The Combinatorial Assignment Problem: Approximate Competitive Equilibrium from Equal Incomes

Eric Budish
Harvard University

Feb 2009

The Combinatorial Assignment Problem

General question: How can we divide a set of indivisible objects amongst a set of agents without using monetary transfers, in a way that is efficient, incentive compatible, and fair?

Specific instance: Course Allocation at Universities

- The indivisible objects are seats in courses
- Each student requires a bundle of courses
- Preferences are private information
- Exogenous restriction against monetary transfers (even at Chicago!)

Another important example is assigning interchangeable workers to tasks or shifts

The Combinatorial Assignment Problem

Relation to the Literature

Combinatorial assignment is one feature removed from problems that have received considerable attention and have compelling solutions

- No restriction on money → Combinatorial Auction Problem (Vickrey 1961; Milgrom 2004; Cramton et al 2006)
- Two-Sided Preferences → Matching Problem (Gale Shapley 1961; Roth 1984; Roth Sotomayor 1990; Roth Peranson 1999; Hatfield Milgrom 2005)
- Single-Unit Demand → School Assignment Problem (Shapley Scarf 1974; Hylland Zeckhauser 1979; Roth 1982; Bogomolnaia Moulin 2001; Abdulkadiroglu et al 2006; Pathak 2006)
- Divisible Goods → Classic Fair Division problem (Steinhaus 1948; Dubins and Spanier 1961; Foley 1967; Varian 1974; Moulin 1995; Brams and Taylor 1996)

Why has Progress on this Problem been Elusive?

Difficulty #1: Impossibility theorems. The only mechanisms that are Pareto efficient and strategyproof are dictatorships.

- What is a dictatorship in this context?
  - Alice chooses her favorite bundle of courses
  - Betty chooses her favorite bundle of courses, out of those not yet at capacity
  - ...
  - Zoe chooses her favorite bundle of courses, out of those not yet at capacity
- Dictatorships seem unfair, even if the choosing order is random. So there is a tension amongst efficiency, incentives, and fairness.

Difficulty #2: What is Fairness? We lack realistic/well-defined criteria of fairness for this environment, in which all objects in the economy are indivisible (in particular, no money).

Budish and Cantillon (2008): Course Allocation at Harvard

HBS Draft Mechanism: students report rank-order lists, which are used to draft courses one at a time over a series of 10 rounds. Choosing order is reversed each round (e.g., ABC CBA ABC CBA)

This mechanism loosely inspires my proposed notions of fairness, but it has two important shortcomings:

1. Preference reporting is constrained. Students cannot express intensity of preference, nor complementarities or substitutabilities. This can harm efficiency.
2. Even for simple preferences, the mechanism is manipulable, even in large markets. Students should overreport their preference for "popular courses", and vice versa.

We find empirically that students strategically misrepresent their preferences, and that this misreporting causes congestion and harms welfare.

This Paper: A New Mechanism for Combinatorial Assignment

This paper proposes a solution to the combinatorial assignment problem: that is, a specific procedure guaranteed to satisfy attractive criteria of efficiency, fairness, and incentive compatibility.

1. Criteria of outcome fairness: *maximin-share guarantee*, and *envy bounded by a single good*
2. Criterion of approximate IC: *strategyproof in a large market*
3. A specific mechanism: *Approximate Competitive Equilibrium from Equal Incomes*
4. Computational Analysis of Approximate CEEI

This paper is something of a balancing act, and I am working around impossibility theorems, and so in some cases there are approximations rather than exact results.

Environment

- Set of N students $\mathcal{S}$ (s_i)
- Set of M courses $\mathcal{C}$ (c_j) with integral capacities $\mathbf{q} = (q_1, ..., q_M)$. No other goods in the economy.
- Each student s_i has a set of permissible schedules $\Psi_i \subseteq \{0,1\}^M$, and a vNM utility function $u_i : \Psi_i \to \mathbb{R}_+$
  - Complementarities, Substitutabilities are allowed
  - No peer effects. No uncertainty about preferences.
  - For the talk, assume all students have the same schedule set, Ψ
  - Maximum number of courses in a schedule is k
- An allocation $\mathbf{x} = (x_i)_{i=1}^N$ is feasible if
  - $x_i \in \Psi$ for each i; and
  - $\sum_{i=1}^N x_{ij} \leq q_j$ for each j
- An economy is a tuple $(\mathcal{S}, \mathcal{C}, \mathbf{q}, \Psi, (u_i)_{i=1}^N)$.

N.B. I often use "students" and "courses" rather than "agents" and "objects"

Competitive Equilibrium from Equal Incomes

Competitive Equilibrium from Equal Incomes has been widely studied in the context of divisible-goods economies (Foley, 1967; Varian, 1974). Here is what CEEI would mean in the context of combinatorial assignment:

1. Agents report preferences over bundles (that is, over elements of Ψ)
2. Agents are given equal budgets of an artificial currency ($= 1$, wlog)
3. We find an item price vector $\mathbf{p}^*$ such that, when each agent i chooses his favorite bundle in his budget set $\{x \in \Psi : \mathbf{p}^* \cdot x \leq 1\}$, the market clears
4. We allocate each agent their demand at $\mathbf{p}^*$

It is easy to see that existence is problematic. Consider the case in which agents have identical preferences.

CEEI as a Design Ideal

Were it not for the existence problem, CEEI would be an attractive mechanism for combinatorial assignment:

- Efficiency: Pareto Efficient
- Fairness: Symmetric, Envy Free, Maximin-Share Guaranteed
- Incentives: Strategyproof in a Large Market

Aren't CEEI-like Mechanisms Already Used in Practice?

Bidding Points Mechanisms

The most widely used course-allocation procedure is the "Bidding Points Mechanism", which works roughly as follows:

1. Each student is given an equal budget of artificial currency, say 1000 points.
2. Students express preferences by bidding for individual classes, the sum of their bids not to exceed 1000
3. For a course with q seats, the q highest bidders get a seat.

- We can interpret bids as expressions of additive-separable preferences over courses
- Schools describe the q^{th} highest bid as the "price", and the procedure as a "market"

This mechanism looks like Competitive Equilibrium from Equal Incomes (CEEI), but ...

Fake Money is Different from Real Money

- Suppose there are four scarce courses, $\{A, B, C, D\}$ and Alice bids $u_{Alice} = (700, 200, 50, 50)$.
- Prices turn out to be $\mathbf{p}^* = (900, 250, 100, 75)$.
- Her demand, with a budget of $b_{Alice} = 1000$, is

$$\arg\max_{x \in \Psi}(u_{Alice}(x) : \mathbf{p}^* \cdot x \leq b_{Alice}) \tag{1}$$
$$= \{A, C\}$$

But the BPM allocates her the bundle

$$\arg\max_{x \in \Psi}(u_{Alice}(x) - \mathbf{p}^* \cdot x) \tag{2}$$
$$= \emptyset$$

- Prices that clear the market according to demands of the form (2) exist, and are easy to compute. But that's not a good reason to use (2)!
  - incentives to misreport, which can harm efficiency.
  - ex-post unfairness (if Alice was left with 1000 *dollars* ...)

Approximate Competitive Equilibrium from Equal Incomes

1. Agents report preferences over bundles
2. Agents are given **approximately equal** budgets of an artificial currency, distributed uniformly on $[1, 1+\beta]$
3. We find an item price vector $\mathbf{p}^*$ such that, when each agent i chooses his favorite bundle in his budget set $\{x \in \Psi : \mathbf{p}^* \cdot x \leq b_i\}$, the market **approximately clears** (minimize sum-of-squares of excess demands)
4. We allocate each agent their demand at $\mathbf{p}^*$

Approximate CEEI: Overview of Main Results

- Theorem 1: For any budget inequality $\beta > 0$, there exist prices that clear the market to within market-clearing error less than $\frac{\sqrt{2kM}}{2}$. (If $\beta = 0$, error can be arbitrarily bad)
- Theorem 2: If we set $\beta < \frac{1}{N}$, the resulting Approximate CEEI satisfies an approximation to the maximin-share guarantee
- Theorem 3: If we set $\beta < \frac{1}{k-1}$, the resulting Approximate CEEI satisfies envy bounded by a single good
- Proposition 7: For any β, the Approximate CEEI Mechanism is Strategyproof in a Large Market

In combination: for β strictly postive but sufficiently small, the Approximate CEEI Mechanism is fair while maintaining attractive approximations of efficiency and IC.

Criteria of Outcome Fairness

"In fair division, the two most important tests of equity are 'fair share guaranteed' and 'no envy'" (Moulin, 1995)

> *Suppose the goods in the economy,* $\mathbf{q}$*, are perfectly divisible. An allocation* $\mathbf{x}$ *satisfies the* **fair-share guarantee** *if* $u_i(x_i) \geq u_i(\frac{\mathbf{q}}{N})$ *for all* i
>
> *An allocation* $\mathbf{x}$ *is* **envy free** *if* $u_i(x_i) \geq u_i(x_j)$ *for all* i, j

In divisible-goods economies, CEEI satisfies both criteria. Unfortunately, indivisibilities complicate fair division:

- Fair share is not well defined with indivisibilities - what is $\frac{1}{N}$ of the endowment?
- Envy freeness will be impossible to guarantee with indivisibilities. What if there are two agents and one indivisible object?

