple version of the HZ model considered here, multiplicity of equilibria is not uncommon. This gives the subjective impression that the problem of finding an EDE has the character of a fully general fixed point problem, which suggests that the corresponding computational problem has similar complexity.

The computational class PPAD is the class of computational problems with the following description. A *directed graph* is a pair (V, E) consisting of a finite set V of *nodes* and a finite set $E \subset V \times V$ of *directed edges*. If $(v, w) \in E$, then v is a *predecessor* of w and w is a *successor* of v. The *indegree* (*outdegree*) of a node is the number of predecessors (successors). A node is a *source* if its indegree is zero and a *sink* if its outdegree is zero. An instance of PPAD has two inputs, the first of which is a Turing machine that defines a directed graph of maximal indegree one and maximal outdegree one, because it takes a node in V as input and computes the predecessor (if it exists) and the successor (if it exists). The other input is a source, and the problem is to find a sink or some other source. The class PPAD abstracts the common features of the Lemke-Howson algorithm for 2-Nash, which is the computational problem of finding a Nash equilibrium of a finite two person strategic form game, as well as various algorithms for computing approximate fixed points. (Cf. Chapter 3 of McLennan (2018).)

A major theoretical development was accomplished by Daskalakis, Goldberg, and Papadimitriou (2006) and Chen and Deng (2006), who proved that 2-Nash is complete for PPAD. Concretely, there is a polynomial time algorithm that takes any problem in PPAD as input and outputs a two person game, and there is another polynomial time algorithm that has the given problem and a Nash equilibrium of this two person game as input, and outputs a solution of the given problem. Thus a polynomial time 2-Nash solver could be turned into a polynomial time algorithm for arbitrary problems in PPAD. It is believed (conjecturally, insofar as it requires at least that $\mathbf{P} \neq \mathbf{NP}$) that there is no polynomial time algorithm for the general problem of computing an approximate fixed point, so this result is compelling evidence that there is no polynomial time algorithm for 2-Nash.

This method of showing that a particular problem is hard because an algorithm for it could be used to solve a problem that is known or believed to be hard is called *reduction*. Subsequently many problems in theoretical economics have been shown to be complete for PPAD, because any instance of 2-Nash, or some other problem that has already been shown to be PPAD complete, can be reduced to an instance of the problem in question. In particular, various seemingly quite elementary versions of general equilibrium have been shown to be PPAD-complete. (E.g., Chen, Dai, Du, and Teng (2009).)

If the problem of finding an equilibrium of the HZ model is in fact PPAD complete, application of general fixed point solvers (as is already outlined in Appendix B of HZA) is likely be as efficient as any method for the general problem. The history of the computational general equilibrium literature suggests looking for additional restrictions on the problem that are practically relevant and allow more efficient algorithms to be employed.

Appendix A: The Vector Field Index

We briefly review the theory of the vector field index (Ch. 15 of McLennan (2018)). Let $M \subset \mathbb{R}^k$ be a smooth n-dimensional manifold with boundary. (Concretely, M is covered by open sets U that are C^∞ diffeomorphic to open subsets of the half space $H_1 = \{ x \in$

$\mathbb{R}^n : x_1 \geq 0\,\}$.) If $C \subset M$ is compact, a *vector field correspondence* Z with domain C is an assignment of a nonempty set $Z(p) \subset T_pM$ of tangent vectors to each $p \in C$. We do not make a formal distinction between a singleton valued correspondence and a function, and if Z is a continuous function we say that Z is a *vector field*.

An *equilibrium* of Z is a $p \in C$ such that $0 \in Z(p)$. We say that Z is *index admissible* is it is upper hemicontinuous, contractible[2] valued, and has no equilibria in the topological boundary $\partial C = C \cap \overline{M \setminus C}$ of C. The *vector field index* assigns an integer $\text{ind}(Z)$ to each index admissible vector field correspondence Z. It is completely characterized by the following properties:

- (Normalization) If z is a smooth vector field on C that has a single equilibrium $p \in C \backslash \partial C$, and $Dz(p) : T_pM \to T_pM$ is nonsingular[3], then $\text{ind}(z) = +1$ if $|Dz(p)| > 0$ and $\text{ind}(z) = -1$ if $|Dz(p)| < 0$, where $|Dz(p)|$ is the determinant of $Dz(p)$.

