Mechanism Design and Allocation Systems in Education

Irene Lo
Columbia University

May 3, 2017

1 Introduction

Context

The following are some of the important issues (and sub-issues) in education policy:

- What is the underlying objective of education, and how do we measure how well we are doing?
  - Welfare metrics and methods of measurement
  - The role of testing and test-based policy (ESSA)
  - Common core and standardized testing
  - Teacher evaluations and incentives
- Socio-economic integration of schools and equal opportunity:
  - Affirmative action, Charter schools with 'diversity missions'
  - Turning around failing schools (ESSA)
- School financing:
  - School vouchers (in light of current political events)
  - Unequal school funding (e.g. Connecticut in the New York Times)

Resources

- Stanford Education Data Archive

This Meeting

- We will discuss some of the context above.
- We will also brainstorm underexplored allocation systems in education.
- We look at two examples:
  - Course allocation at Wharton
  - School choice in Boston

2 High-Level Concerns

2.1 Welfare Metrics

We first ask, what are we trying to achieve in education? At a high level, we can say our objective is that students learn things so that they can have better lives.

One problem is, how do you measure welfare in this system? This brings us to the issue of testing and test-based policy. We currently have ESSA, which was signed by President Obama in 2015. ESSA is part of the push to have a common core and have standardized testing. The goal is to have a better way to keep people within the system accountable and to meet certain objectives. There has been a lot of debate about whether standardized testing is the correct way to evaluate quality of education. What is the common core teaching students? In terms of incentives, this raises the question of whether this incentivizes teachers to teach to the test instead of helping students to learn things. Another issue is what the scores are used for, e.g. to decide what schools get more funding or to determine teachers' salaries.

2.2 Socio-Economic Integration

There is a growing amount of work looking at socio-economic integration, including affirmative action, charter schools that have diversity missions, and so on. One line of work is whether these diversity initiatives are helpful to the students. On the issue of equal opportunity, one of the focal policies of ESSA is to identify under-performing schools and turn them around. From a policy perspective, equalizing opportunity by improving under-performing schools is a central issue.

2.3 School Financing

This has been brought back into the spotlight due to Betsy Devos, the current U.S. Secretary of Education. A question that has attracted a lot of attention and is a major policy concern is the question of whether school vouchers and school choice, in the sense of giving people subsidies to attend private schools, and whether this is a policy that benefits school children. We want to understand the impact of school voucher systems on students and families.

Related to school financing, there is also a lot of inequality in school funding. For example, a legal ruling in Connecticut resulted in a filed lawsuit claiming that the state was shortchanging the poorest districts in school funding, thereby increasing inequality.

3 Allocation Systems in Education

Mechanism designers working in the realm of education have largely focused on school choice. Here are some other allocation problems in education that are under-explored:

- Assigning teachers to schools and classes. The problem of assigning teachers to schools is not that salient in the US since the market is very decentralized. But, once you're within the school, this becomes an important problem.

- *Question:* Is the decentralized system normal in other countries outside of the U.S. and is it successful?

Answer: There is some literature on assigning teachers to schools in France. There's also some work in New York (link below). From conversations with some academics who work in this area, there is some sense that centralizing would be better, but there are political barriers to implementing that in the U.S.

  - There is work in economics that shows that there are teacher quality gaps (TQGs): teachers of different qualities sort, where better teachers want to teach better (i.e., high-achieving) students. This occurs both across schools and within schools.
    * A consequence is rising inequality, as the better students get better teachers, and thereby perform better. There is empirical evidence for this.
      · Low-quality teachers are not proportionally distributed (e.g., Clotfelter et al., 2005 [5]; Kalogrides et al., 2013 [8]; Lankford et al., 2002 [10])
      · There is inefficiency in terms of added value from having a good teacher (e.g., Goldhaber et al., 2015 [6]; Isenberg et al., 2013 [7])
  - Evidence of teacher 'trickle-down' chains:
    * High levels of teacher reassignment have negative impact on students (e.g. Ronfeldt, Loeb and Wyckoff [11]). What you see is that, when a teacher vacates a position, then the next best teacher moves up, then the next one also moves up, and so on. This creates a lot of movement and impacts a lot of students.
  - An attempt at centralization: NYC DoE Online Open Market Transfer System
    * The system is fairly broken: most listings are filled outside of the system, and not through the market. But, it is an attempt to centralize the system for teacher assignment. There are likely a lot of interesting questions to ask here about why this market is not working well and what can be done to improve it.
  - *Question:* Is there an obvious difference between this and resident matching? What makes teacher assignments special?

