

Market Design Final Paper: Two-Sided Course Allocation

Gideon Moore

Fall 2019

Abstract

I develop a modified version of Budish (2011)'s *Approximate Competitive Equilibrium from Equal Incomes* (or A-CEEI) which allows school administrators to correct market failures unaddressed in the original mechanism by taxing student enrollment. This system preserves A-CEEI's incentive compatibility and lack of congestion while maintaining reasonable fairness and efficiency bounds, and adds the ability for matches to contain a degree of two-sidedness. While the lack of an outside good prevents computation of an optimal tax rate, I outline the impact of various types of tax schema changes, as well as advise on how to correct tax rates when presented with an undesired equilibrium. Finally I derive a boundary on market-clearing error when computing this allocation given a set tax schema. As far as I am aware, this is the only mechanism developed which is incentive compatible, uncongested, and two-sided that also provides bounds on fairness and efficiency losses.

1 Background

1.1 The Problem

Each term, students must vie against their peers for seats in overbooked courses. However, universities struggle to create a system which allocates courses both efficiently and fairly. While a price-based

market would likely be efficient here, universities and students alike are hesitant to implement varying tuition across courses. Further, students' preferences over courses are nuanced, and rife with both complementarities and substitutabilities. Thus, traditional deferred-acceptance based approaches to matching without transfers are liable to fail. These two issues combined make course allocation a difficult problem to solve. Many institutions have built their own mechanisms for allocating courses, but all either create incentive issues or are prone to unravelling.

When considering mechanisms for matching courses to students, there are several characteristics which are desirable. An ideal market would be:

- Incentive Compatible—Students should feel safe revealing their true preferences over courses.
- Uncongested—Students should be able to gather all the relevant information about classes and submit their requests over a reasonably calm time period.
- Efficient—The mechanism should allocate courses in a way that maximizes social welfare.
- Fair—Students should receive course bundles which are approximately equally desirable.
- Two-Sided—Just as students may prefer certain courses, courses may prefer certain varieties of students, and thus may wish to incentivize their enrollment.

The last of the above, two-sidedness, is often neglected in the course matching literature, instead treating courses as indifferent objects. However, there are many reasons why professors or administrators may prioritize certain students, and those priorities may be idiosyncratic over courses. For example, administrators may desire that a senior who requires a course to graduate be given registration priority over sophomores who will have opportunities later in their careers. On the other hand, a creative writing department may wish to prioritize freshmen, who could go on to major in the field, over upperclassmen from other departments taking the subject as an elective. In addition, it is possible students may impose externalities on others. For example, a freshman enrolling in an upper level seminar may require a disproportionate quantity of instructor attention at the expense of their more senior peers.

1.2 Current Solutions

1.2.1 First-Come, First-Served

Many universities employ a simple “first-come, first-served” approach, opening registration simultaneously for all students and barring registration for further students once enrollment hits the cap. However, these markets suffer extreme congestion immediately after the start time, leading students to pursue measures such as writing scripts to enroll in their courses for them.¹ Thus, courses are not allocated to those who need them most, but instead to the students who are most capable of automation.

To satisfy this demand, startups such as “BU Ninja” are now licensing their automation services to many students for a fee,² collecting rents off of the market frictions present in this system. In addition, the congestion in the system undermines the incentive compatibility; if a student anticipates a risk of not getting a highly desired course, and the time required to attempt to enroll in this course would allow other courses they prefer to fill, they may not bother attempting to enroll in the highly-desired course and instead focus on the less-desirable courses. Lastly, this system is in no way two-sided, as course slots are simply taken by the students who get their first.

1.2.2 Lotteries

Other schools use variants of course lottery systems, where any number of students may attempt to enroll for a course, but only a subset of randomly selected students are then enrolled. Unmatched students then either enter successive lottery rounds or enter a free-for-all add-drop period. This system has some advantages over the first-come first-served system above. First and foremost, it prevents the rampant congestion, as students may input their preferences at any point over a pre-determined interval prior to a centralized lottery date. Further, with some modifications this system can be made two-sided, as administrators can request that certain classes of students all be placed before opening the lottery to students outside this class.