- (Additivity) If Z with domain C is an index admissible vector field correspondence, $C_1, \ldots, C_r$ are pairwise disjoint compact subsets of C, and Z has no equilibria in the closure of $C \setminus (C_1 \cup \cdots \cup C_r)$, then $\text{ind}(Z) = \sum_k \text{ind}(Z|_{C_k})$.

- (Continuity) If Z is an index admissible vector field correspondence with domain C, then there is a neighborhood U of the graph of Z such that if Z' is an index admissible vector field correspondence with domain C whose graph is contained in U, then $\text{ind}(Z') = \text{ind}(Z)$.

Note that Additivity implies that $\text{ind}(Z) = 0$ if Z has no equilibria, so that equilibria necessarily exist if $\text{ind}(Z) \neq 0$. Additivity allows us to unambiguously define the *index* of an isolated equilibrium of Z to be the index of the restriction of Z to any compact neighborhood of the equilibrium that contains no other equilibria.

Consider a point $p \in M$ and a C^∞ coordinate chart $\varphi : U \to V$ where U is an open subset of M containing p and $V \subset H_1$ is open. We say that $v \in T_pM$ is *not outward pointing* (*inward pointing*) if $v_1 \geq 0$ ($v_1 > 0$). We say that v is *not inward pointing* (*outward pointing*)

[2]A topological space X is *contractible* if there is a continuous function $c : X \times [0,1] \to X$ such that $c(\cdot, 0)$ is the identity function on X and $c(\cdot, 1)$ is a constant function.

[3]One way to see that the image of $Dz(p)$ is contained in T_pM is to consider that if $\varepsilon : (-\varepsilon, \varepsilon) \to M$ and $\nu : (-\varepsilon, \varepsilon) \to \mathbb{R}^k$ are smooth functions such that $\nu(t) \perp T_{\varepsilon(t)}M$ for all t, then differentiating $0 = \langle z(\varepsilon(t)), \nu(t)\rangle$ gives $0 = \langle Dz(\varepsilon(t))\varepsilon'(t), \nu(t)\rangle + \langle z(\varepsilon(t)), \nu'(t)\rangle$ and thus $0 = \langle Dz(p)\varepsilon'(0), \nu(0)\rangle$.

if $-v$ is not outward pointing (inward pointing). A vector field correspondence Z on M is *not outward pointing* (*inward pointing*, *not inward pointing*, *outward pointing*) if, for all $p \in M$ and $z \in Z(p)$, z is not outward pointing (inward pointing, not inward pointing, outward pointing). The Poincaré-Hopf theorem (e.g., Milnor (1965)) asserts that the index of a continuous outward pointing vector field on a compact C^∞ manifold with boundary is the manifold's Euler characteristic. Normalization and the fact that the properties above characterize the index imply that the index of the negation of a vector field is the vector field's index multiplied by $(-1)^{\dim q}$. Theorem 15.11 of McLennan (2018) generalizes the Poincaré-Hopf theorem to C^∞ manifolds with corners, and to index admissible vector field correspondences that are not outward pointing, but which need not be inward pointing.

Appendix B: Omitted Proofs

Proof of Proposition 1. By construction I is nonempty valued, and it is easy to see that it is upper hemicontinuous. Since everything is translation invariant it suffices to prove that I is lower hemicontinuous at the origin when the origin is an element of Q, and an element of $I(0)$, which is to say that $0 \in P_1 \cap P_2$. Concretely, for a given sequence $\{q^r\}$ in Q converging to 0 we need to produce a sequence $\{x^r\}$ with $x^r \in I(q^r)$ for all r and $x^r \to 0$.

Let C_1 and C_2 be the cones given by the inequalities defining P_1 and P_2 that hold with equality at 0. If $q \in Q$, then $(C_1 + q) \cap C_2 \neq \emptyset$. Moreover, there is a constant $A > 0$ such that if $q \in Q$ and x is the point in $(C_1 + q) \cap C_2$ that is closest to the origin, then $\|x\| \leq A\|q\|$. If $q \in Q$ and $\|q\|$ is sufficiently small, then the point in $(C_1 + q) \cap C_2$ that is closest to the origin is an element of $(P_1 + q) \cap P_2$. For sufficiently large r we let x^r be the point in $(C_1 + q^r) \cap C_2$ that is closest to the origin. ■

In preparation for the proof of Proposition 2 we establish the relevant continuity properties of B_i^V, F_i^V, and G_i^V.