    Answer: If you're talking about assigning teachers to schools, then the length of time the teacher spends in the school system is much longer than the length of time a resident spends as a hospital resident. Another thing that's very salient is the concern of equality of opportunity, and we have to take that into consideration in these allocation, diversity-enhancing, and financing problems.

- Teacher residency assignments:
  - Examples: Teach for America, Urban Teacher Residency United, Sposato, NYC Teaching Fellows (NYC)
  - The length of stay is more comparable to hospital residency. It's not clear that the most important thing is to learn from people what they want and use that information to assign them, which seems to be what is done in hospital residency. In teacher residency, there are other, more global objectives which may be more important.
  - There is some work that shows that improving matching in teacher residency impacts retention.
- *Comment 1:* There is an interesting point in that if you assign a bad teacher to good students, that might end up being unpleasant for both the teachers and the students. Similarly, if you

assign bad teachers to bad students, then that exacerbates inequality. So, perhaps, a sensible solution would require that teachers meet some sort of competency and then allow teachers' and students' preferences to determine the assignments.

- *Comment 2:* One argument is that the education production function is supermodular, which would encourage this positive assortative matching. But, on the other hand, one of the goals of using matching mechanisms instead of prices is to create equality. We have public education for this reason, so even if a great teacher is less effective with bad students, they would still be more effective than a bad teacher.
- *Comment 3:* You have a finite set of teachers and you must assign them to students. You have some sort of quality measure for teachers based on how good they are, and a performance measure for the students. The outcome is determined by the teacher-student matching based on these values. You want to assign teachers to students. What would be a sensible objective here? Is it minmax or do you want some sort of aggregate measure? How do equality constraints and value-added constraints factor in? It's also not clear how you would determine this outcome based on the quality-performance pair.
- *Comment 4:* How does the heterogeneity of the classroom in terms of student performance impact teacher assignment? Do we want to have a more heterogeneous class so that good teachers have to teach a wider-range of students? Do you also end up having positive peer-effects if you have a more heterogeneous group? It does seem that teachers prefer more homogeneous classes in terms of performance. By having more heterogeneous classrooms, you might also have the ability to allow students to pick up skills related to working in groups with mixed-ability, which prepares them for the real-world, as opposed to a curated environment of equal ability.
- Classroom allocation:
  - Fair allocation of unused classrooms in public schools to charter schools (e.g. Kurokawa, Procaccia and Shah [9]).
  - Public schools have spare classrooms, perhaps due to enrollment, and charter schools need classrooms. These charter schools also have dichotomous preferences over the public schools (e.g., whether the public school is close enough or not). They also have a demand for the size of the classroom, and the public schools have a capacity they can handle. You want to assign each charter school to at most one school. (In theory, you could assign them to more, but anecdotally, charter schools seem to always reject these offers.) You want to assign charter schools to public schools while respecting proportionality, envy-freeness, pareto-optimality, and strategy-proofness. They implement this in a school system in California.
- Other assignment problems include:
  - Assigning TAs to classes.
  - Matching students with thesis advisors.
  - Assigning students to dormitories.

- One question to brainstorm is: how much transfer of ideas is there between matching teachers, as a human capital, and physical facilities like dorms and classes? Are these problems somehow linked outside of the fact that both of these are about education or do they somehow influence one another?

- *Question:* For assigning courses to schedules, is there a commonly accepted solution that is just not yet adopted by schools or is it an open question?

  Answer: I don't think there is a common solution to this, or that there is even a lot of literature on this. So, this is definitely an allocation problem to continue to think about.

- A problem somewhat linked to this is course recommendation-systems for MOOCs. There are some interesting questions that come up here such as, in Coursera, they found that women tend to not want to take STEM courses. So, should the system go with what the individuals want and not recommend STEM courses or try to correct for any biases that there might exist? What are the objectives here?

- Related to MOOCs, there is also the issue of determining the best sequence of material to get through some course, which might vary person to person. There is a bunch of internal research that Duolingo has done to determine what order to present a language concept to best aid individuals learning a language, which might be impacted by a language that you already know.

4 Course Allocation at Wharton

This is based on three papers by Eric Budish. We will talk about the theory and then discuss how they implemented it in practice.

In the course allocation problem, you have students who have preferences for courses that they want to take. You have constraints such as: (1) class capacity constraints, (2) simultaneous enrollment constraints, (3) prerequisite constraints, and so on. You want to assign students to courses respecting these constraints and optimizing certain fairness and efficiency objectives.