¹McKenzie, Lindsay (2019). “A Class Registration Bot Backfires.” <https://www.insidehighered.com/news/2019/01/15/computer-program-automatically-registers-students-classes-has-unintended>

²<https://buninja.com/>

However, lotteries are not without their flaws. This system does not satisfy incentive compatibility, as students must consider what their “realistic” first choice course is, as otherwise their second choice might fill in the first round. Further, it obscures for school administrators which courses are in greatest demand, as students who choose not to enter the lottery for their desired courses are never observed by the system. In addition, while ex-ante fair, it is possible for one student to win every lottery and another to lose, making the potential allocations quite unfair ex-post.

This system also struggles to handle complementarities and substitutabilities. Given two complementary courses, it is possible a student will win the lottery for one but not the other, making the course the student “won” no longer desirable. Alternately, a student who strongly desires one course in a heavily enrolled department may enter and win multiple lotteries, but only actually desire one of the relevant courses.

1.2.3 Random Serial Dictatorship

The simplest canonical example meeting our incentive compatibility, uncongestedness, and efficiency constraints is a repeated serial dictatorship. In an arbitrary order, students may choose their ideal course schedule from the remaining available seats. Under this mechanism, each student on their turn has no reason to choose a bundle other than their most preferred, and each actor may take as much time in their choice as desired. Further, given each student takes their most preferred bundle available, you cannot improve any student’s bundle without making another student worse off, meaning the mechanism is guaranteed to be Pareto efficient. This mechanism is also able to handle complementarities and substitutabilities, unlike those above; if any parts of a complement are unavailable, the student simply will not choose any, and no one will choose multiple courses within the same substitution group.

Unfortunately, this system is extremely unfair—it is possible for one student to receive a schedule of nothing but overbooked courses, while another student receives none of their desired courses. Lastly, this mechanism is wholly single-sided; since students choose whatever schedule they desire, there is no capacity for courses to influence the allocation.

1.2.4 Draft

One way to make the serial dictatorship more fair is to have students choose one course at a time rather than their whole bundle at once. This mechanism is known as a “draft,” and is studied in the context of Harvard Business School by Budish and Cantillon (2012). Under the HBS draft, the authors found students received their first choice course 63% of the time, compared to 49% under a serial dictatorship. Further, mean average course rank fell from 8.74 under the dictatorship to 7.99 under a draft.

However, the draft has many problems of its own; first and foremost, it is no longer strategyproof. Under a draft, students must decide whether to take their first choice course first, or choose a more-popular, less-preferred course in hopes their true first choice will come back around the table. This introduces two avenues for deadweight loss. First, the most popular courses face even higher demand, making the students who desire those courses worse off. Second, it is possible students miscalculate whether their favorite course will still be open on their following turn, meaning they end up with a less desirable course despite the fact the more desirable course was at one point available. Thus, the non-strategyproofness of the mechanism imposes real costs on the students.

This mechanism also struggles with complementarities; in order to get both parts of a complement, a student must choose the first course before they are guaranteed enrollment in the second course, with no guarantee the second course will still be available during their next turn. Lastly, there is no course agency present in this mechanism; as students act as a dictator on their turn, courses have no input over who enrolls.

1.2.5 Point Auction

Rather than adapting a dictatorship mechanism, several business schools such as Haas and Kellogg use point auction systems. Under this design, students are allocated equal quantities of points, then use those points to bid on courses against their classmates. Advocates of this system claim the resemblance to a market causes the course to be efficient. However, Sönmez and Ünver (2010) shows this mechanism is inefficient due to the fact bids are expected to serve a dual purpose: they are both an indicator of desire and an allocation mechanism.

Consider a course expected to have relatively low demand. Under this system, students should bid as little as possible for this course in order to preserve their points for more desirable courses. However, this is true *regardless* of how much students may desire this course. In the event demand for the course is higher than expected, the student will have lost out on their first choice course. This difficulty is worsened by the existence of complementarities and substitutabilities among courses. Students cannot adjust their bids conditional on their enrollment in other courses, making it likely for students to receive only one half of a complement or multiple versions of a substitute. A tax mechanism could theoretically make this mechanism two-sided; however, at the moment I am not aware of any universities which use such a system.

1.3 Second Welfare Theorem

According to the Second Welfare Theorem as defined in Arrow et al. (1951), markets can achieve any Pareto efficient assignment through the use of prices assuming that agents are locally non-satiated, as is the case in a world with cash transfers. However, the lack of money in the world of course assignment makes it unclear whether the Second Welfare Theorem still applies.