Lemma 2 *B_i^V is a continuous convex valued correspondence.*

Proof. The definition of U_i^V implies that $B_i^V(p)$ is nonempty whenever $p \in U_i^V$, and it is obviously convex. Evidently $w_i \in B_i^V(p)$ whenever $\{p^r\}$ is a sequence in V converging to $p \neq 0$, $w_i^r \in B_i^V(p^r)$ for each r, and $w_i^r \to w_i$, so the correspondence is upper hemicontinuous.

It remains to demonstrate lower hemicontinuity. Suppose that $\{p^r\}$ is a sequence in U_i^V converging to $p \in U_i^V$ and $w_i \in B_i^V(p)$. Fixing a convex neighborhood $U \subset \omega_i + V$ of w_i, it suffices to show that U has a nonempty intersection with $B_i^V(p^r)$ for sufficiently large r. Since $p \in U_i^V$, there is a point $w_i'' \in B_i^V(p)$ in the interior of A_i^V. Let w_i' be a point in the interior of the line segment between w_i and w_i'' that is contained in U. Then $w_i' \in B_i^V(p)$, and, because A_i^V is convex, W_i' is in the interior of A_i^V. Since $\langle p, \cdot \rangle$ is not constant on A_i^V there are w_i^0, w_i^1 in the interior of $A_i^V \cap U$ such that $\langle p, w_i^0 \rangle < \mu_i(p) < \langle p, w_i^1 \rangle$. By continuity $\langle p^r, w_i^0 \rangle < \mu_i(p^r) < \langle p^r, w_i^1 \rangle$ for large r, in which case some convex combination of w_i^0 and w_i^1 is an element of $B_i^V(p^r)$ in $A_i^V \cap U$. ■

Lemma 3 *If X_i is a polytope, then F_i^V is a continuous convex valued correspondence.*

Proof. The definition of U_i^V implies that F_i^V is nonempty valued. The proof of upper hemicontinuity is as usual: the definition is a matter of weak inequalities of continuous functions. To prove that F_i^V is lower hemicontinuous fix $p \in U_i^V$ and $x_i \in F_i(p)$, and suppose that $\{p^r\}$ is a sequence in U_i^V converging to p. Choose $w_i \in B_i^V(p)$ such that $x_i \in w_i - C$. Since B_i^V is lower hemicontinuous there is a sequence $\{w_i^r\}$ with $w_i^r \in B_i^V(p^r)$ and $w_i^r \to w_i$. Proposition 1 implies that there is a sequence $\{x_i^r\}$ with $x_i^r \in X_i \cap (w_i^r - C)$ and $x_i^r \to x_i$, and the definition of F_i^V implies that $X_i \cap (w_i^r - C) \subset F_i^V(p^r)$. ■

Lemma 4 *If X_i is a polytope, then G_i^V is an upper hemicontinuous convex valued correspondence.*

Proof. In view of the last result this is a consequence of Berge's theorem. ■

Proof of Proposition 2. By definition D_i^V is nonempty valued. To show that D_i^V is convex valued fix $p \in U_i^V$, and suppose that $w_i, w_i' \in D_i^V(p)$ and $0 \le t \le 1$. There are $x_i, x_i' \in X_i$ with $x_i \le w_i$, $x_i' \le w_i'$, and $u_i(x_i) = u_i(x_i')$. Of course $(1-t)w_i + tw_i' \in B_i^V(p)$, $(1-t)x_i + tx_i' \le (1-t)w_i + tw_i'$, and $u_i((1-t)x_i + tx_i') \ge (1-t)u_i(x_i) + tu_i(x_i')$, but x_i and x_i' are maximizers, so this inequality holds with equality. Thus $(1-t)w_i + tw_i' \in D_i^V(p)$.

To prove upper hemicontinuity suppose that $\{p^r\}$ is a sequence in U_i^V converging to $p \in U_i^V$, $w_i^r \in D_i^V(p^r)$ for all r, and $w_i^r \to w_i$. For each r choose $x_i^r \in X_i$ such that $w_i^r - x_i^r \in C$ and $u_i(x_i^r) = \max_{x_i \in F_i^V(p^r)} u_i(x_i)$. Passing to a subsequence, we may assume that $x_i^r \to x_i$. Then $w_i - x_i \in C$ (since C is closed) and $x_i \in G_i^V(p)$ because G_i^V is upper hemicontinuous, so $w_i \in D_i^V(p)$. ■

References

Abdulkadiroğlu, A., Y.-K. Che, and Y. Yasuda (2011): "Resolving Conflicting Preferences in School Choice: The 'Boston Mechanism' Reconsidered," *American Economic Review*, 101, 399–410.