Theory: New solution concept (A-CEEI)

Budish [2]: "The Combinatorial Assignment Problem: Approximate Competitive Equilibrium from Equal Incomes" proposes the following solution:

- A-CEEI: Approximate Competitive Equilibrium from Equal Incomes

  Definition 1. *An* (α, β)*-CEEI* $(\mathbf{x}^*, \mathbf{b}^*, \mathbf{p}^*)$*:*
  x_i^* *= allocation of agent* i*,* b_i^* *= budget of agent* i*,* p_j^* *= price of item* j*, satisfying, for* q_j *= the capacity of item* j*,*

  1. $x_i^* = \max_{\succ_i}\{x \ : \ \mathbf{p}^* \cdot \mathbf{x} \leq \mathbf{b_i^*}\}$
  2. *(*α*-approximately CE)* $||\mathbf{z}^*||_\mathbf{2} \leq \alpha$*, where* z_j^* *is 'excess demand' for item* j*:*
     (a) $z_j^* := \sum_i x_{ij}^* - q_j$ *if* $p_j^* > 0$ *(excess demand at over-demanded)*
     (b) $z_j^* := \max\{\sum_i x_{ij}^* - q_j, 0\}$ *if* $p_j^* = 0$ *(no excess demand at under-demanded)*

3. *(β-approximately EE)* $1 \leq \min_i b_i^* \leq \max_i b_i^* \leq 1 + \beta$

- A-CEEI is a variant of CEEI. Under CEEI, each agent has equal income and they trade freely until a competitive equilibirum is reached. That is, you set prices and agents buy items they can afford and you want to find market-clearing prices. CEEI is not guaranteed to exist and it's not hard to show this.

- Note that there are two approximations: the α says that this is approximately an equilibrium and the β says the incomes are approximately equal.

- Properties of A-CEEI:

  – Existence result:

    Theorem 1. *For any $\beta > 0$ there exists a $(\sqrt{\sigma M}/2, \beta)$-CEEI, where*

    * *$\sigma = \min\{2k, M\}$,*
    * *k is the maximum number of items a single agent can own, and*
    * *M is the number of distinct items.*

    Moreover, for any $\mathbf{b}$ satisfying $1 \leq \min_i b_i \leq \max_i b_i \leq 1 + \beta$ and $\varepsilon > 0$, there exists a $(\sqrt{\sigma M}/2, \beta)$-CEEI with budgets $\mathbf{b}^*$ such that $||\mathbf{b} - \mathbf{b}^*|| < \varepsilon$.

  – Guarantees each agent an approximate maximin share:

    Theorem 2. *If $(\mathbf{x}^*, \mathbf{b}^*, \mathbf{p}^*)$ is an (α, β)-CEEI where*

    * *at price $\mathbf{p}^*$, goods endowment costs at most $(1 + \delta)$ times budget endowment for some δ (i.e. $\sum_j p_j^* q_j \leq \sum_i b_i^*(1 + \delta)$)*
    * *β is sufficiently small, $\beta < (1 - \delta N)/N(1 + \delta)$ (where N is the number of agents)*

    then $\mathbf{x}^*$ satisfies the $(N + 1)$-maximin share guarantee, i.e.

    $$x_i^* \succ_i \max_{\succ_i}\{\min_{\succ_j}\{x_1, x_2, \ldots, x_N, x_{N+1}.\}\}$$

  – Guarantees that envy is bounded by a single good:

    Theorem 3. *If $(\mathbf{x}^*, \mathbf{b}^*, \mathbf{p}^*)$ is an (α, β)-CEEI with $\beta \leq 1/(k - 1)$, where k is the maximum number of items a single agent can own, then for all agents i, i', removing at most one good from agent i''s allocation eliminates any envy that agent i has for i'.*

  – Strategy-proof in the large:

    * When agents are price takers, it is a dominant strategy to report truthfully. This is important in contrast to course auctions. Previously, they used to have bidding mechanisms similar to CEEI. This was not strategy proof.
    * In large markets, agents are approximately price takers.

- *Question:* In Theorem 1, why can't we just choose β arbitrarily close to 0? What are the consequences for doing so?

  Answer: There might be computational tradeoffs. It seems like you really can set β to be any value, but you can't have it be 0. This is highlighted in Theorem 2, which says that you can perturb any budget even by a small amount and still find an equilibrium.

- *Question:* Why do you need α? If everyone has a different budget, then the problem would be solved.

  Answer: If you want the claim to be true for all β, then you would have to relax something else. You can get something with $\alpha = 0$ but with high budget inequality.

- Note that Theorem 2 also guarantees max-min share. Suppose an agent goes first and they split each course into $N+1$ different courses, where N is the worst one out of those. Suppose you get the worst out of those. The endowment that you get in A-CEEI is at least as good as that.