According to Miralles and Pycia (2014), the Second Welfare Theorem remains applicable in many-one matching markets with no transfers, despite the potential for satiation and lack of currency. They claim that “in market design contexts, our characterization of efficient assignments allows one to restrict attention to price mechanisms at least in settings, such as large markets, where such mechanisms are incentive compatible” (Miralles and Pycia 2014, p. 2).

By breaking each of our many-seat courses into several one-seat courses, it is trivial to transform the course allocation problem into a form of the school choice problem they use as their example such that each seat is assigned to only one student, but each student may be assigned many seats. Thus, conditional on the existence of a strategyproof, price-based mechanism—which I develop in section 2.2—I claim that this mechanism can reach any Pareto-efficient allocation of courses through sufficient budget variation.

2 The Math

2.1 The combinatorial assignment problem

An instance of the combinatorial assignment problem is defined by the following parameters:

- \mathcal{S} —The set of students. Let $N = |\mathcal{S}|$.
- \mathcal{C} —The set of courses. Let $M = |\mathcal{C}|$. A schedule for student i is thus defined as vector $x_i \in \{0, 1\}^M$ where each index j of the vector is a Boolean indicating enrollment in course j .
- (q_j) —A vector defining the size of each course j .
- (\succ_i) —A strict preference relation for each student over the bundles of all courses.

Given the above, an allocation $(x_i)_{i \in \mathcal{S}}$ assigns a schedule x_i for each student. The sole restriction on this problem is that given an allocation, courses may not exceed their capacity; that is, $\forall j \in \mathcal{C} : \sum_{i \in \mathcal{S}} x_{ij} \leq q_j$.

Students' preferences are defined over bundles, and may include both complementarities and substitutes. This is reflective of real markets for courses, which are full of these interactions. For example, a student may not wish to take four classes in the same department in one semester, but may be happy with any of them individually. Alternately, a senior two classes short of a minor may derive great utility from any pair of courses in the given department, but would derive much less given each course individually. Thus, the objective value $\sum_{i \in \mathcal{S}} U_i(\vec{x}_i)$ cannot be broken down further into utilities of individual course enrollments.

To use a competitive equilibrium approach to this problem, we can endow each student with a budget b_i . From there, when given ordered preferences, we can use the algorithm developed in Othman, Sandholm, and Budish (2010) to generate market clearing prices, then assign each student their most preferred affordable bundle. This is the approach taken in Budish (2011). Given all benefits and costs of enrollments are realized privately, this approach generates an outcome that is both stable and efficient from the perspective of each student.

2.2 Markets with Externalities

To expand the problem defined above, suppose that each student-course pair imposes some externality s_{ij} onto society. These externalities may be positive or negative. For example, having an older student in a course may be beneficial for their classmates if the older student can explain concepts to others. Alternately, if an older student is taking up seats in an introductory course, the department may prefer a first year student with the hope the first year will go on to major in the field. In this case, social utility can be modeled as

$$\sum_{i \in \mathcal{S}} \left(U_i(\vec{x}_i) + \sum_{j \in \vec{x}_i} s_{ij} \right)$$

A social planner, then, would wish to maximize the above sum rather than simply private utility as in 2.1. In a traditional market, a Pigouvian tax equivalent to the externality could be imposed on each transaction, bringing students' incentives into line with the rest of society. However, this approach relies upon the existence of an outside good.

When implementing a competitive equilibrium, it is important to note that the currency used has no value when saved; anything not spent is lost. This lack of outside option causes prices to be severed from marginal utility. Instead, price is simply used to define a choice set. Consider a student who is somewhat, but not particularly, interested in a very expensive course c . Suppose further that the rest of the student's schedule is relatively inexpensive, meaning they can still afford c . In a market with outside goods, this student would likely not enroll in c , as the price is high and the utility of enrollment is low. However, because leftover currency evaporates at the end of allocation, they are likely to enroll in this course, despite the price being significantly higher than their marginal utility would indicate, as the alternative is to receive no course and no money.

Given price and marginal utility are no longer linked, it is much more difficult to determine the optimal tax rate, as it is unclear how much a given tax may impact aggregate demand. This creates headaches for administrators attempting to target specific equilibria.

One natural response to this dilemma would be to allow students to save unspent currency.