The Maximin Share Guarantee

I explicitly accept that indivisibilities complicate fair division and propose weaker criteria

Definition 1. *Fix an economy. Agent s_i's* **maximin share** *is*

$$\underline{u}_i = \max_{(x_k)_{k=1}^N} \left[\min(u_i(x_1), ..., u_i(x_N))\right]$$
$$\text{s.t. } (x_k)_{k=1}^N \text{ feasible}$$

- In words: an agent's maximin share is the maximum utility level she can guarantee herself as divider in divide-and-choose against opponents with preferences identical to her own (or adversarial opponents)
- Rawlsian guarantee from behind a "thin veil of ignorance" (Moulin, 1991)
- Coincides with fair share if goods divisible, prefs convex and monotonic

Envy Bounded by a Single Good

Definition 2. *An allocation* **x** *satisfies* **envy bounded by a single good** *if*

For any s_i, s_j: There exists some object $c_{j'}$ in bundle x_j such that:

$$u_i(x_i) \geq u_i(x_j \backslash \{c_{j'}\})$$

- In words: if student s_i envies s_j, the envy is bounded: by removing some single good from s_j's bundle we could eliminate s_i's envy
- Justification: any envy should be no larger than is necessary due to the level of indivisibility in the economy
- Coincides with envy-freeness in a limit as consumption bundles become perfectly divisible

Diamonds and Rocks Example

Example 1. Two agents. Four objects: two Diamonds (Big, Small) and two Rocks (Pretty, Ugly). At most two objects per agent.

$$\begin{aligned}\text{Maximin Share} &= \min[u(\{\text{Big Diamond, Ugly Rock}\}), \\ &\quad u(\{\text{Small Diamond, Pretty Rock}\})] \\ &= u(\{\text{Small Diamond, Pretty Rock}\})\end{aligned}$$

- The allocation in which one agent obtains {Small Diamond, Pretty Rock} and the other obtains {Big Diamond, Ugly Rock} also satisfies envy bounded by a single good
- The procedural fairness of requirement of symmetry requires randomization over who gets which bundle

Dictatorships and Fairness

- Dictatorships are procedurally fair if the choosing order is uniform random
- However, dictatorships fail the outcome fairness criteria
  - Whichever student chooses first gets both diamonds
- The criteria help to formalize why dictatorships are unfair in multi-unit assignment

Remark 1: In single-unit assignment (e.g., one diamond, one rock), dictatorships satisfy the maximin-share guarantee and envy bounded by a single good.

- Dictatorships are frequently used in practice for single-unit assignment problems (school choice, housing assignment)

Strategyproof in a Large Market

Definition 4. *The continuum replication of* $(\mathcal{S}, \mathcal{C}, \mathbf{q}, (\Psi_i)_{i=1}^N, (u_i)_{i=1}^N)$, *written* $(\mathcal{S}^\infty, \mathcal{C}, \mathbf{q}, (\Psi_i^\infty)_{i=1}^N, (u_i^\infty)_{i=1}^N)$ *is constructed as follows*

- *The set of students is* $\mathcal{S}^\infty = \{\widetilde{s}_i\}_{i \in (0,N]}$
- *The set of courses and their capacities are left as is, except now we understand a course's capacity constraint as a Lebesgue measure of students in the course.*
- *Students numbered* $(0,1]$ *in the continuum are identical to* s_1 *in the original, students numbered* $(1,2]$ *in the continuum are identical to* s_2 *in the original, etc.*

Definition 5. *A mechanism is* **strategyproof in a large market** *if, for any economy, it is strategyproof (dominant-strategy IC) in the continuum replication* $(\mathcal{S}^\infty, \mathcal{C}, \mathbf{q}, (\Psi_i^\infty)_{i=1}^N, (u_i^\infty)_{i=1}^N)$

A Simple Interpretation of SP in a Large Market

An agent's report serves two distinct functions

1. Subtly influences the agent's "opportunity set"
    - Competitive Eqm: i's opportunity set is $\{x_i : \mathbf{p}^* \cdot x_i \leq b_i\}$.
        - i's report might affect $\mathbf{p}^*$
    - Gale-Shapley Male-Proposing Deferred Acceptance: female i's opportunity set is the set of proposals she receives.
        - i's report affects which proposals she *rejects*, and each rejection might cause a "rejection chain" that causes new proposals to occur (Kojima and Pathak, 2008)
2. Acts on the realized opportunity set to select an outcome
    - CE: choose favorite bundle in $\{x_i : \mathbf{p}^* \cdot x_i \leq b_i\}$
    - GS M-P DA: i is allocated her most preferred proposal

Continuum replication turns off 1. A mechanism is SP in a Large Market if for any *fixed* opportunity set, truthful reporting is best.

Which Market Designs are Strategyproof in a Large Market?

Manipulable in Large Markets	SP in Large Markets
Bidding Points Mechanism	Deferred Acceptance
HBS Draft Mechanism	Double Auctions
Boston Mechanism	Assignment Exchange
All-Pay Auctions	Probabilistic Serial
Discriminatory Auctions	Uniform Price Auctions

- All course-allocation mechanisms currently found in practice are manipulable even in large markets
  - Budish and Cantillon (2008): empirical evidence that this matters for welfare
- By contrast, many widely used non-strategyproof mechanisms are strategyproof in a large market
  - Of course, in several instances we have a much more highly detailed understanding of incentives away from the limit (e.g., Cripps and Swinkels 2006)

Approximate CEEI

Definition 6. *Fix an economy. The allocation* $\mathbf{x}^* = (x_1^*, ..., x_N^*)$*, non-negative budgets* $\mathbf{b}^* = (b_1^*, ..., b_N^*)$ *(with* $\min b_i^* = 1$*,wlog), and non-negative prices* $\mathbf{p}^* = (p_1^*, ..., p_M^*)$ *constitute an* (α, β)-**approximate competitive equilibrium from equal incomes** *(Approximate CEEI) of this economy if:*

(i) $x_i^* = \arg\max_{x_i' \in \Psi} [u_i(x_i') : \mathbf{p}^* \cdot x_i' \leq b_i^*]$ *for all* $i = 1, ..., N$

(ii) $||\mathbf{z}^*||_2 \leq \alpha$ *where* $\mathbf{z}^* = (z_1^*, ..., z_M^*)$ *and*

$z_j^* = \sum_i x_{ij}^* - q_j$ *if* $p_j^* > 0$

$z_j^* = \max(\sum_i x_{ij}^* - q_j, 0)$ *if* $p_j^* = 0$

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

- α is the approximation error in market clearing, measured in Euclidean distance (CE: $\alpha = 0$)
- β is the maximum inequality in the budgets, measured as a proportion (EI: $\beta = 0$)

Demand Sensitivity Parameter

- Students' demands are discontinuous with respect to price
- The largest possible individual demand discontinuity occurs when a tiny change in price causes a student's demand to change from some bundle of k courses to an entirely disjoint bundle of k courses (k is the maximum number of courses in any permissible schedule)
  - E.g. a tiny change in price might cause demand to change from {Big Diamond, Ugly Rock} to {Small Diamond, Pretty Rock}
  - If $2k > M$ then the largest discontinuity will involve the student's demand changing for each of the M courses
- Let $\sigma = \min(2k, M)$ denote the magnitude of this maximum-possible demand discontinuity (measured as a square of Euclidean distance, i.e., as sum-of-squares)

Theorem 1

Existence of Approximate CE from Approximate EI

Theorem 1. *Fix an economy. For any* $\beta > 0$*, there exists a* $(\frac{\sqrt{\sigma M}}{2}, \beta)-$*Approximate CEEI.*

In particular, for any budget vector $\mathbf{b}'$ *that satisfies* $\max_i(b_i') \leq 1 + \beta$ *and* $\min_i(b_i') = 1$*, and any* $\varepsilon > 0$*, there exists a* $(\frac{\sqrt{\sigma M}}{2}, \beta)-$*Approximate CEEI with budgets of* $\mathbf{b}^*$ *that satisfy* $|b_i^* - b_i'| < \varepsilon$ *for all* i*.*

Remark 1: $\frac{\sqrt{\sigma M}}{2}$ is small in two senses

- $\frac{\sqrt{\sigma M}}{2}$ does not grow with N (number of students) or $\mathbf{q}$ (number of copies of each good). As $N, \mathbf{q} \to \infty$, we get exact market clearing (error goes to zero as a fraction of the endowment)
- $\frac{\sqrt{\sigma M}}{2}$ is a small number in practical problems (especially for a worst case bound)
    - For instance, in a semester at HBS, $k = 5$ and $M = 50$, and so $\sigma = 10$ and $\frac{\sqrt{\sigma M}}{2} \approx 11$ (11 seats in one class, or 3 seats in each of 13 classes, etc.)
    - Contrast with 4500 course seats allocated per semester

N.B. a first welfare theorem indicates that an approximate CEEI allocation is Pareto efficient with respect to the set of goods actually allocated. Market-clearing error is of course inefficient.

Remark 2: The market-clearing error bound is tight.

Remark 3: This is an approximate existence result for item prices in an economy with complex preferences including complementarities.

- A difference versus the general combinatorial auction setting is that here the complementarities are "small", because students want at most one of each good. (Captured in σ)

Proof of Theorem 1: Overview

Consider a tâtonnement price-adjustment function of the form

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

A fixed point $f(\mathbf{p}') = \mathbf{p}'$ would be a competitive equilibrium, but need not exist.

1. Mitigate discontinuities in $f(\cdot)$ using budget perturbations
2. "Convexify" $f(\cdot)$ into a correspondence $F(\cdot)$, and then obtain a fixed point $\mathbf{p}^* \in F(\mathbf{p}^*)$
3. Map from price space to demand space in a neighborhood of $\mathbf{p}^*$. What is the structure of demand discontinuities?
4. Bound market-clearing error, using the structure of demands near to $\mathbf{p}^*$. Use an exact fixed point of $F(\cdot)$ to find an approximate fixed point of $f(\cdot)$

Step 1: Mitigate Discontinuities in f()

Choice Sets

How do students' choice sets change with price?

- Consider a student s_1 whose budget is 1000 and whose favorite bundle x' at some price vector costs 999.
- Increase cost of x' to 1001. s_1's choice set is different.
- Her demand will change. Any time her demand changes it does so 'discontinuously' because all goods in the economy are discrete. The size of this discontinuity is bounded by $\sqrt{\sigma}$.
- Suppose students $s_2, \ldots, s_N$ also have a budget of 1000. Then their choice sets vary identically with s_1's, and their demands might vary together as well. The discontinuity might be size $N\sqrt{\sigma}$
- Suppose instead s_2 has a budget of 1002, ... Now, the change in price that changed s_1's choice set need not change s_2's choice set.