——— (2015): "Expanding 'Choice' in School Choice," *American Economic Journal: Microeconomics*, 7(1), 1–42.

Ashlagi, I., and P. Shi (2016): "Optimal Allocation Without Money: An Engineering Approach," *Management Science*, 62, 1078–97.

Benedetti, R., and J.-J. Risler (1990): *Real Algebraic and Semi-algebraic Sets*. Hermann, Paris.

Bergstrom, T. C. (1976): "How to discard 'free disposability' at no cost," *J. Math. Econom.*, 3, 131–134.

Bochnak, J., M. Coste, and M.-F. Roy (1998): *Real Algebraic Geometry*. Springer, New York.

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

Bogomolnaia, A., H. Moulin, F. Sandomirskiy, and E. Yanovskaya (2016a): "Dividing Goods and Bads Under Additive Utilities," arXiv:1610.03745 [cs.GT].

——— (2016b): "Dividing Goods or Bads Under Additive Utilities," arXiv:1608.01540v2 [cs.GT].

——— (2017): "Competitive Division of a Mixed Manna," *Econometrica*, 85, 1847–1871.

Budish, E. (2011): "The Combinatorial Assignment Problem: Approximate Competitive Equilibrium from Equal Incomes," *J. Polit. Economy*, 119, 1061–1103.

Budish, E., and E. Cantillon (2012): "The Multi-unit Assignment Problem: Approximate Theory and Evidence from Course Allocation at Harvard," *Amer. Econ. Rev.*, 102, 2237–2271.

Budish, E., Y.-K. Che, F. Kojima, and P. Milgrom (2013): "Designing Random Allocation Mechanisms: Theory and Practice," *Amer. Econ. Rev.*, 103, 585–623.

Budish, E., and J. B. Kessler (2016): "Bringing Real Market Participants' Real Preferences into the Lab: An Experiment that Changed the Course Allocation Mechanism at Wharton," working paper, Booth School, University of Chicago.

Chen, X., D. Dai, Y. Du, and S.-H. Teng (2009): "Settling the Complexity of Arrow-Debreu Equilibria in Markets with Additively Separable Utilities," in *Proceedings of the 2009 50th Annual IEEE Symposium on Foundations of Computer Science*, pp. 273–282.

Chen, X., and X. Deng (2006): "Settling the Complexity of Two-Player Nash Equilibrium," in *Proceedings of the 47th Annual IEEE Symposium on Foundations of Computer Science*, pp. 261–272.

Codenotti, B., S. Pemmaraju, and K. Varadarajan (2004): "Algorithms column: The computation of market equilibrium," *ACM SIGACT News*, 35.

Daskalakis, C., P. Goldberg, and C. Papadimitriou (2006): "The Complexity of Computing a Nash Equilibrium," in *Proceedings of the 38th ACM Symposium on the Theory of Computing*.

Debreu, G. (1970): "Economies with a Finite Set of Equilibria," *Econometrica*, 38, 387–392.

Devanur, N., C. Papadimitriou, A. Saberi, and V. Vazirani (2008): "Market Equilibrium via a Primal-Dual Algorithm for a Convex Program," *Journal of the ACM*, 55, 22:1–22.

Echenique, F., A. Miralles, and J. Zhang (2018): "Fairness and Efficiency for Probabilistic Allocations with Endowments," Working paper, Caltech.

Eisenberg, E. (1961): "Aggregation of Utility Functions," *Management Science*, 7, 337–350.

Eisenberg, E., and D. Gale (1959): "Consensus of Subjective Probabilities: The Pari-mutuel Method," *Annals of Mathematical Statistics*, 30, 165–168.

Featherstone, C., and M. Niederle (2008): “Ex Ante Efficiency in School Choice Mechanisms: An Experimental Investigation,” National Bureau of Economic Research (NBER) Working Paper 14618.

Gale, D. (1960): *Theory of Linear Economic Models*. McGraw Hill, New York.

Gale, D., and A. Mas-Colell (1975): “An Equilibrium Existence Theorem for a General Model Without Ordered Preferences,” *J. Math. Econom.*, 2, 9–15.