However, this causes the mechanism to no longer be strategyproof. Rather than listing their true preferences, students are incentivized to prefer less desirable courses in order to save credit for later semesters. While this could, in theory, be rectified by performing a single allocation process for all semesters, this would require students to express preferences over a much larger universe of possible schedules, and anticipate their preferences much further in advance. In addition, in a larger, all-semester allocation problem, there remains no outside good, returning us to the problem at hand.

When prices become severed from utility, it becomes difficult to formulate a solution in the price space (the tax rate) with a predictable relationship to utilities (the externalities). However, we can predict how changing the tax rates in different ways will impact allocations, as seen in section 2.6.

2.3 Strategyproofness

2.3.1 Preference Manipulation

Claim: Submitting true preferences is a weakly dominant strategy for each individual student.

When N is large, students act as price takers in a competitive equilibrium, as shown in Budish (2011). Therefore, we can take the price vector \vec{p} as exogenous.

Consider a student i . From our equilibrium definition, we know they have the match they most prefer from their submitted preferences such that $\vec{x}_i^T \cdot \text{diag}(\vec{k}_i) \cdot \vec{p} \leq b_i$.

Preference manipulations can fall into one of four categories:

1. Ranking \vec{x}_i below its true position
2. Ranking \vec{x}_i above its true position
3. Changing the ordering of preferences above \vec{x}_i
4. Changing the ordering of preferences below \vec{x}_i

I will now show none of these four manipulations are profitable.

1. Suppose i were to manipulate their preferences such that some other \vec{x}_i' were falsely ranked greater than \vec{x}_i . There are two possibilities:

- (a) \vec{x}'_i is not achievable—in this case, this allocation is infeasible for i , and so the allocation algorithm will attempt to assign them their next most preferred bundle, \vec{x}_i , yielding the same outcome.
- (b) \vec{x}'_i is achievable—here, the mechanism would see that \vec{x}'_i is seemingly preferred to x_i , and so assign it to student i . However, we know that i truly prefers \vec{x}_i to \vec{x}'_i ; therefore, i is strictly worse off.

Thus, truthtelling weakly dominates any manipulation where \vec{x}_i is artificially low

2. Suppose i were to manipulate their preferences such that \vec{x}_i was falsely ranked above another \vec{x}'_i . We already know from our equilibrium condition that \vec{x}'_i is unachievable given the budget constraint. Thus, i has not changed their outcome.
3. Suppose i were to manipulate their preferences above \vec{x}_i . We know none of these bundles are achievable; thus, their ordering is irrelevant, and i will still end up with bundle \vec{x}_i .
4. Suppose i were to manipulate their preferences below \vec{x}_i . As \vec{x}_i is achievable, i will never be assigned bundles of a lower rank. Thus, i has not changed their outcome.

As seen above, truthtelling weakly dominates all four types of possible manipulations. Therefore, truthtelling is a weakly dominant strategy.

2.3.2 Outside Strategy: Professorial Override

In the status quo, a significant portion of course allocation mechanisms allow professors to override the system and enroll students into their courses. This can create unravelling effects; if students are better off getting overrides from professors than enrolling formally, they are disincentivized from using the primary mechanism.

While one solution would be to simply prevent professorial overrides, this would likely be unpopular. My proposed alternative would be to give students exactly what they want and commit them to the class; however, they must still pay for it out of their endowment. Thus, a course override makes a student weakly worse off.

The proof of this claim is straightforward; it is possible they would have been allocated the course through standard operation of the mechanism, in which case the schedule they receive under this modification is identical to that which they would have received in the absence of an override. Alternatively, they receive their most preferred bundle of courses *containing* the mandated course. This bundle may be less preferred than some other bundle which would have been feasible; however, by forcing a system override, the student has made this alternative bundle infeasible. In this case, the student is worse off. Therefore, under this modification, students cannot improve their assignment through professorial overrides.

An extreme case of this failure would occur when students get their entire bundle of courses through professorial overrides. In this case, their bundle of courses may exceed their budget, making this bundle superior to that which they would have received from A-CEEI, while the market administrator has no retributive mechanism. In this scenario, it is likely better to address these issues on the professor side rather than the student side, perhaps by only allowing professors a highly limited number of course overrides. It is also unclear what incentive the professors would have to support this form of unravelling; if their preference is for a general class of students, they could simply indicate preference for those students in their tax vectors rather than dealing with individual appeals.