- *Question:* Is this a full-information equilibrium or is it enough just to know the prices? In the non-large market version where you can potentially affect the price, are you bidding taking into account what you expect others to bid, or do you know exactly what they're reporting?

  Answer: They don't have any results on what happens with strategy-proofness when you can affect prices. There is some discussion in the paper but not formal results.

- Proposed mechanism:

  1. Each student i reports preferences over permissible schedules.
  2. Assign student i a budget $b_i^* \sim U[1, 1+\beta]$, with $0 < \beta < \min\{1/N, 1/(k-1)\}$ (where N is the number of students and k is the max number of courses a student can take).
  3. Compute a $(\sqrt{\sigma M}/2, \beta)$-CEEI $(\mathbf{x}^*, \mathbf{b}^*, \mathbf{p}^*)$.

- There are some issues with implementing the proposed mechanism, which we discuss in the next section.

Buy-in: Lab experiments

Budish and Kessler [4]: "Can Agents 'Report Their Types'?: An Experiment that Changed the Course Allocation Mechanism at Wharton"

- Problem description:
  $N \approx 1700$ students, $M \approx 300$ courses, students request up to $k \approx 8$ courses.

- Issue: Reporting preferences over all 2^M (or even just $\binom{M}{k}$) permissible schedules is infeasible
  .

  - Is there a simple preference-reporting language where agents are able to report their preferences 'accurately enough'?
  - The language they use assumes additive cardinal utility, with pairwise substitutes and complements.

    1. Favorite course $\leftarrow$ value $= 100$
       All other courses $\leftarrow$ value in $[1, 100]$

2. Assign extra positive (or negative) value to certain pairs of classes

Figure 1: Screenshot of CEEI User Interface

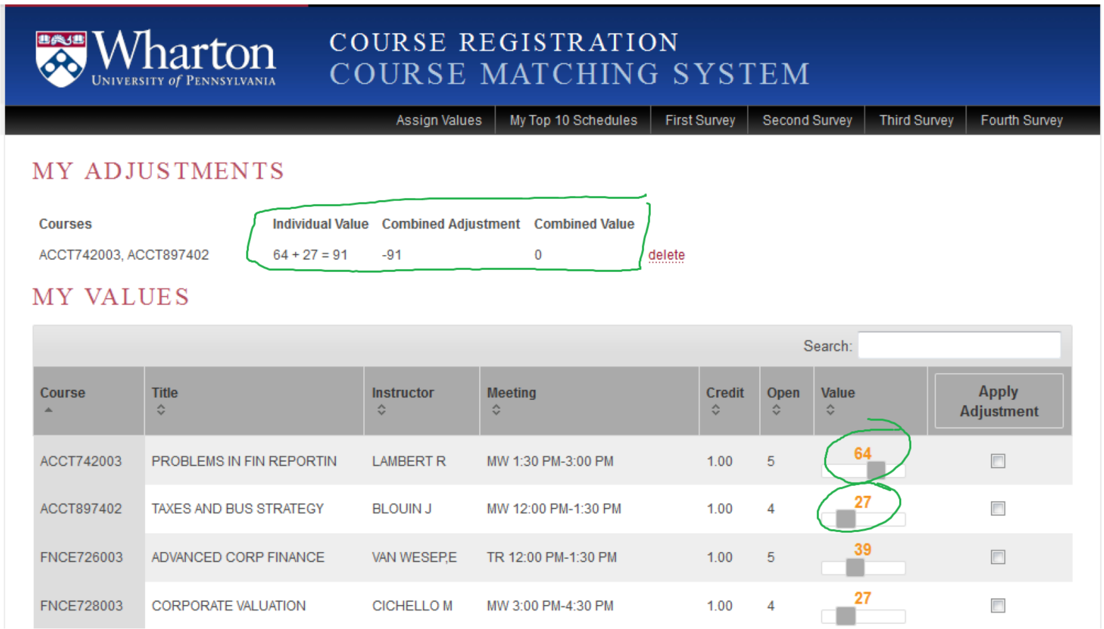

3. A high-level question is whether this is a reasonable preference-reporting language. One way they checked this is using the top-ten widget. Once you put in your preferences, it shows you your top 10 schedules, shown in Figure 2, and you could check whether this is in line with your preferences.