Step 1: Mitigate Discontinuities in f()

Budget-Constraint Hyperplanes

- The story is a bit more complicated because we use item prices, not bundle prices. By changing the price of bundle x' we necessarily change the price of other bundles.
- Let $H(s_i, x) = \{\mathbf{p} : \mathbf{p} \cdot x = b_i\}$ denote the hyperplane along which student s_i can exactly afford bundle x.
- Every time price crosses a budget-constraint hyperplane there is a potential discontinuity in demand, hence in $f(\cdot)$
- The number of such "budget-constraint hyperplanes" is finite.
- I define a perturbation scheme – a tiny "tax / credit" specific to each agent-bundle pair – that ensures
    - No two of the hyperplanes are identical
    - No more than M hyperplanes intersect at any one price vector (M is the number of courses, hence prices)
- This perturbation is the reason for ε in the Theorem statement

Budget-Constraint Hyperplanes

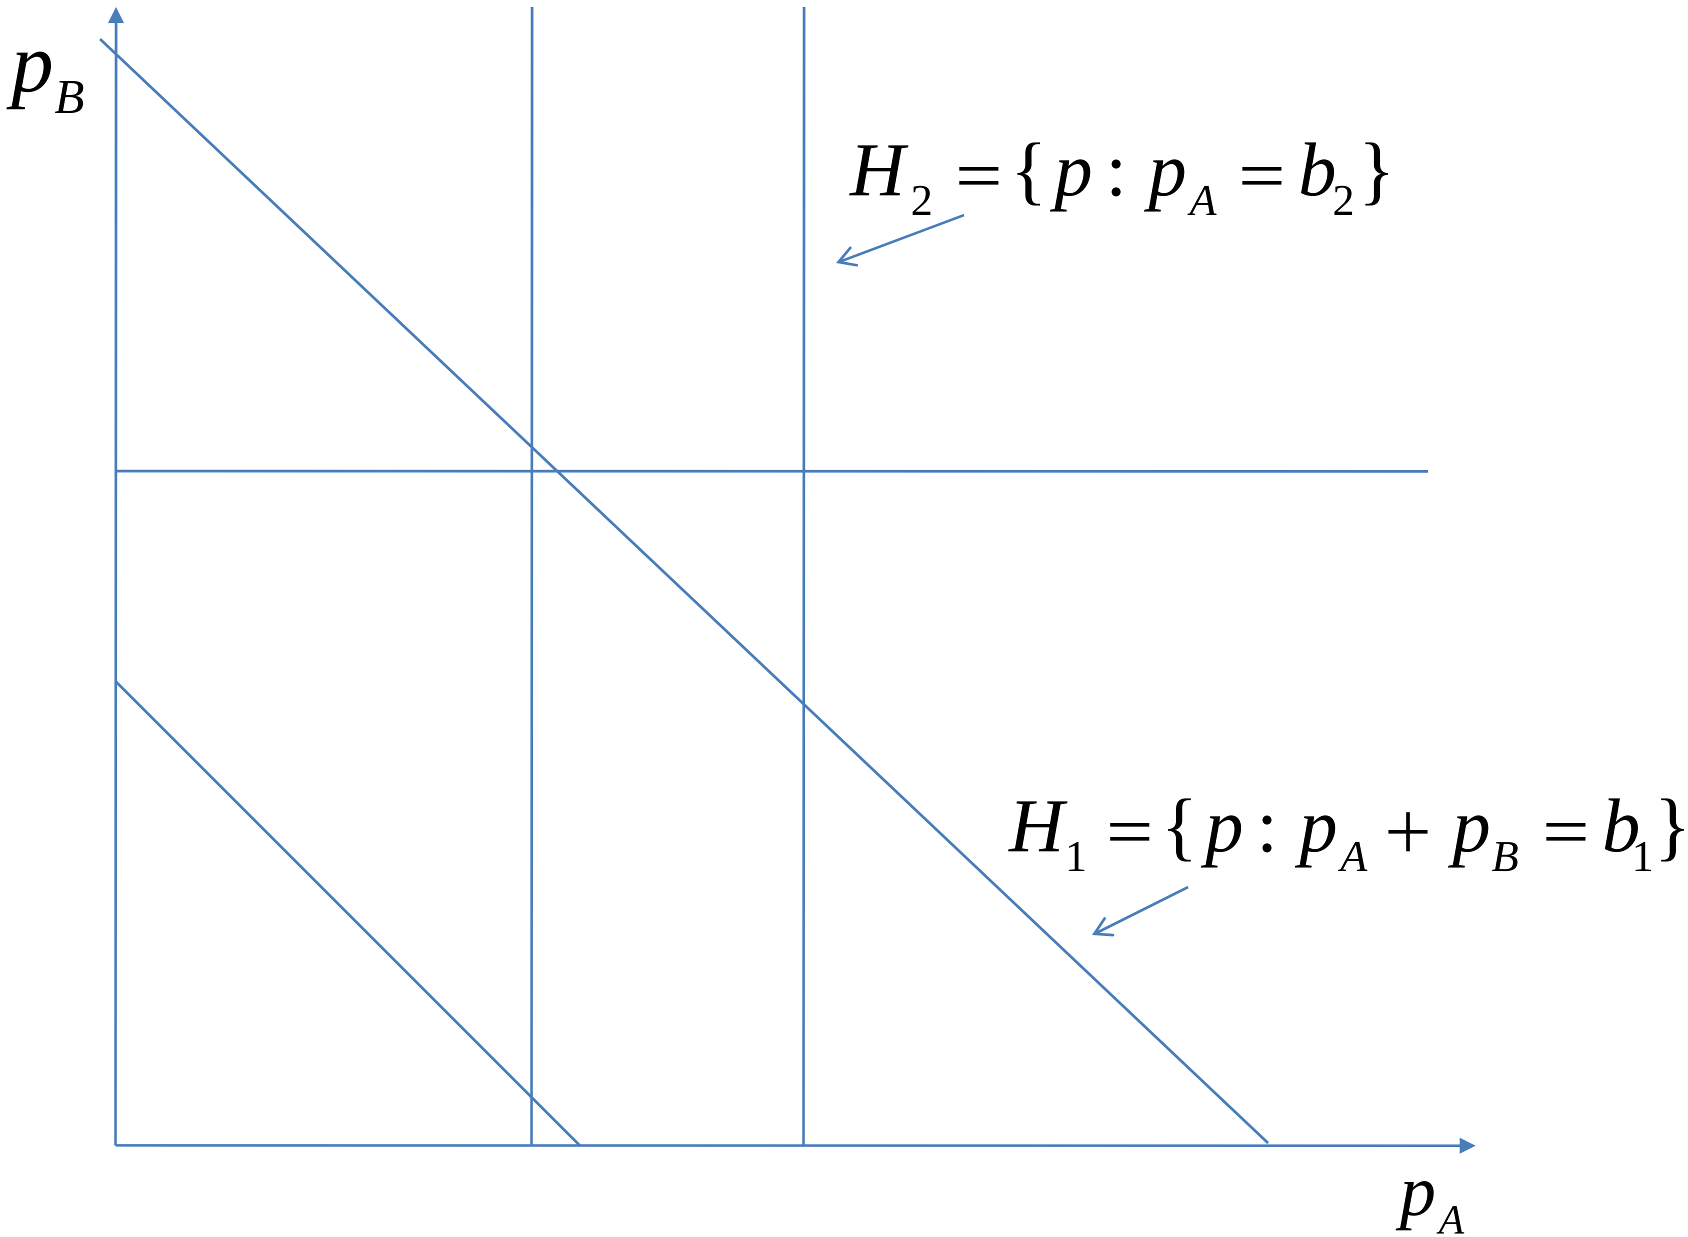

Step 2: Convexify f() and Obtain a Fixed Point

Consider the following "convexification" of $f(\cdot)$:

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

- Cromme and Diener (1991, Lemma 2.4): for any map f on a compact and convex set, the correspondence F constructed as above has a Kakutani fixed point
  - (Our price space is $[0, 1 + \beta + \varepsilon]^M$, which is compact and convex)
- What does $\mathbf{p}^* \in F(\mathbf{p}^*)$ mean? A convex combination of excess demands close to $\mathbf{p}^*$ exactly clears the market.
- We will use the exact fixed point of $F(\cdot)$ to find an approximate fixed point of $f(\cdot)$

Step 3: Map from Price Space to Demand Space

- We can put a lot of structure on demands near to $\mathbf{p}^*$
- Suppose that $\mathbf{p}^*$ lives on L budget-constraint hyperplanes. Step 1 ensures $L \leq M$. If $L = 0$, we are done. There are two key ideas:

1. If a price vector $\mathbf{p}'$ is close enough to $\mathbf{p}^*$, the only differences in choice sets will arise from which side $\mathbf{p}'$ is on of the L hyperplanes that intersect at $\mathbf{p}^*$.
   - Out of a whole neighborhood, we can limit attention to 2^L points
2. Each agent's demand depends only on which side of their *own* b-c-h price is on
   - A "change-in-demand vector" $v_i \in \{-1, 0, 1\}^M$ describes how agent s_i's demand changes as price crosses her particular b-c-h.
   - L such vectors describe how demand changes over the 2^L points

(Case of some students owning multiple of the L b-c-h's is handled in the proof)

Fixed Point of Convexified Tatonnement F(·)