2.3.3 Outside Strategy: Appeals

Allocation mechanisms frequently require some sort of appeals process. It is often the case these appeals are justified; however, this also opens up the possibility for student abuse. If students find the appeals process too generous, it may incentivize them to do much of their course selection after the initial allocation. One example of this strategy would be a senior who put off mandatory classes until the last possible semester, and now finds that all bundles which allow them to complete their degree are too expensive. If the appeals court then provides them with a bundle sufficient for graduation, they have now exceeded their fair allocation.

The obvious solution to this issue is to make the appeals process draconian; in the above example, the senior would simply need to enroll for another semester. However, this solution could end up

throwing out valid appeals, and would also likely be unpopular.

I would argue that a proper implementation of my mechanism would have relatively few appeals of the above variety, as administrators would likely want to implement taxes such that seniors face lower prices for mandatory courses than their younger counterparts. In fact, this is one scenario where my modified mechanism surpasses the original, as this possibility does not exist without the addition of taxes.

In addition, if this problem becomes common, this may be indicative of a larger issue of undersupply of mandatory courses. By making a highly overbooked course mandatory, a university implicitly places a cap on the total number of majors in that department. This is fundamentally incompatible with allowing all who wish to major in a department to graduate. There is no allocation mechanism which can solve this problem; there simply aren't enough seats for the number of students.

In the general case, where a student appeals to receive a non-mandatory course, it would be reasonable to allow changes given extenuating circumstances. One solution would be to allow students to take "advances" on their endowments from later semesters. This would solve the problem in the short term, but may raise similar issues down the road if the student then doesn't have enough credit remaining in the later semester. From my perspective, moderate advances seem permissible at the jurisdiction of the administration conditional on demonstrating the lack of credit in later semesters will not be problematic.

Having said this, it is inadvisable to allow inter-semester transfers generally. Allowing borrowing over time compromises the strategyproofness of the mechanism; if in both freshman and sophomore year I can borrow against my endowment senior year, I may not wish to list my true first choice bundle freshman year, as I would then have less credit available to borrow the following year, as outlined in section 2.2.

Further, inter-temporal transfers undermine the administrator's power to impose time-variant taxes; if the tax scheme is generally favorable to seniors and unfavorable to freshmen, a utility-maximizing student would borrow freshman year against their senior endowment. While this is optimal from that student's perspective, this makes it more difficult to determine optimal taxation, since administrators must anticipate the next four years of tax rates rather than just the current

semester's.

2.4 Efficiency

I cannot guarantee Pareto efficiency across brackets; in some sense, this is by design. Under an optimal tax schema, a student's enrollment in a course is taxed because they impose some social cost; thus, even if their enrollment would be privately preferred, and thus a trade for this seat would make both them and their trading partner better off, the market administrator has decided the social cost of this trade outweighs the private benefit.

Among those who impose the same social costs by enrolling in a course, however, market administrators are indifferent about who enrolls, while the students may not be. Thus, an ideal mechanism will be Pareto efficient within a bracket. Budish (2011) shows in Proposition 2 that A-CEEI is Pareto efficient in the special case where all students are in a single tax bracket. Note that when examining any subset of players who share a tax bracket, my modified CEEI becomes equivalent to Budish's A-CEEI. Therefore, we can show that this mechanism must be Pareto efficient within a bracket.

Suppose there exists a trade between two members of the same tax bracket which would make both parties better off after completion of the modified CEEI allocation process. If this is the case, one of the traded courses must have a greater price, or the courses are equal in price. If this is so, the student who possesses the weakly more valuable course could have simply bought the course he desires at the market price instead of the course he possesses. Therefore, this could not have been a CEEI allocation to begin with. Given this contradiction, we know there cannot exist trades within a bracket which would improve both parties, so modified CEEI must be Pareto efficient within a bracket.

2.5 Fairness

2.5.1 Envy

One objective for a course allocation is that it minimizes “envy;” that is, it minimizes how many students would prefer another student’s schedule to their own, as well as how much they prefer the other student’s schedule. Basic A-CEEI is “envy-bounded by a single good,” meaning envy may exist, but by removing a single good from the envied bundle, the envier now prefers their own bundle. This imperfection comes from the budget dispersion necessary to ensure existence; if incomes were wholly equal and an equilibrium still existed, the allocation would be envy-free.