Figure 2: Screenshot of Top Ten Widget

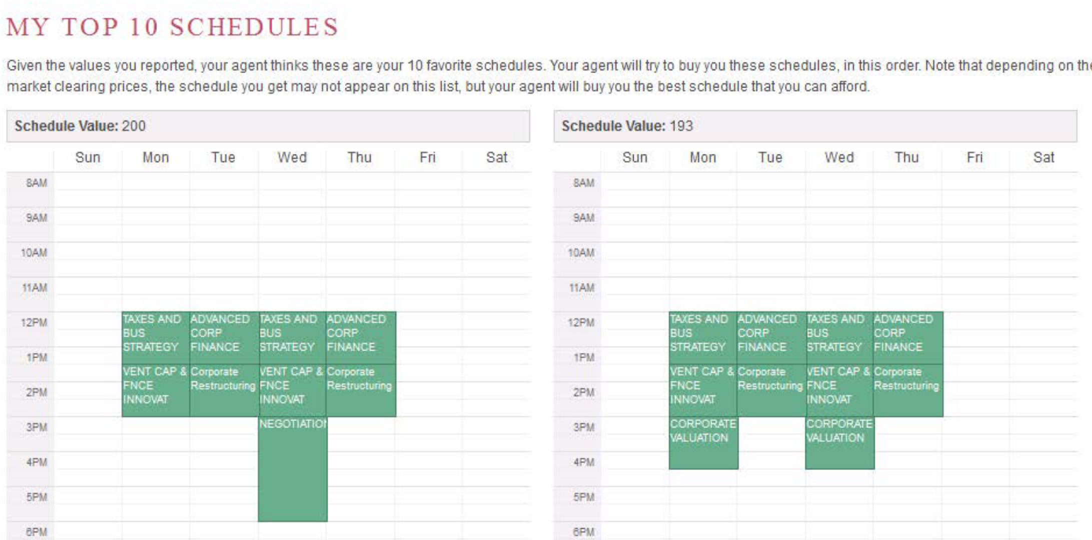

– **Remark**: This was used in the actual implementation of A-CEEI in CourseMatch at Wharton. One notable change is that the 'top ten' widget was modified to show more than the top ten schedules since not all students were being assigned schedules in their top 10, so you want to check that the reporting is accurate even further down the list.

- *Question:* Are there priorities for students (say who are about to graduate)?

  Answer: In the equilibrium concept, there is no priority, but you can skew the budgets of students you want to prioritize in implementation.

- Properties of A-CEEI: Efficiency, fairness, strategy-proof in the large, simplicity.
  - **Experimental Design:** (8 sessions of 14-19 students each)
    1. Study Instructions:
       * 25 Wharton spring semester course sections.
       * Select schedule of 5 courses according to the individual's own preferences. Note, this is different from the norm where students would be told what their preferences are for the purposes of the given experiment.
       * Note, these students are already used to bidding in complex systems like the previous allocation mechanism used by Wharton.
    2. They test two mechanisms: The Wharton course auction, which was the initial system used, versus the CEEI course matching mechanism
       * Half of sessions ran the Auction first, other half ran CEEI first.
       * Give the instructions, run the mechanism, survey the users about their experience with the mechanism.
    3. Binary comparisons:
       * Users were asked to report which of the two schedules they prefer, and the strength of their preference.
       * The first and last binary comparisons were the schedule they would have been given under CEEI and the Auction (with the order reversed).
       * In between, they did ≈ 12 comparisons on envy (fairness): the student's own CEEI with another's CEEI, the student's own Auction with another's Auction.
       * 5 comparisons on accuracy of preferences were captured through this language: The student's own CEEI with another CEEI under 10% or 30% higher or lower budgets.
    4. General survey questions:
       * How did the student feel about strategic simplicity and overall satisfaction?
       * How did the student feel about transparency and understanding?
  - **Efficiency:**

Table 3: Allocative Efficiency — Results from Binary Comparisons of CEEI Schedule and Auction Schedule

Session	Subjects in the Session	Prefer CEEI	Prefer Auction	Indifferent	Identical	Majority Voting Result
1	18	8	5	4	1	CEEI
2	15	5	5	2	3	Tie
3	19	10	4	3	2	CEEI
4	16	5	4	3	4	CEEI
5	18	8	8	2	0	Tie
6	14	6	5	1	2	CEEI
7	18	8	6	2	2	CEEI
8	14	6	5	0	3	CEEI
All	132	56	42	17	17	6-0-2

Table 4: Results on Pareto Efficiency: Reported Preferences

	Auction	CEEI	Probability Ratio Test (one-sided)
# of Pareto-improving trades detected (% of course seats)	260 (32.8%)	44 (5.6%)	$p < 0.001$
# of students involved in at least one trade (% of students)	98 (74.2%)	22 (16.7%)	$p < 0.001$

- You can see that it's not a clear domination, but if you had to do a majority voting, every session would have CEEI as the winner. You can think of this as, CEEI is more fair in that Auction would give a small number of students a really good schedule in their opinion, but the majority would prefer CEEI.
- **Fairness:**