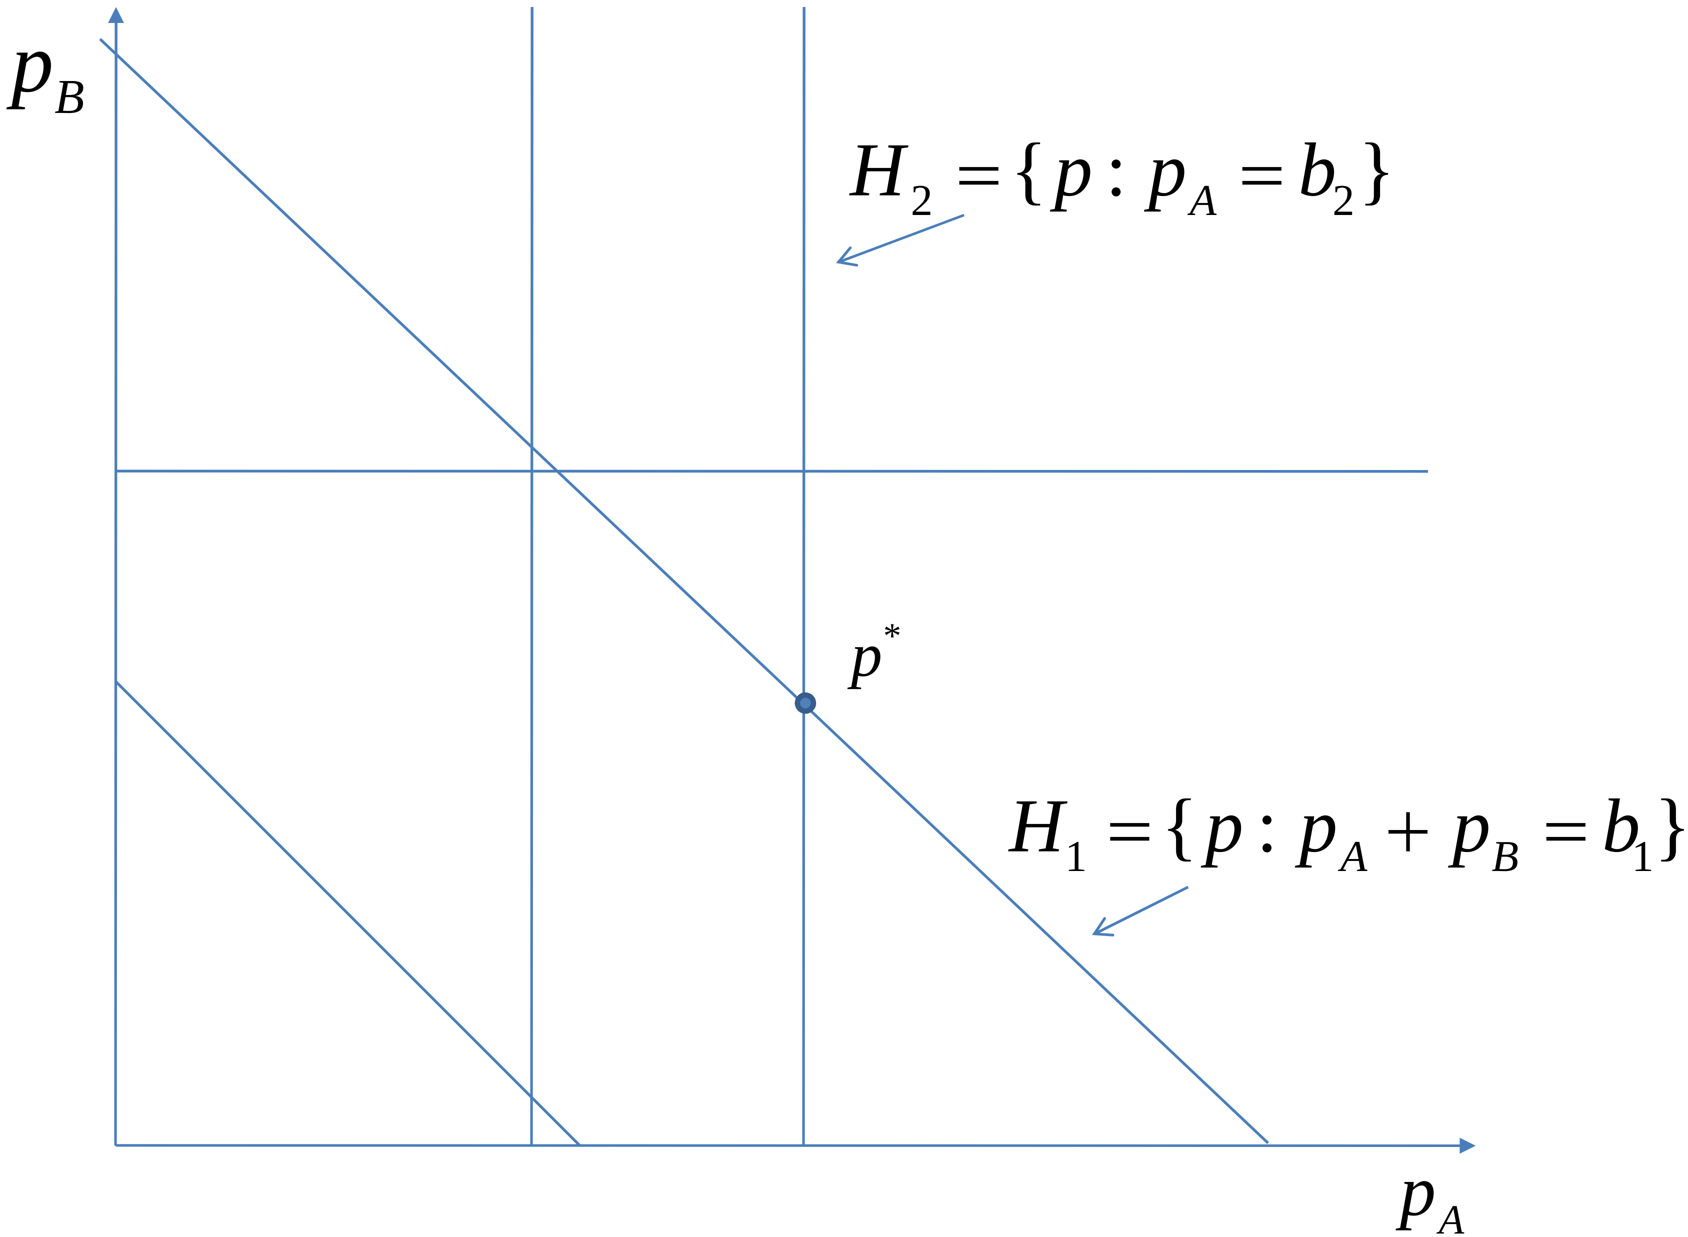

Map from Price Space to Demand Space I
Ball around p*

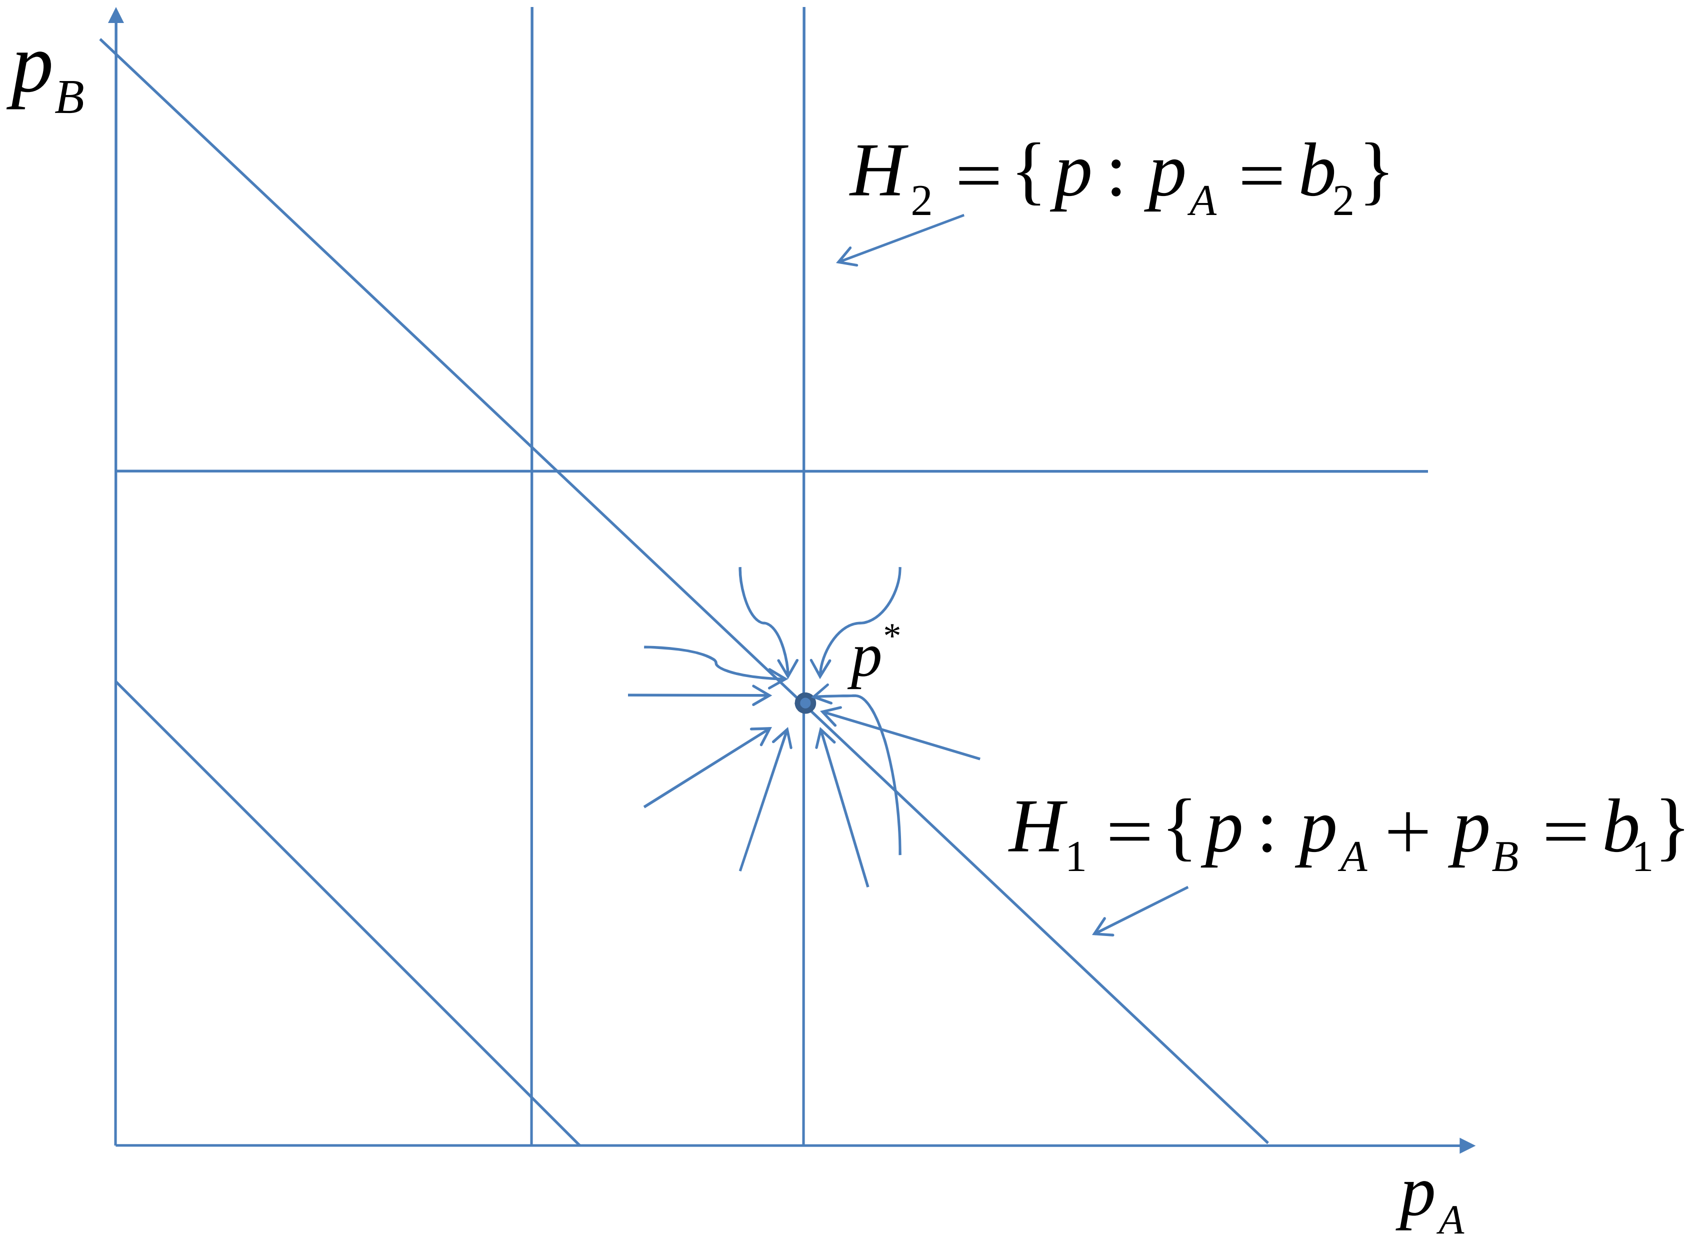

Map from Price Space to Demand Space II

2^L points near p*

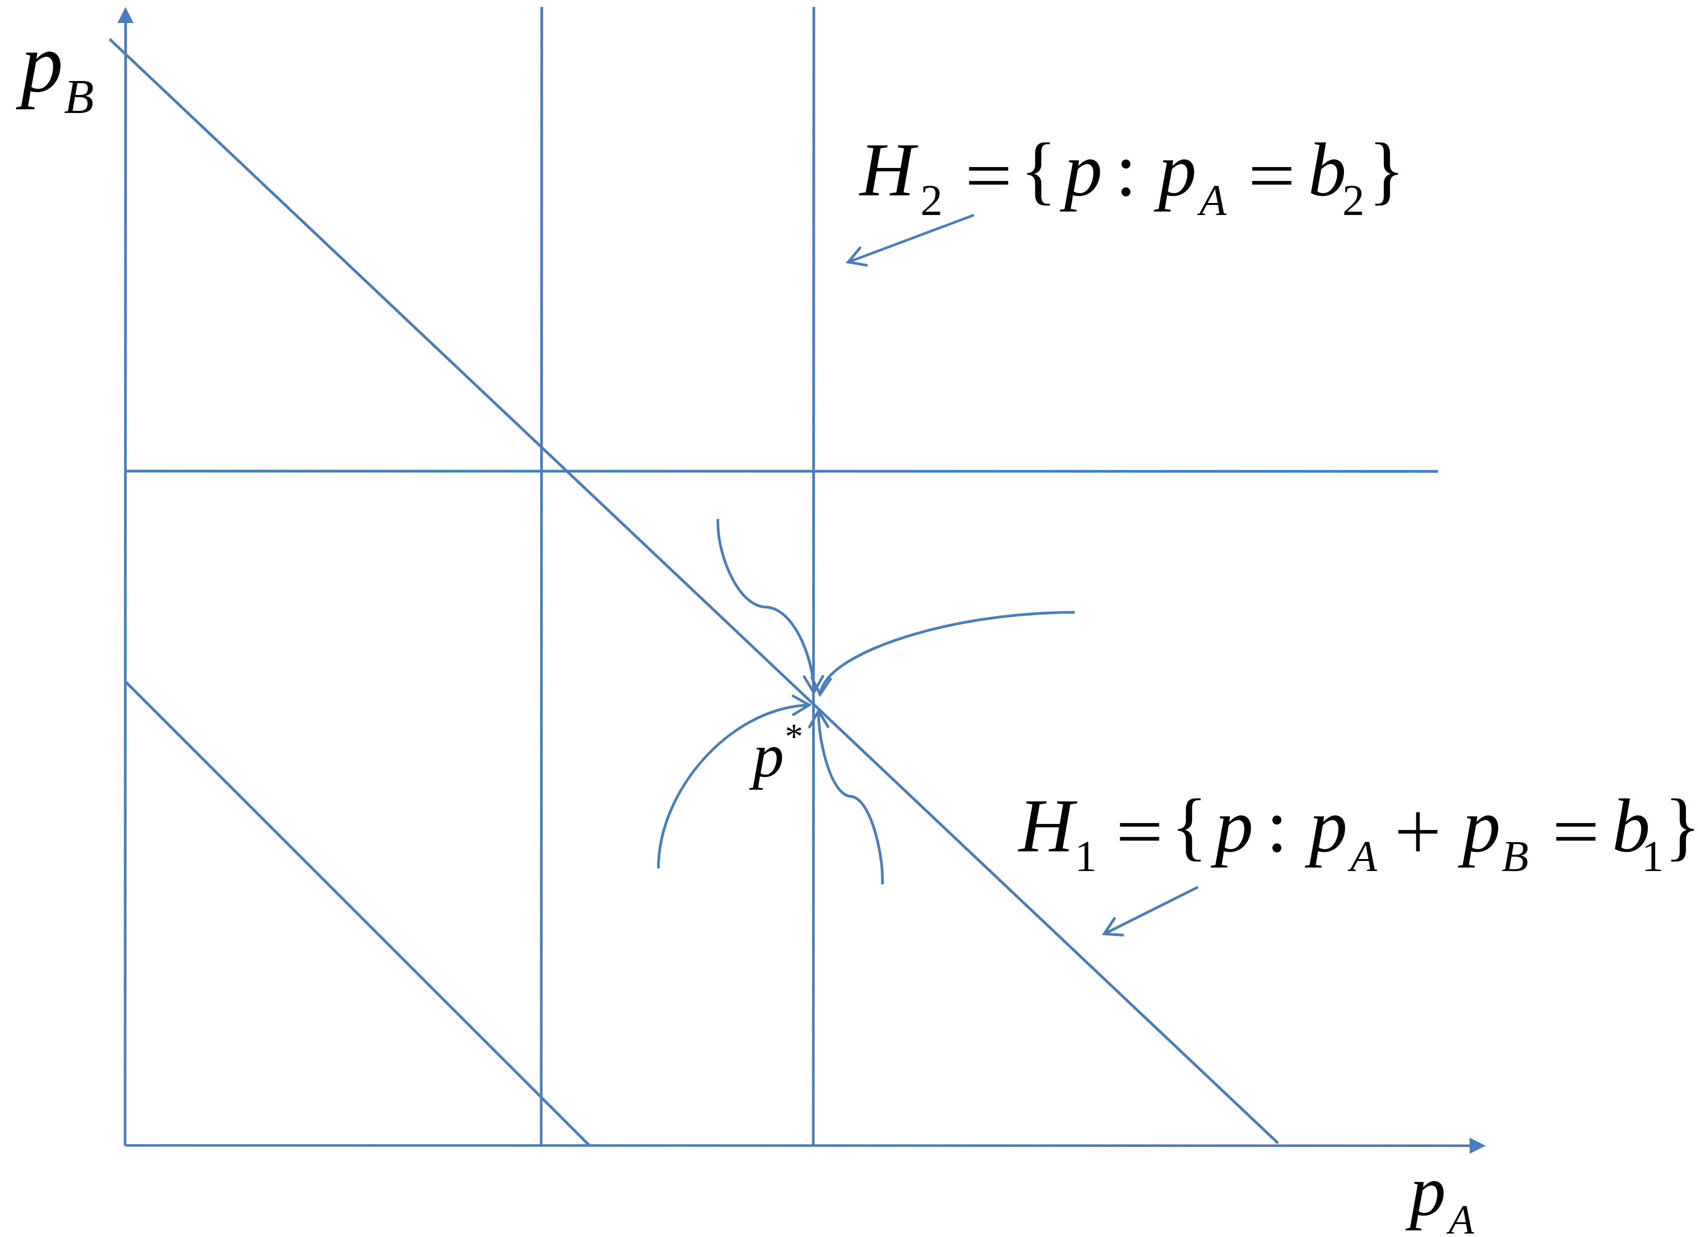

Map from Price Space to Demand Space III
L Change-in-Demand Vectors near p*

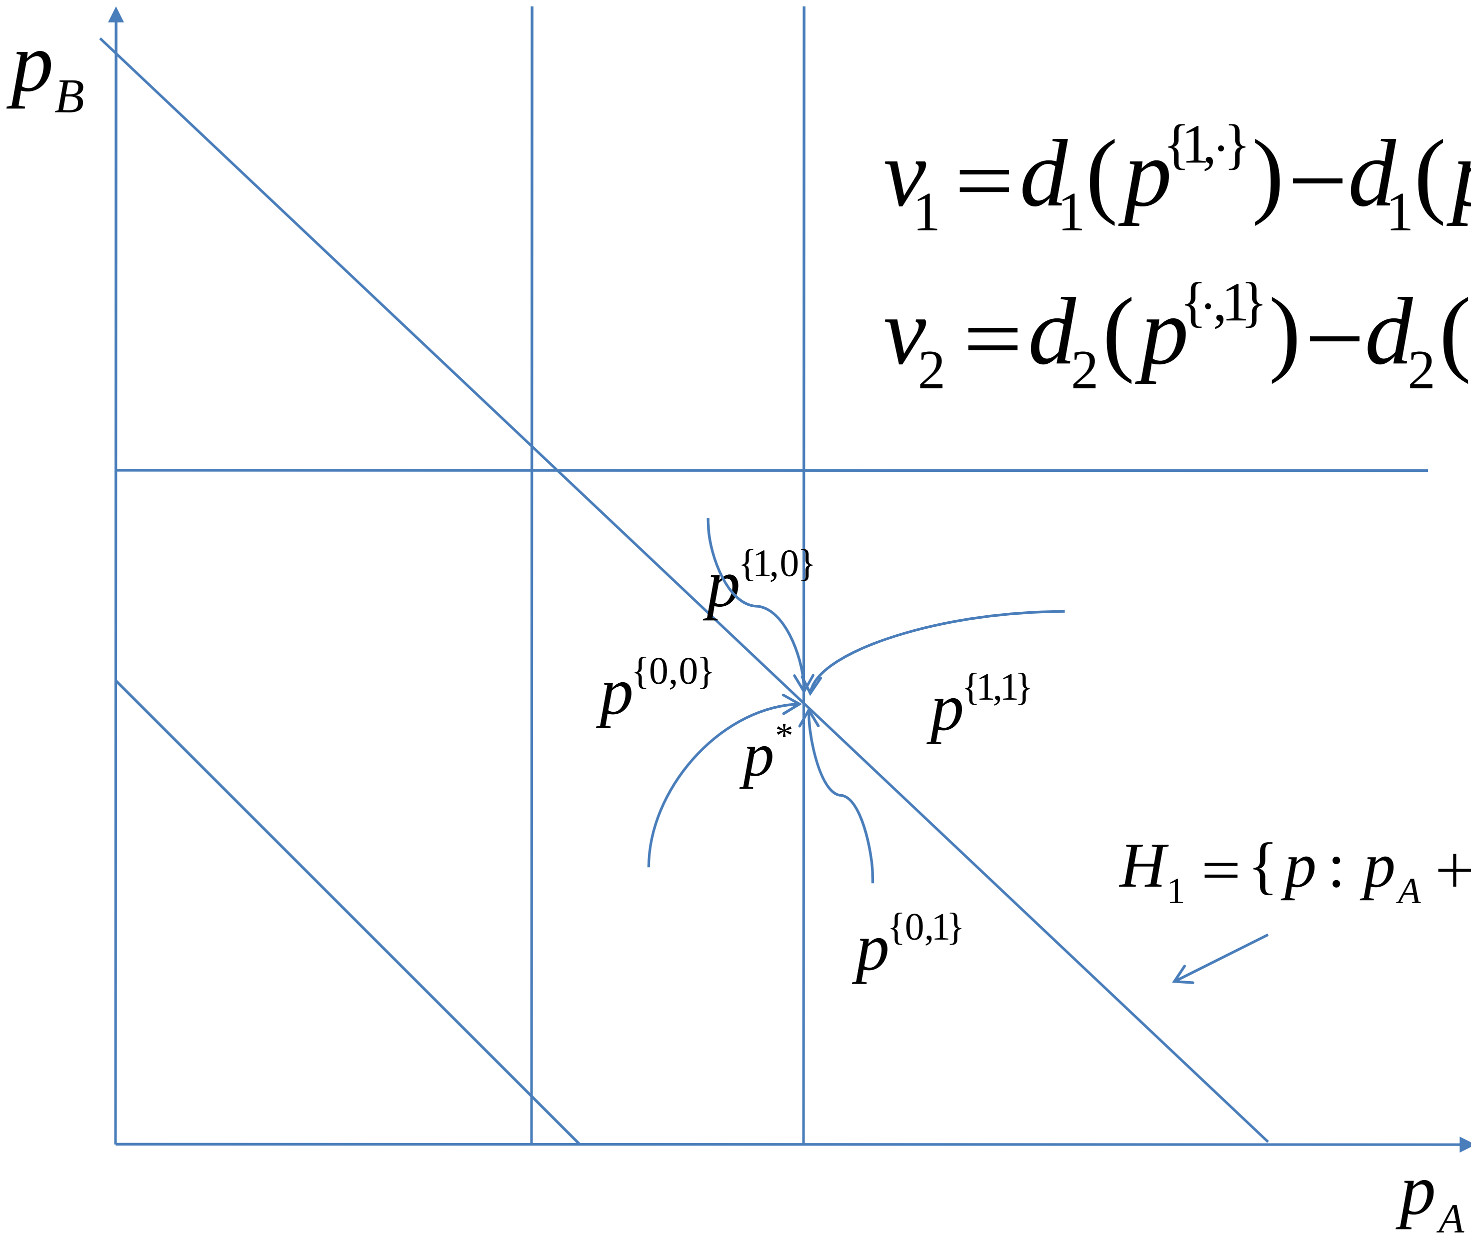

$$v_1 = d_1(p^{\{1,\cdot\}}) - d_1(p^{\{0,\cdot\}})$$

$$v_2 = d_2(p^{\{\cdot,1\}}) - d_2(p^{\{\cdot,0\}})$$

Step 4: Obtaining the Bound

From step 3, the set of feasible excess demands in an arbitrarily small neighborhood of $\mathbf{p}^*$ is

$$\{a \in \{0,1\}^L : z(\mathbf{p}^*) + \sum_{(i)=1}^{L} a_{(i)} v_{(i)}\} \qquad (*)$$

- Now $\mathbf{p}^* \in F(\mathbf{p}^*)$ tells us something much more useful! Perfect market clearing is in the convex hull of (*).
- Our market-clearing error bound is the maximum-minimum distance between a vertex of (*) and a point in its interior
- The worst case occurs when $L = M$, the M change-in-demand vectors $\{v_{(i)}\}_{i=1}^{L}$ each have length $\sqrt{\sigma}$ and are mutually orthogonal, and the perfect-market clearing ideal is equidistant from all 2^M vertices of (*)
- Half the diagonal length of an M-dimensional hypercube of side length $\sqrt{\sigma}$ is $\frac{\sqrt{\sigma M}}{2}$
- Note that $z(\mathbf{p}^*)$ itself need not be within $\frac{\sqrt{\sigma M}}{2}$ of perfect market clearing.

Fairness Properties of Approximate CEEI

To what extent do approximately equal budgets guarantee that students will receive fair outcomes ex-post? We might worry for several reasons