The proof of bounded envy is trivial; given incomes are (approximately) equal, a student who prefers another student’s bundle is free to have purchased the desired bundle instead provided their income could have covered the price of the courses. If the student cannot afford the bundle, we can remove the most expensive course from the bundle, which comprises at least $(\text{max bundle size})^{-1}$ of the total cost. As incomes are arbitrarily close, this meaningful reduction in the cost of the bundle has made the remaining courses affordable. Thus, equilibrium is envy bounded by a single good.

The modified A-CEEI developed in this paper clearly cannot be bounded in envy for all tax schema, as serial dictatorships are achievable using this mechanism and are unbounded in envy. However, we can make claims about envy within a bracket.

We know that the subset of students within a given tax schema experience the game as A-CEEI; for example, if all sophomore economics majors experience the same tax rates, they all face the same prices, and therefore have the same relationships to each other as in A-CEEI. Thus, while envy is unbounded between tax groups, we do know that allocations are envy-bounded by a single good within a given tax group.

2.5.2 Theoretical Bounds

Beyond considering intra-bracket fairness, administrators may place some value on inter-bracket fairness. An ideal mechanism would have students receive course bundles which are approximately equally desirable. Given this is a modified A-CEEI, students begin with endowments which are

dispersed uniformly on the interval $[1, 1 + \beta]$, where $0 < \beta < \min(N^{-1}, (\text{max bundle size})^{-1})$ is arbitrarily small. This means the ratio of the wealth held by the most wealthy person to wealth held by the least wealthy person is bounded by $1 + \beta$.

Once we introduce taxes, however, this inequality may grow. Let the minimum possible tax rate be 1. In the case where the person with the greatest base endowment only desires courses with the minimum tax rate, their functional endowment is still $1 + \beta$. However, the least fortunate possible person—who has the lowest possible raw endowment and desires only courses taxed at the maximum possible rate—functionally has an endowment of $\max(\vec{k})^{-1}$. Therefore, the ratio of the endowments of the most wealthy person to the least wealthy person would be $\max(\vec{k}) \cdot (1 + \beta)$.

2.5.3 Empirical Measures

Of course, theoretical bounds may be significantly worse than true performance. Administrators may wish to find equilibria that balance inequality with market correction, so it would be useful to have a measure of market inequality.

One benefit of using a competitive equilibrium mechanism is that prices can be used as a measure of a bundle's value. Given scarcer courses have higher prices, the sum of the prices of a course are a good indication of the wealth collected by a student. This allows us to construct a Gini index to examine the distribution of wealth over the student body, as suggested in Budish, Cachon, et al. (2016).

When constructing the Gini index, however, it is important to consider which prices to use for our calculations. Given that prices are determined endogenously with tax rates, it would be difficult to compare inequality across tax schemes if we used the prices determined by the modified A-CEEI. Instead, I would suggest using the untaxed prices (found using the Budish A-CEEI) as a baseline for measuring bundle desirability.

Suppose administrators wish to balance the efficiency gains of taxing externalities with potential equity losses from changing students' relative endowments. Comparing the Gini index of Budish's A-CEEI with the index of the modified A-CEEI could provide a metric of equality losses. If the distribution of courses ends up being more unfair than desired, administrators could then raise taxes

on the students with the most valuable bundles, functionally lowering their endowments, and then re-run the process.

Note the Gini index is not a perfect measure of equality in this context. It is possible for a student's first choice bundle to be entirely composed of zero-cost courses. In this case, they appear much worse off in the Gini index than they are in reality. However, there are no students who receive higher priced courses that would have preferred lower priced courses; thus, no one is worse off than they appear. Therefore, the Gini index can act as a worst-case upper bound on inequality within an observed system.

2.6 Choices of Tax Rates

Let each student i face some vector of tax rates \vec{k}_i . When attempting to enroll in course j , the course's price p_j is multiplied by k_{ij} before the student attempts to make their purchase. By setting k vectors intelligently, school administrators can exert significant control over the final allocation.

2.6.1 Relative vs. Absolute k

When determining allocations, the absolute magnitudes and directions of the tax vectors are functionally irrelevant. If a set of vectors K produces allocation X , the set of vectors cK , where each vector is scaled by the same constant, will produce the same allocation X . This is because prices are determined endogenously; if everyone shifted their accounting from dollars to cents, prices would change but allocation of goods would remain identical.