——— (1979): “Correction to an Equilibrium Existence Theorem for a General Model Without Ordered Preferences,” *J. Math. Econom.*, 6, 297–298.

Harsanyi, J. C. (1973): “Oddness of the Number of Equilibrium Points: A New Proof,” *Int. J. Game Theory*, 2, 235–250.

Hart, O., and H. Kuhn (1975): “A Proof of the Existence of Equilibrium Without the Free Disposal Assumption,” *J. Math. Econom.*, 2, 335–343.

He, Y., A. Miralles, M. Pycia, and J. Yan (2018): “A Pseudo-Market Approach to Allocation with Priorities,” *Amer. Econ. J.: Microeconomics*, 10, 272–314.

Hirsch, M., C. Papadimitriou, and S. Vavasis (1989): “Exponential lower bounds for finding Brouwer fixed points,” *J. Complexity*, 5, 379–416.

Hylland, A., and R. Zeckhauser (1979a): “The Efficient Allocation of Individuals to Positions,” *J. Polit. Economy*, 87, 293–314.

——— (1979b): “‘The Efficient Allocation of Individuals to Positions’: Appendices,” Harvard University.

Immorlica, N., B. Lucier, J. Mollner, and E. G. Weyl (2017): “Raffles,” working paper, Microsoft Research.

Jain, K., and V. V. Vazarani (2010): “Eisenberg-Gale Markets: Algorithms and Game-Theoretic Properties,” *Games and Economic Behavior*, 70, 84–106.

Kreps, D., and R. Wilson (1982): “Sequential Equilibrium,” *Econometrica*, 50, 863–894.

Le, P. (2018): "Competitive Equilibrium in the Random Assignment Problem," *Int. J. Econ. Theory*, 13, 369–385.

Liu, Q., and M. Pycia (2008): "Ordinal Efficiency, Fairness, and Incentives in Large Markets," mimeo, UCLA.

Mas-Colell, A. (1992): "Equilibrium Theory with Possibly Sated Preferences," in *Equilibrium and Dynamics: Essays in Honor of David Gale*, ed. by M. Majumdar, pp. 201–213. St. Martin's Press, New York.

McLennan, A. (2018): *Advanced Fixed Point Theory for Economics.* Springer, New York.

McLennan, A., and S. Takayama (2018): "Multiple Equilibria in the Hylland and Zeckhauser Model," in preparation, University of Queensland.

Milnor, J. (1965): *Topology from the Differentiable Viewpoint.* University Press of Virginia, Charlottesville.

Miralles, A. (2009): "School Choice: The Case for the Boston Mechanism," in *Auctions, Market Mechanisms and Their Applications*, ed. by S. Das, M. Ostrovsky, D. Pennock, and B. Szymanski, pp. 58–60. Springer, New York.

Nash, J. F. (1950): "The Bargaining Problem," *Econometrica*, 18, 155–162.

Polemarchakis, H., and P. Siconolfi (1993): "Competitive Equilibria without Free Disposal or Nonsatiation," *J. Math. Econom.*, 22, 85–99.

Pratt, J. W. (2007): "Fair (and Not So Fair) Division," *J. Risk and Uncertainty*, 35, 203–236.

Pratt, J. W., and R. J. Zeckhauser (1990): "The Fair and Efficient Division of the Winsor Family Silver," *Management Science*, 36, 1293–1301.

Pycia, M. (2014): "The Cost of Ordinality," mimeo, UCLA.

Roberts, J., and A. Postlewaite (1976): "The Incentives for Price-Taking Behavior in Large Exchange Economies," *Econometrica*, 44(1), 115–127.

Sönmez, T., and M. U. Ünver (2010): "Course Bidding at Business Schools," *Int. Econ. Rev.*, 51, 99–123.

Troyan, P. (2012): "Comparing School Choice Mechanisms by Interim and Ex-ante Welfare," *Games and Economic Behavior*, 75, 936–47.

Varian, H. R. (1974): "Equity, Envy, and Efficiency," *J. Econom. Theory*, 9, 63–91.

Weller, D. (1985): "Fair Division of a Measurable Space," *J. Math. Econom.*, 14, 5–17.

Whitney, H. (1957): "Elementary Structure of Real Algebraic Varieties," *Ann. of Math.*, 66, 545–556.