- In single-unit demand case, cardinal budget information is meaningless; all that matters is the order of the budgets
  - e.g., two students and two objects, no difference between budgets of $(1000, 999)$ and $(1000, 1)$. In either case, the budget of 1000 gets his favorite object.
- More generally, since goods are indivisible, students' optimal consumption bundles might not exhaust their budgets.
  - e.g., a student whose favorite bundle costs 1000 and whose second favorite bundle costs 1 doesn't care if her budget is 999 or 1.

Theorem 2: Approximate CEEI Guarantees Approximate Maximin Shares

Theorem 2: if $\beta < \frac{1}{N}$ then $\mathbf{x}^*$ guarantees each agent their $N+1$-maximin share (maximin share in a hypothetical economy with one additional agent)

Intuition for proof

1. If $\beta < \frac{1}{N} \Rightarrow$ even poorest student has $> \frac{1}{N+1}$ of the income endowment
2. If $\mathbf{p}^*$ is an exact c.e. $\Rightarrow$ goods endowment costs weakly less than the income endowment.
3. So if $\mathbf{p}^*$ is an exact c.e., each student must be able to afford some bundle in any $N+1$-way split.
4. Hence, each student must be able to afford some bundle weakly preferred to her $N+1$-maximin share.

Argument is a bit messier because $\mathbf{p}^*$ might be an approximate c.e. The proof exploits the Kakutani fixed-point step from the proof of Theorem 1.

Theorem 3: Approximate CEEI Guarantees that Envy is Bounded by a Single Good

We know that exactly equal incomes guarantees exact envy-freeness, because all students have the same choice set.