Beyond questions of magnitude, direction is also only meaningful when it provides one consumer with a comparative advantage. Compare the allocations derived when all consumers face the tax rates $\langle 1, 1, 1 \rangle$ vs. when all face tax rates $\langle 0.1, 1, 10 \rangle$. In this case, the prices will simply scale to counteract the tax rate, so each student will still spend the same fraction of their endowments on each course, yielding the same final allocation.

Provided with the above, administrators need not consider absolute magnitudes and directions of tax vectors, but only how the vectors provide certain students with advantages over others.

2.6.2 Vector Magnitude and the Second Welfare Theorem

As discussed in section 1.3, we know from Miralles and Pycia (2014) that conditional on the existence of a strategyproof, price-based allocation mechanism, it is possible to reach any Pareto-efficient allocation through variance of student endowments. Further, as shown in section 2.3.1, modified CEEI is both functionally strategyproof and price-based. Thus, we know that by transforming the initial allocations, the market administrator is able to achieve any allocation which is Pareto-optimal from the perspective of each student (that is, which ignores externalities s_{ij}).

Consider the set of \vec{k} vectors such that each student faces a constant multiplier for all courses, while different students still face different values. By varying the students' constants, a market administrator can generate any desired initial endowment, and can therefore reach any desired Pareto-efficient allocation.

2.6.3 Vector Direction and Externalities

Given it is possible to reach any Pareto-efficient allocation by simply adjusting magnitudes, why would an administrator modify the direction of students' \vec{k} vectors? Any set of vectors that is not wholly linearly dependant shifts the market away from Pareto efficiency. When we consider k values as taxes, this becomes intuitively clear; by imposing a tax on the market, the administrator has created deadweight loss. This is true so long as enrollments do not impose externalities on other students. However, in the presence of externalities as described in section 2.2, varying tax vector directions can help to correct market failures, leading to gains in overall welfare.

2.6.4 Notable values of \vec{k}

- *A-CEEI*—To generate Budish (2011)'s A-CEEI allocation, have all student-course rates be constant.
- *Serial Dictatorship*—To implement a serial dictatorship, we must implement large budget inequalities. More specifically, if vectors are constant for each student, and each student's constant is a unique element of the sequence $\left((1 + \text{largest permissible bundle size})^{-i}\right)_{i=1}^N$, then

students play the dictatorship game in order of rising tax rate.

- *Absolute Preferences*—Suppose a professor wishes to enroll all freshmen interested before admitting any underclassmen. It is straightforward to show this allocation constraint is not well-defined under this mechanism.

Suppose we had some nonzero tax rates k_f and k_u for freshmen and upperclassmen, respectively. For simplicity, let $k_f = 1$. Consider a freshman and upperclassman with nonzero willingness-to-pays of w_f and w_u . To satisfy the above constraint, we would require a k_u such that

$$\forall w_f, w_u \in (0, \max(\vec{b})] : \frac{w_u}{k_u} < w_f \iff \frac{w_u}{w_f} < k_u$$

Given $w_f \in (0, \max(\vec{b})]$, w_f can become arbitrarily close to zero, making $\frac{w_u}{w_f}$ an arbitrarily large lower bound on k_u , which is impossible. Therefore, it is impossible to implement absolute preferences in this system.

3 Market Clearing Error Given K

The Othman, Sandholm, and Budish (2010) price search algorithm can come close to perfect market-clearing demand; however, due to goods being non-divisible, it is impossible to guarantee perfect market clearing. In this section I expound on the proof provided in Budish (2011) to accommodate varying tax rates, and find that market clearing demand is bounded by $\frac{h\sqrt{\sigma M}}{2}$, where $\sigma = \min(M, 2 \cdot \text{max schedule size})$ and h is the maximum number of distinct tax rates faced by a single course. As shown below, σ is the maximum number of courses which can experience a demand discontinuity at a single price.

3.1 Tax Rates and Market Clearing Error

Consider all the students with the same tax rate k_{ic} when facing course c . We know from the dispersion of incomes at the beginning of the process that none of these students have equal incomes; however, these students all face the same price, as their tax rates are identical. Therefore, from

the basic A-CEEI proof, we know at most M agents have intersecting hyperplanes, meaning the maximum discontinuity in aggregate demand with respect to price is $M\sqrt{\sigma}$.

