Costly Information Acquisition and Stable Matching Mechanisms

Nicole S. Immorlica[*], Jacob D. Leshno[†], Irene Y. Lo[‡], Brendan J. Lucier[§]

October 8, 2018

PRELIMINARY AND INCOMPLETE

Abstract

This paper studies how mechanisms affect information acquisition by agents in matching markets. We consider a tractable "Pandora's box" model where students hold a prior over their value for each school and can pay an inspection cost to learn their realized value. The model captures how students' decisions to acquire information depend on priors and market information, and can rationalize a student's choice to remain partially uninformed. In such a model students need market information in order to optimally acquire their personal preferences, and students benefit from being "last to market", i.e. waiting for the market to resolve before acquiring information. We extend the definition of stability to this partial information setting and define regret-free stable outcomes, where the matching is stable and each student has acquired information optimally, as if she were last to market.

We show that regret-free stable outcomes have a cutoff characterization, and the set of regret-free stable outcomes is a non-empty lattice. However, there is no mechanism that always produces a regret-free stable matching, as there can be information deadlocks where every student finds it suboptimal to be the first to acquire information. In settings with sufficient information about the distribution of preferences, we provide mechanisms that exploit the cutoff structure to break the deadlock and approximately implement a regret-free stable matching.

[*]Microsoft Research, New England
[†]University of Chicago Booth School of Business
[‡]Columbia University, Department of Industrial Engineering & Operations Research
[§]Microsoft Research, New England

1 Introduction

Matching markets have been the subject of much academic research as well as substantial interest from practitioners. In these markets agents have preferences over the individuals they are matched to, and the assignment is not determined simply by monetary transfers. Matching theory investigates the role of marketplace rules in determining the allocation, and elucidates how matching markets can and should compute the overall assignment from individual agent preferences. Such models allow us to better understand decentralized markets, such as college admission, or to facilitate a better design of centralized assignment mechanisms, such as the medical match and school choice (see, e.g. Roth and Sotomayor, 1992; Roth, 2015).

In this paper, we investigate how mechanisms affect how agents form their preferences. The prevalence of incomplete information is well-studied in the context of auction markets (see e.g. Eso and Szentes (2007); Milgrom and Weber (1982)), but is relatively unexplored in matching market settings. This is despite the fact that in matching settings such as medical residency matching and school choice, it is common not only for agents to have incomplete information about their preferences, but also for them to spend a significant amount of effort investigating potential placements before forming their final preferences. For example, in NYC public high school admissions students must submit their preferences over more than 700 programs at more than 400 high schools. Moreover, costly information acquisition is also an important equity problem in school choice, as students from underprivileged backgrounds are often inadequately informed about their options and must exert the most effort to determine their preferences (see, e.g. Hassidim et al., 2015; Kapor et al., 2016).

Thus motivated, we study the effects of market design on costly information acquisition in a many-to-one school choice market. In our model school priorities are common knowledge, and students can acquire costly information about their preferences over school. We model each agent's information acquisition problem using the "Pandora's box" framework of Weitzman (1979) in the tractable continuum matching market of Azevedo and Leshno (2016). Each student knows a prior distribution for each school's utility to them, and must pay a cost in order to learn actual utility realization. The utility realizations are independent for each student, and students individually decide on their information acquisition process. The student information acquisition problem admits an optimal solution via a simple index policy, and allows for students to only partially collect information.

We define *stability under incomplete information* for this setting: an outcome for the market is stable with respect to acquired information and information acquisition costs if there is no blocking pair, i.e. a (student, school) pair such that the student (i) has higher priority at a given school than another student assigned to that school or the school is undercapacitated, and the student either (ii) prefers the school to their assigned school, given their acquired information, or (ii') does not have enough information to make a decision and is willing to pay the cost to collect further information. This definition extends the standard definition of stability, and is equivalent to the standard definition when students do not incur information acquisition costs and collect all preference information. However, in the presence of information acquisition costs it is possible for different acquired information to lead to different stable outcomes. Hence the design of the market mechanism can induce beliefs that lead students to acquire information differently and implement different outcomes, even when there is a unique stable matching under full information.