Table 1: Envy Under CEEI and Auction — Joint Test of the Mechanism and the Reporting Language, using Binary Comparisons

	Auction	CEEI	Probability Ratio Test (one-sided)
Panel A: By Subject (Auction: N = 119, CEEI: N = 117)			
% of subjects who display any envy of another subject's schedule	42%	31%	$p = 0.036$
% of subjects who display any large envy ("prefer" or "strongly prefer") of another subject's schedule	34%	21%	$p = 0.008$
Panel B: By Comparison (Auction: N = 499, CEEI: N = 475)			
% of comparisons in which the subject displays any envy of the other subject's schedule	19%	12%	$p = 0.002$
% of comparisons in which the subject displays any large envy ("prefer" or "strongly prefer") of the other subject's schedule	14%	8%	$p = 0.002$

- We see that the amount of envy is reduced when you go from Auction to CEEI.

Table 2: Envy Under CEEI and Auction — Isolated Test of the Mechanism, Assuming Perfect Preference Reporting

	Auction	CEEI	Probability Ratio Test (one-sided)
Panel A: By Subject (Auction: N = 119; CEEI: N = 117)			
% of subjects who envy another subject's schedule according to CEEI	29%	4%	$p < 0.001$
Panel B: By Comparisons (Auction: N = 499; CEEI: N = 475)			
% of comparisons in which one subject envies the other subject's schedule according to CEEI	15%	2%	$p < 0.001$

Table 2 reports the envy results based on the utility measure constructed from the CEEI preference reports.

- We can see that if the preferences are perfectly reported, then CEEI does even better.
- **Strategy-proof in the large, simple:**

Table 10: Qualitative Responses About CEEI and Auction

Questions about each System	CEEI Average	Auction Average	Wilcoxon p-value
Panel A: Questions Regarding Strategic Simplicity and Overall Satisfaction			
I enjoyed participating in this course allocation system.	4.72	4.37	$p = 0.095$
I like this course allocation system.	4.55	4.18	$p = 0.095$
This course allocation system is simple.	4.45	3.73	$p = 0.001$
I had to think strategically about what other students would do in this course allocation system.	2.93	6.42	$p < 0.001$
Someone with perfect knowledge of the historical supply and demand for courses could have had an advantage over me in this system.	3.67	6.04	$p < 0.001$
Panel B: Questions Regarding Transparency and Understanding			
I understand how this course allocation system works.	4.83	5.92	$p < 0.001$
I felt like I had control over my schedule in this course allocation system.	3.95	4.45	$p = 0.073$

- There are differences in strategy-proofness and simplicity, e.g., look at differences in "I had to think strategically about what other students would do in this course allocation system." (The scale is 1-7.)

Implementation: Engineering Issues

Budish, Cachon, Kessler and Othman [3]: "Course Match: A Large-Scale Implementation of Approximate Competitive Equilibrium from Equal Incomes for Combinatorial Allocation"

- Main obstacles in implementation:
  - Reporting complexity
  - Computational complexity
  - Violations of capacity constraints
- *Question:* Presumably, every course allocation procedure has some over/under-subscription issue. Is this a specifically an A-CEEI problem? Specifically, for the under-demand, if no one wants to take a course, then this will be an issue regardless of the mechanism you use.

  Answer: Other mechanisms have the constraints more baked in to the system. Here, the professors are giving you a target capacity and a hard capacity. This procedure might violate the hard capacity because you're giving it some lee-way.

  For the issue of under-demand, the type that they are trying to solve is if there are students who like a course but the final allocation has the number of students enrolled under the target capacity. They solve the over/under-demand problem in different ways, but we don't go into that here.
- Algorithm Description:

1. Price vector search:
   – Gradient neighbor (tatonnement) and individual adjustment neighbors.
   – 'Tabu' search - remove price vectors that do not produce new course allocations.
   – Not hill climbing: allow up to five consecutive steps where there is no improvement in the best clearing error.
2. Eliminate over-subscription:
   – Raise the price of the most over-subscribed course to eliminate half of over-subscription.
   – Binary search to find the price.
3. Reduce under-subscription:
   – Increase each student's budget by 10%.
   – Iteratively: select a student, form a new choice set from the student's initial allocation and set of courses that have open seats, assign the student the best schedule they can afford from new choice set.

- Overcoming reporting complexity:
  – Same reporting language as in [4].
- Overcoming computational complexity:
  – Parallelizing searches.
- Overcoming capacity constraints:
  – Stages for eliminating over-subscription and reducing under-subscription.