However, we know nothing about the budgets of students with tax rate $k_{jc} \neq k_{ic}$. By the same argument as the k_{ic} case, there may be as many as M agents' hyperplanes intersecting at a given point where those agents have tax rate k_{jc} . There is nothing preventing this intersection from occurring on top of the intersection in the k_{ic} case; thus, there is now a new upper bound of $2M$ intersecting hyperplanes, yielding a discontinuity of $2M\sqrt{\sigma}$. In the general case, this yields a pattern that given h distinct values of k , the maximum aggregate demand discontinuity will be $hM\sqrt{\sigma}$. Note, however, that there exist only N agents in the problem, meaning no more than N hyperplanes may intersect at any given point. Therefore the true maximum aggregate demand discontinuity is actually $\min(hM, N)\sqrt{\sigma}$.

Let us examine some special cases to verify this result. Budish's A-CEEI is the case where all students have the same k for all classes. In this case, $h = 1$, so the maximum demand discontinuity is $M\sqrt{\sigma}$, as found in his paper. In the most degenerate case, each agent's k_{ij} for each class is the exact reciprocal of their budget, perfectly offsetting the budget dispersion. Therefore, it is possible for all students to have the exact same preferences, yielding a maximal demand discontinuity of $N\sqrt{\sigma}$, as seen in standard CEEI.

3.2 Market clearing demand and the convex hull

We know from Budish's proof that some market clearing demand exists within the convex hull of the feasible demands. The maximum possible discontinuity in demand for a single course is $h\sqrt{\sigma}$. However, this only occurs when all h possible tax groups have identical budget hyperplanes. Thus, given the hM intersecting hyperplanes, there are only $\frac{hM}{h} = M$ feasible demands.

As demand is an M -dimensional space, we know then from Budish's proof that the worst-case distance from a feasible demand to the perfect market-clearing demand occurs when the feasible demands form a hypercube with side length $h\sqrt{\sigma}$ and the market-clearing demand is equidistant from all points. In this case, the distance from the perfect demand to all feasible demands is half the diagonal of the hypercube, which equals $\frac{h\sqrt{\sigma M}}{2}$. Note an important trait of this bound which

carries over from A-CEEI: the error bound does not grow with the number of students enrolled in the university. Therefore, as $N \rightarrow \infty$, the value of the error relative to the general population will decrease to zero.

4 Conclusion

I have shown that introducing a tax scheme into Budish’s Approximate Competitive Equilibrium from Equal Incomes mechanism provides market administrators with significant power to correct market failures while maintaining many of the desirable properties present within the original mechanism. These properties include incentive compatibility, error independent of the number of seats to allocate, and bounds on envy within groups of students. Further, this infrastructure is able to unify several existing mechanisms—CEEI, A-CEEI, and the serial dictatorship—under a single mechanism. While I cannot provide a concrete answer for optimal tax rates, I do establish several properties of desirable tax rates, and highlight the costs of modifying tax schema in various dimensions.

References

- Arrow, Kenneth J et al. (1951). “An extension of the basic theorems of classical welfare economics”. In: *Proceedings of the second Berkeley symposium on mathematical statistics and probability*. The Regents of the University of California.
- Budish, Eric (2011). “The combinatorial assignment problem: Approximate competitive equilibrium from equal incomes”. In: *Journal of Political Economy* 119.6, pp. 1061–1103.
- Budish, Eric, Gérard P Cachon, Judd B Kessler, and Abraham Othman (2016). “Course match: A large-scale implementation of approximate competitive equilibrium from equal incomes for combinatorial allocation”. In: *Operations Research* 65.2, pp. 314–336.
- Budish, Eric and Estelle Cantillon (2012). “The multi-unit assignment problem: Theory and evidence from course allocation at Harvard”. In: *American Economic Review* 102.5, pp. 2237–71.

Miralles, Antonio and Marek Pycia (2014). “Prices and efficient assignments without transfers”. In: *Available at SSRN*.

Othman, Abraham, Tuomas Sandholm, and Eric Budish (2010). “Finding approximate competitive equilibria: Efficient and fair course allocation”. In: *Proceedings of the 9th International Conference on Autonomous Agents and Multiagent Systems: volume 1-Volume 1*. International Foundation for Autonomous Agents and Multiagent Systems, pp. 873–880.

Sönmez, Tayfun and M Utku Ünver (2010). “Course bidding at business schools”. In: *International Economic Review* 51.1, pp. 99–123.