Theorem 3: if $\beta < \frac{1}{k-1}$ then $\mathbf{x}^*$ satisfies envy bounded by a single good

- Intuition: suppose s_i envies s_j. Then

$$1 \leq b_i^* < \mathbf{p}^* \cdot x_j^* \leq b_j^* \leq \frac{k}{k-1}$$

- Since x_j^* contains at most k goods, one of them must cost at least $\frac{1}{k-1}$. s_i can afford the bundle formed by removing this good from x_j^*
- By revealed preference, s_i must weakly prefer her own bundle to the bundle formed by removing this single good from x_j^*, so her envy is bounded.

Notice that the level of budget inequality required for Theorem 3 is different from that required for Theorem 2

Putting it all Together: The Approximate CEEI Mechanism

The Approximate CEEI Mechanism:

1. Agents report utility functions
2. Set β suitably small and choose non-identical budgets uniform randomly from $[1, 1+\beta]$. Compute the set of $(\frac{\sqrt{\sigma M}}{2}, \beta)-$Approximate CEEI price vectors.
3. Choose uniform randomly from amongst those price vectors with the smallest market-clearing error.
4. Announce the prices $\mathbf{p}^*$, budgets $\mathbf{b}^*$ and allocation $\mathbf{x}^*$.

Note 1: choosing prices uniform randomly ensures that the procedure is strategyproof in a large market. There are other possible tie-breaking rules that preserve incentives in this way.

Note 2: it is possible to add a step in which we first seek an exact CEEI.

Putting it all Together: The Approximate CEEI Mechanism

Properties of the Approximate CEEI Mechanism.

Efficiency
- Ex-post efficient with respect to the allocated goods.