Table 2. Stagewise results from the spring 2014 Course Match production run

	Stage 1: Price vector search α^2				Stage 2: Oversubscription elimination α^2				Stage 3: Undersubscription reduction α^2			
Compute server	+	−	Seats	Loss (%)	+	−	Seats	Loss (%)	+	−	Seats	Loss (%)
1	51	40	32	0.19	0	141	67	0.80	0	31	17	0.07
2	23	28	24	0.20	0	205	97	1.26	0	30	16	0.07
3	75	35	25	0.31	0	138	74	0.96	0	21	17	0.07
4	36	125	71	0.68	0	228	106	1.00	0	83	33	0.33
5	59	48	36	0.28	0	201	87	1.15	0	52	24	0.09
6	53	48	32	0.22	0	186	86	0.87	0	32	22	0.07
7	47	56	42	0.54	0	202	86	1.13	0	17	13	0.02

- Efficiency and fairness:

Table 3. Quantitative summary statistics from the fall 2013 and spring 2014 allocations

	Fall 2013	Spring 2014
Students	1,650	1,700
Courses	285	344
Courses with maximum above target capacity, i.e., $q_j < \hat{q}_j$	78	262
Total capacity overhead $(\sum_j(\hat{q}_j - q_j)/\sum_j q_j)$	0.8%	2.1%
Number of courses with demand above target capacity	13	49
Fraction of total capacity allocated	71%	74%
Number of courses with a positive price	154	199
Number of undersubscribed positive price courses	15	18
Deadweight loss as percentage of economy	0.19%	0.07%
Highest course price as fraction of average budget	1.31	0.88

Table 6. Gini coefficients of student allocations

Mechanism	Semester	First-year students	Second-year students	All students
Auction	Fall 2012	0.33	0.36	0.54
Course match	Fall 2013	0.13	0.22	0.32
Auction	Spring 2013	0.25	0.39	0.34
Course match	Spring 2014	0.10	0.12	0.15

Note. Fall 2012 allocations are taken as round 5 of the auction.

- It's interesting to note that this solution could be applied to other domains as well. One close match would be assigning program committee members to papers for review. You have a lot of the same ingredients.
- It's surprising also how well they're able to do with these additive utilities. If we think about the domain, if students are in their last year of their program, there might be a lot of complements and substitutes imposed by graduation requirements, but the paper shows that these pair bids show up rarely. It might be because we're in a business school setting, but it's still surprising.
- *Question:* Even with the simplifications that they make with bidding, it is still somewhat complicated. Has there been thought given to seeding these with natural complementaries or substitutes or using recommender systems?

  Answer: You would need data to figure this out, so a good place to do this might be on MOOCs.

5 School Choice in Boston

- **Theory and Implementation**:
  Ashlagi and Shi [1]: "Optimal Allocation Without Money: An Engineering Approach"
  Shi [13]: "Assortment Planning in School Choice"
  - Theory: Lottery-plus-cutoff mechanism in large market approximation.
  - Utility model: multinomial logit

    $$u_{is} = \kappa_s - d_{ts} + \omega \text{Walk}_{ts} + \beta\epsilon_{is}$$

    for student i with geocode t at school s, where κ_s is the quality of school s, d_{ts} is distance, Walk_{ts} is an indicator for walk zone, and ϵ_{is} is idiosyncratic match value. κ_s, ω, β can be estimated from historical data.
  - Objective: A linear combination of average utility and max-min welfare

    $$W = \alpha \sum_t \left(w_t v_t + (1-\alpha) \min_{t'} v_{t'} \right)$$

    where t is the geocode, w_t is the proportion of students who live in geocode t, v_t is the expected utility of a student from geocode t.

- Computation:
  * Relaxation of 'nested menus' formulation gives an LP.
  * The LP has $exp(|S|)$ number of variables.
  * The dual can be solved efficiently in $poly(|S|, |T|)$ time and gives nested menus.
- Alternatively: use an assortment optimization framework.

- **Buy-in**: Interactions with BPS
  Shi [12]: "Guiding School-Choice Reform through Novel Applications of Operations Research"
  - Key steps in obtaining buy-in from BPS:
    1. Feb - Nov 2012: BPS external advisory committee (EAC) meetings.
    2. October 2012: Presented optimization-based proposal.
    3. November 2012: Presented simplified optimization-based proposal.
  - Solution approach:
    1. Build a demand model from previous choice data
    2. Precisely define outcome metrics:
       * Minimum access to quality schools (equity of access to quality)
       * Median expected number of neighboring students co-assigned (community cohesion)
       * Minimum access to capacity (supply and demand shortage)
       * Median access to dream school (availability of desirable choice)
       * Average walk distance (proximity to home)
       * Median expected choice rank (predictability of assignment)
       * Standard deviation across students of expected percentage of peers of each race or lunch status (racial and socioeconomic segregation)