Fairness
- Symmetric
- N+1 Maximin Share Guaranteed
- Envy Bounded by a Single Good

Incentives
- Strategyproof in a Large Market

Properties are especially attractive in a continuum economy

Relationship to Random Serial Dictatorship

Single-Unit Demand

- The Approximate CEEI Mechanism coincides with Random Serial Dictatorship
- Both satisfy maximin-share guarantee and envy bounded by a single good
- Dictatorships frequently used in practice (school choice, housing assignment)

Multi-Unit Demand

- The mechanisms are importantly different.
- Suppose students require at most k objects. RSD corresponds to an exact competitive equilibrium ($\alpha = 0$) from budgets of

$$\mathbf{b}^{RSD} = (1, k+1, (k+1)^2, (k+1)^3, ..., (k+1)^{N-1})$$

- Dictatorships not observed in practice

Table 2: Comparison of Alternative Mechanisms

Mechanism	Efficiency (Truthful Play)	Fairness (Truthful Play)	Incentives	Preference Language
Approximate CEEI Mechanism	Pareto Efficient w/r/t Allocated Goods Allocation error is small for practice and goes to zero in the limit	N+1 – Maximin Share Guaranteed Envy Bounded by a Single Object	Strategyproof in the Limit	Ordinal over Schedules
HBS Draft Mechanism	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 Object	Manipulable	Ordinal over Items
Bidding Points Mechanism	If preferences are additive-separable, Pareto Efficient but for quota issues described in Unver and Sonmez (2008)	Worst Case: Get Zero Objects	Manipulable	Cardinal over Items
Unver-Sonmez Enhancement to Bidding Points Mechanism	If preferences are additive-separable, Pareto Efficient	Worst Case: Get Zero Objects	Bidding Phase: Manipulable Allocation Phase: SP in the Limit	Bidding Phase: Cardinal over Items Allocation Phase: Ordinal over Items
Random Serial Dictatorship	Pareto Efficient	Worst Case: Get k worst Objects	Strategyproof	Ordinal over Schedules
UChicago Primal-Dual Linear Program Mechanism	Pareto Efficient when preference-reporting limits don't bind	Worst Case: Get Zero Objects	Manipulable	Cardinal over a Limited Number of Schedules
Pratt Geometric Prices Mechanism	If preferences are additive-separable, Pareto Efficient	Worst Case: Get Zero Objects	Strategyproof in the Limit	Cardinal over Items
Brams and Taylor Adjusted Winner	If preferences are additive-separable, Pareto Efficient	Worst Case: Get Zero Objects	Manipulable	Cardinal over Items
Herreiner and Puppe Descending Demand Procedure	Pareto Efficient	Does not satisfy Maximin Share Guarantee or Envy Bounded by a Single Object	Manipulable	Ordinal over Schedules

References: HBS – Budish and Cantillon (2008); Bidding Points Mechanisms – Unver and Sonmez (2008); RSD – Budish and Cantillon (2008); UChicago Primal-Dual Linear Program Mechanism – Graves et al (1993); Pratt(2007); Brams and Taylor (1996); Herreiner and Puppe (2002)

Ex-Post versus Ex-Ante Evaluation of Efficiency and Fairness

- Fairness: ex-post is actually the more stringent requirement.
  - For instance, RSD is ex-ante envy free even though it is very unequal ex-post.
  - Additionally, ex-post is likely the perspective that matters to real-life market administrators
- Efficiency: ex-ante is the more stringent perspective
  - a necessary but not sufficient condition for a lottery over allocations to be ex-ante Pareto efficient is that all its realizations are ex-post Pareto efficient
  - impossibility theorems are even more severe (Zhou, 1990)
- I leave theoretical analysis of the tradeoff between ex-post fairness and ex-ante efficiency to future work (Budish, Che, Kojima and Milgrom, in progress)
- Here, I show that the proposed mechanism has attractive ex-ante efficiency performance in a specific course-allocation environment

Computational Analysis - Algorithm

Theorem 1 is non constructive, and implementing the Approximate CEEI Mechanism is non-trivial. There are two key challenges:

1. Calculating excess demand at a particular price ($z(\mathbf{p})$) is NP Hard – each agent must solve a set-packing problem
2. Price space is large. So even if $z(\mathbf{p})$ were easy to compute, finding an approximate zero is a difficult search problem

Othman, Budish and Sandholm (2008) develop a computational procedure that overcomes these challenges in life-size problems.

1. Demands are calculated using an integer program solver, CPLEX
2. We use a method called "Tabu Search" to find an approximate zero. Departure point is the Tatonnement process $\mathbf{p}^{t+1} = \mathbf{p}^t + z(\mathbf{p}^t)$

The algorithm can currently handle "semester-sized" economies in which students consume 5 courses. Each run takes 1 hour.

Computational Analysis - Data and Key Assumptions

Budish and Cantillon (2008) data on course allocation at Harvard Business School:

- 50 Fall courses, 47 Spring Courses. About 20 per semester are scarce.
- 456 students: Actual stated preferences and underlying truthful preferences (from surveys – out of 916 students total)
- HBS preferences are ordinal over individual courses.
- To convert into utilities from bundles I make two substantive assumptions
  - A1: Additive-Separable Preferences (no comps or subs)
  - A2: Students care about the "Average Rank" of the courses they receive (e.g. 2nd + 3rd favorite better than 1st + 5th favorite)
  - Theory can handle more complex preferences but A1 and A2 seem reasonable given data incompleteness
- A3: Students report their preferences truthfully under Approximate CEEI
  - Caveat: no way to validate whether 916 students is "large"

Average Rank Preferences: Example

Suppose students require 2 courses each, and in the data I observe that student i's true preference ordering over individual courses is:

$$1, 2, 3, 4, 5, \ldots$$

Then I assume that her true preference ordering over schedules satisfies:

$$u_i(\{1,2\}) > u_i(\{1,3\}) > u_i(\{1,4\}) = u_i(\{2,3\}) > u_i(\{1,5\})\ldots$$

I don't need to make any further assumptions about the magnitudes of these inequalities.

Figure 1: Distribution of Market-Clearing Error

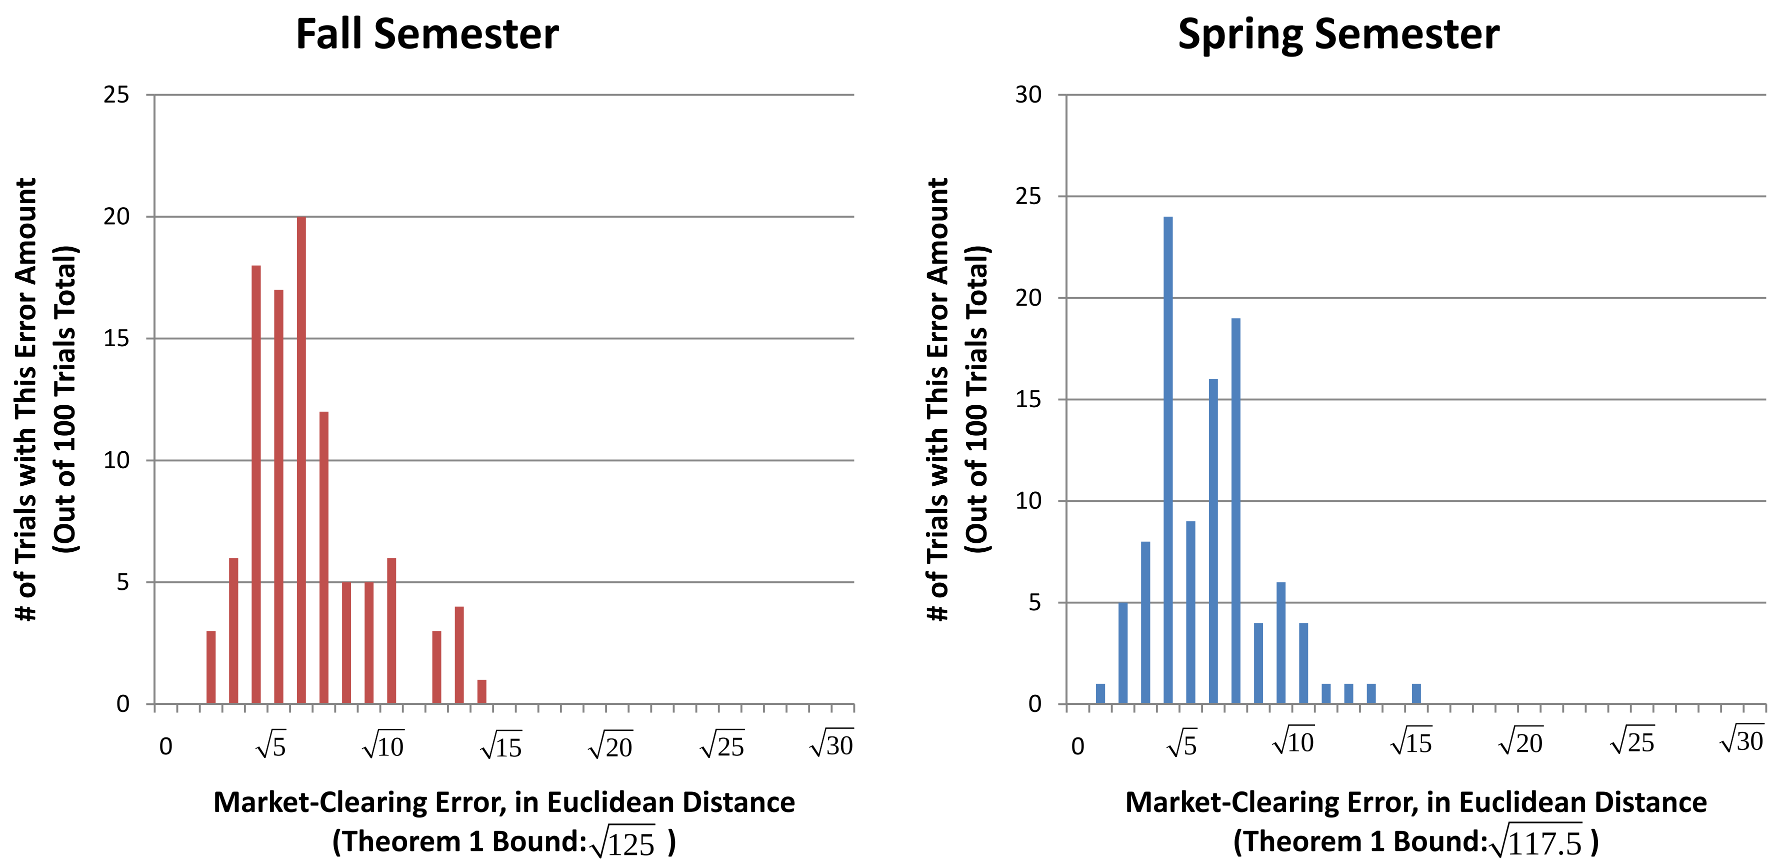

Description: The Othman, Budish and Sandholm (2008) Approximate CEEI algorithm is run 100 times for each semester of the Harvard Business School course allocation data (456 students, ~50 courses, 5 courses per student). Each run uses randomly generated budgets. This table reports the distribution of the amount of market-clearing error per trial, measured in Euclidean Distance (square-root of sum of squares). 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).

Figure 2: Distribution of Difference Between Best and Worst Outcomes

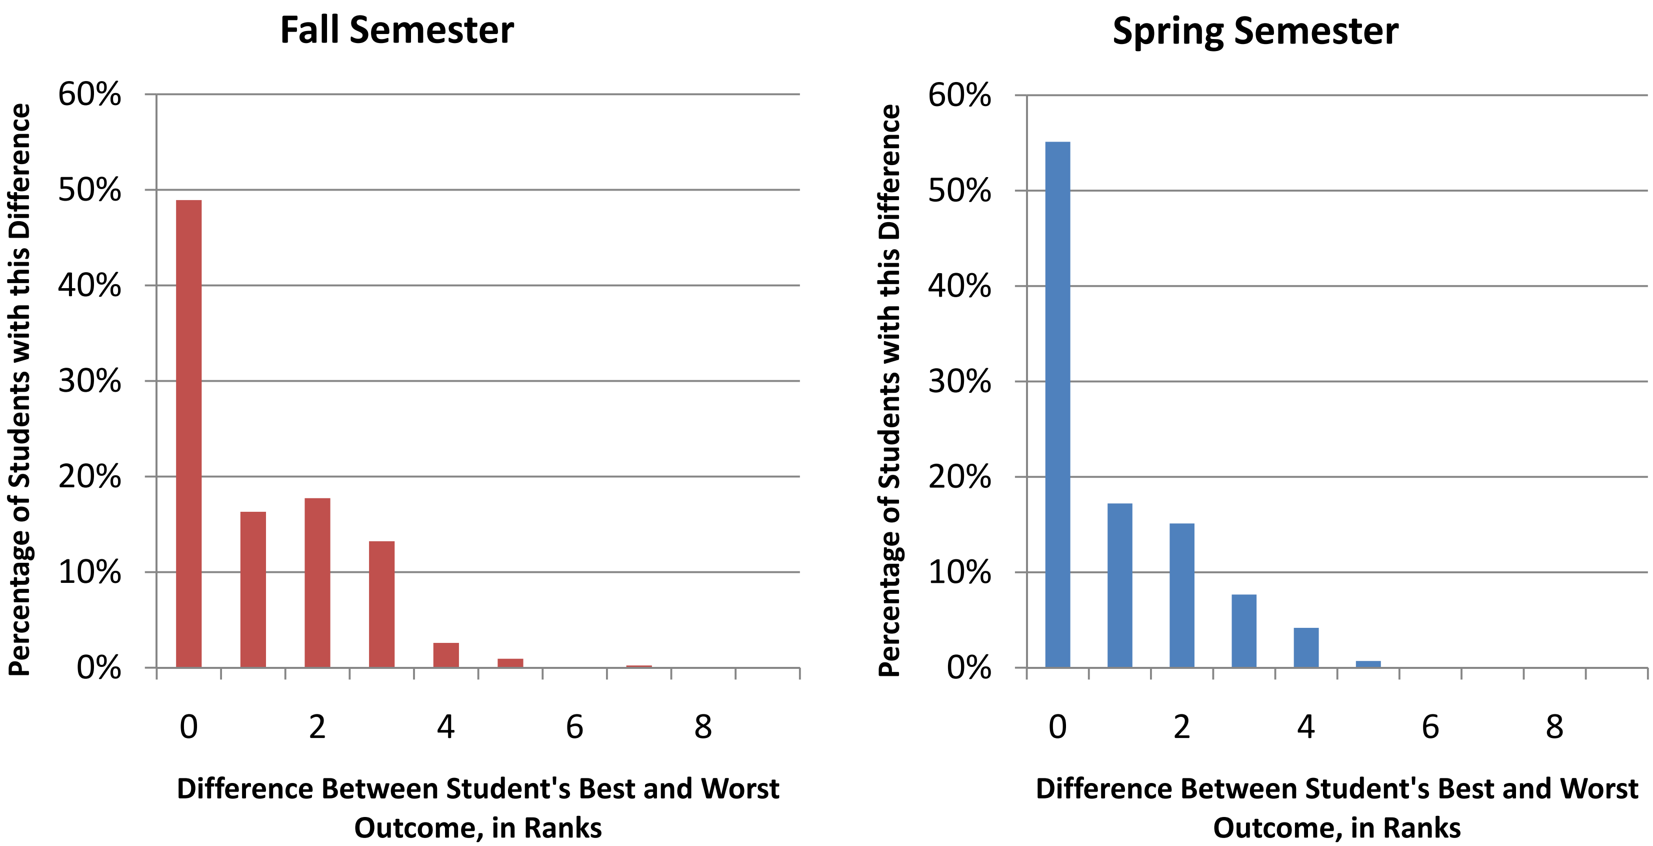

Description: The Othman, Budish and Sandholm (2008) Approximate CEEI algorithm is run 100 times for each semester of the Harvard Business School course allocation data (456 students, ~50 courses, 5 courses per student). Each run uses randomly generated budgets. This table reports the distribution of the difference between a student's single best and single worst outcome over the 100 trials, in ranks. Here is an example calculation: a student whose best received bundle consists of his 1,2,3,4,5^{th} favorite courses, and worst bundle consists of his 2,3,4,6, 7^{th} favorite courses has a difference of (2+3+4+6+7) - (1+2+3+4+5) = 7

Figure 3: Average Rank Comparison
Approximate CEEI vs. HBS Draft Mechanism

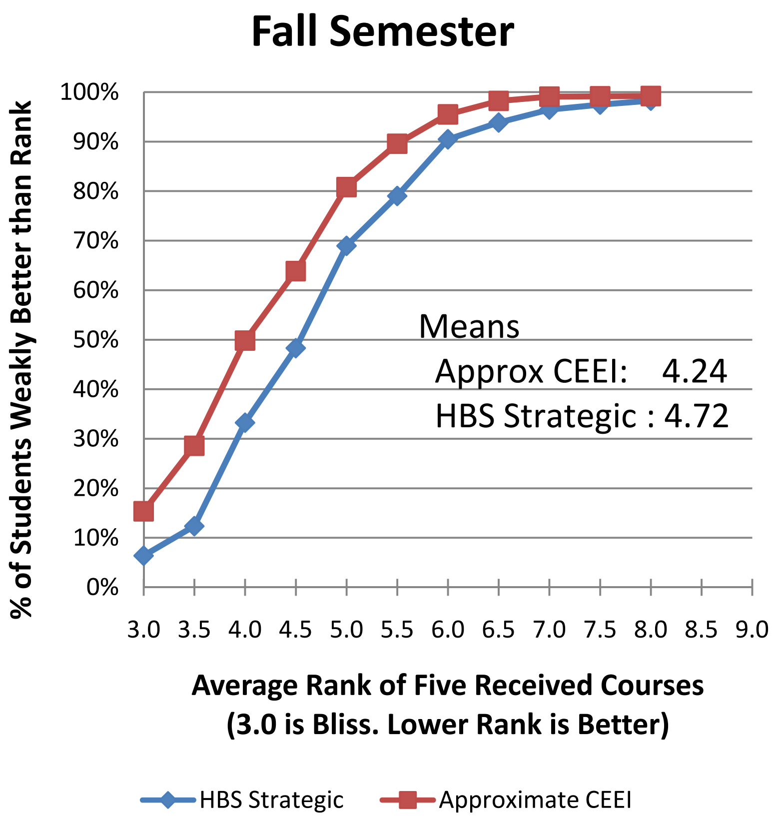

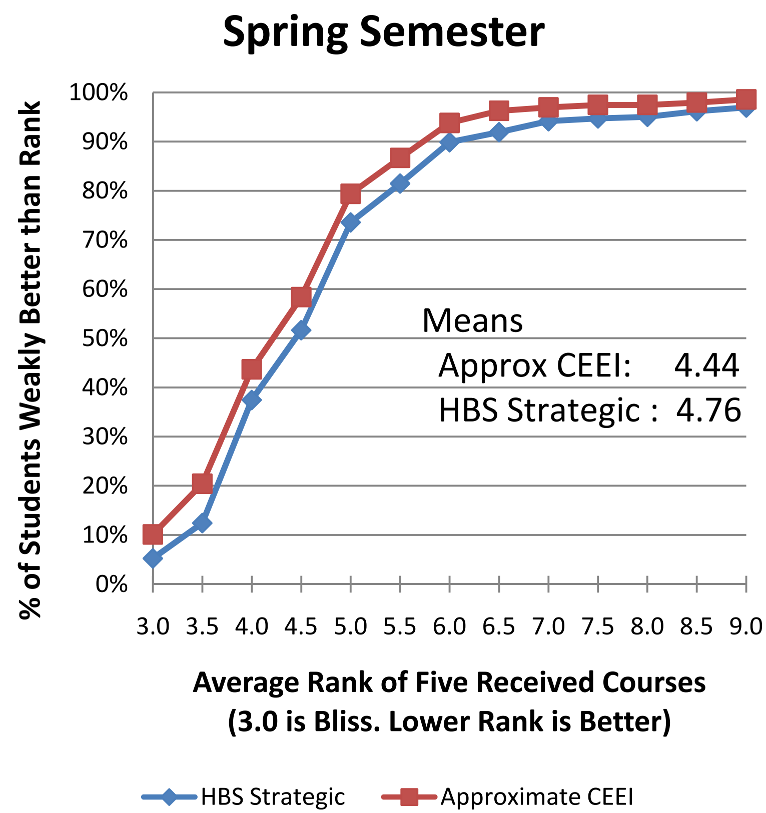

Description: The Othman, Budish and Sandholm (2008) Approximate CEEI algorithm is run 100 times for each semester of the Harvard Business School course allocation data (456 students, ~50 courses, 5 courses per student). Each run uses randomly generated budgets. For each random budget ordering I also 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*100 = 45,600 student-trial pairs. Average rank is calculated based on the student's true preferences. For instance, a student who receives her 1,2,3,4,5th favorite courses has an average rank of (1+2+3+4+5)/5 = 3.

Conclusion: Summary

- Combinatorial assignment is a problem of practical and theoretical importance
- Dictatorship theorems: there is no perfect mechanism
- I propose criteria that constitute an attractive compromise of competing objectives
  - Outcome fairness: Maximin Share, Envy Bounded by a Single Object
  - Incentives: Strategyproof in a Large Market
- I construct a specific mechanism that satisfies the criteria while maintaining approximate efficiency
  - Adapt the CEEI to an indivisible-goods environment
  - Existence requires approximating both "CE" and "EI"
- Performance compelling on data

Conclusion: Open Questions

- Are there classes of preferences with a better bound?
- How closely can we approximate EI if exact CE is a design requirement?
- Sequential problem: what if there are multiple sets of objects that cannot be allocated simultaneously?
- Two-sided preferences problem: what if professors or managers have preferences too?
- In what other environments is CEEI an attractive framework for market design?