       Build a simulation engine
    3. Work with BPS to propose a set of plans with small choice menus
       * Home-based A plan:
         · All schools in 1-mile walk zone
         · 2 closest schools in top 25% (BPS rank)
         · 4 closest schools in top 50% (BPS rank)
         · 6 closest schools in top 75% (BPS rank)
         · 3 closest schools with excess capacity
       * Home-based B plan:
         · All schools in 1-mile walk zone
         · 3 closest schools in top 25% (BPS rank)
         · 6 closest schools in top 50% (BPS rank)
         · 9 closest schools in top 75% (BPS rank)
         · 3 closest schools with excess capacity
       * BPS plans: Status quo (3-zone), 10-zone, 11-zone

4. Detailed analysis of plans to BPS and EAC

Plan	Status quo	10-zone	Modified 11-zone	Home-based-A	Home-based-B
Minimum access to quality					
Top 33% of schools (%)	14.6	0.7	0.9	12.7	13.2
Top 50% of schools (%)	19.5	22.7	22.7	22.5	27.8
Top 67% of schools (%)	31.4	22.7	22.7	32.9	32.9
Median no. of coassigned neighbor students	3.12	3.84	4.20	3.92	3.80
Minimum access to capacity (%)	47.2	31.8	31.9	36.0	35.6
Median access to dream school (%)	42.4	32.0	31.4	31.0	32.1
Average distance (miles)	2.03	1.24	1.19	1.25	1.30
Median rank obtained	2.83	2.65	2.62	2.78	2.91
Standard deviation across students					
% peers: Free lunch	9.0	11.5	11.5	11.2	11.1
% peers: Black	14.1	15.4	15.5	15.5	15.4
% peers: White	9.4	12.2	12.2	11.2	11.0
% peers: Hispanic	17.7	19.2	19.3	18.5	18.2
% peers: Asian	9.0	10.0	10.0	10.3	10.4
Bus coverage area per school (sq. miles)	24.71	6.96	6.41	6.58	8.63

Table 1: The table shows simulated performance of the assignment plans that the EAC considered for its final vote in February 2013.

References

[1] Itai Ashlagi and Peng Shi. Optimal allocation without money: An engineering approach. *Management Science*, 62(4):1078–1097, 2015.

[2] Eric Budish. The combinatorial assignment problem: Approximate competitive equilibrium from equal incomes. *Journal of Political Economy*, 119(6):1061–1103, 2011.

[3] Eric Budish, Gérard P Cachon, Judd B Kessler, and Abraham Othman. Course match: A large-scale implementation of approximate competitive equilibrium from equal incomes for combinatorial allocation. *Operations Research*, 2016.

[4] Eric B Budish and Judd B Kessler. Changing the course allocation mechanism at wharton. 2014.

[5] Charles T Clotfelter, Helen F Ladd, and Jacob Vigdor. Who teaches whom? race and the distribution of novice teachers. *Economics of Education review*, 24(4):377–392, 2005.

[6] Dan Goldhaber, Lesley Lavery, and Roddy Theobald. Uneven playing field? assessing the teacher quality gap between advantaged and disadvantaged students. *Educational Researcher*, 44(5):293–307, 2015.

[7] Eric Isenberg, Jeffrey Max, Philip Gleason, M Johnson, J Deutsch, and M Hansen. Do lowincome students have equal access to effective teachers. *Evidence from*, 26, 2016.

[8] Demetra Kalogrides, Susanna Loeb, and Tara Béteille. Systematic sorting: Teacher characteristics and class assignments. *Sociology of Education*, 86(2):103–123, 2013.

[9] David Kurokawa, Ariel D Procaccia, and Nisarg Shah. Leximin allocations in the real world. In *Proceedings of the Sixteenth ACM Conference on Economics and Computation*, pages 345–362. ACM, 2015.

[10] Hamilton Lankford, Susanna Loeb, and James Wyckoff. Teacher sorting and the plight of urban schools: A descriptive analysis. *Educational evaluation and policy analysis*, 24(1):37–62, 2002.

[11] Matthew Ronfeldt, Susanna Loeb, and James Wyckoff. How teacher turnover harms student achievement. *American Educational Research Journal*, 50(1):4–36, 2013.

[12] Peng Shi. Guiding school-choice reform through novel applications of operations research. *Interfaces*, 45(2):117–132, 2015.

[13] Peng Shi. Assortment planning in school choice. Working paper, 2017.