In settings with costly information acquisition students need information about their possible matches in order to optimally acquire information, and students may benefit from waiting for the market to resolve before acquiring information. We refine the set of stable outcomes to the set of *regret-free stable outcomes*, under which the information acquired by each student is the same as if they performed their optimal information acquisition process knowing the preferences and information acquisition processes of all other students. In other words, each student acquires information as if she were the last the enter the market, and no student regrets not waiting for further information about other students' preferences before learning her own preferences. This means that regret-free stable matchings do not depend on student beliefs. We furnish the surprising result that *the set of regret-free stable matchings has a lattice* *structure*, which it inherits from of the set of stable matchings under complete information (attributed to John Conway in Knuth, 1976), and hence is non-empty and has an outcome that is unambiguously the best for all students.

We then turn to the problem of providing matching mechanisms for implementing regret-free stable outcomes. We show that as regret-free stable matchings are characterized by cutoffs, the student-optimal regret-free stable outcome can be implemented by learning and posting the appropriate admissions cutoffs. For example, given sufficient market structure, school-proposing Deferred Acceptance can be implemented in a sequential manner to learn the regret-free stable cutoffs with regret-free information acquisition. However, we also demonstrate that there exist economies where regret-free stable matchings cannot be computed without incurring additional costly information acquisition, and also where the student-optimality of a regret-free stable matching cannot be verified without incurring additional costly information acquisition. In general settings, standard mechanisms can result in *information deadlocks*, where no information is gathered because every student finds it strictly optimal to wait for others to acquire information first. Hence the presence of costly information acquisition does not affect the structure of the set of stable outcomes but rather the algorithmic questions of computing a regret-free stable outcome and verifying its optimality.

We show how to approximately compute the market-clearing admissions cutoffs when we have historical information about demand or can estimate it by subsampling, and in such settings provide mechanisms that implement outcomes that are student-optimal regret-free stable with respect to perturbed school capacities. Our results illustrate that, *given sufficient information about aggregate student demand for schools, it is possible to approximately implement a regret-free stable matching.*

1.1 Prior Work

This paper contributes to the literature of matching markets with incomplete information. The stream of work that is closest to ours is that of Aziz et al. (2016); Rastegari et al. (2013, 2014), which analyze a matching model where there is partial ordinal information on both sides of the market that can be refined through costly interviews. They ask computational questions regarding the minimal number of interviews required to find a stable matching, and find that under a tiered structure an iterative version of DA minimizes the number of required interviews. Our finding that a sequential version of DA implements a regret-free stable matchings when agents are willing to inspect all schools they can attend is a particular case of this result where the preferences of one side are known. Drummond and Boutilier (2013, 2014) consider more general algorithms that acquire information through both interviews and comparisons and provide algorithms that achieve approximately stable matchings with low information costs. Lee and Schwarz (2009); Kadam (2015) also study information sharing through interviews.

Several papers consider other aspects of imperfect information in matching markets, without allowing agents to search for information. Liu et al. (2014) suggest a notion of stability under asymmetric information between agents. Chakraborty et al. (2010) consider agents with incomplete information who update their preferences after seeing the matching. Ehlers and Massó (2015) demonstrate that there is a strong connection between ordinal Bayesian Nash equilibria of stable mechanisms under incomplete information and Nash equilibria of the mechanism under corresponding complete information settings. We similarly define a notion of stability under incomplete information and find a strong parallel with the structure of stable matchings under complete information.

Empirical work demonstrates that incomplete information is important in the school choice setting. Kapor et al. (2016) provides empirical evidence that many students participating in a school choice mechanism are not well informed, and make mistakes when reporting their preferences, and Dur et al. (2016) provides evidence that different parents exert different levels of efforts in learning about school choice.

There is also a growing literature about information acquisition in market design. In an auction setting, Kleinberg et al. (2016) shows that descending price auction create optimal incentives for value discovery. Chen and He (2015) study how the DA and Boston mechanisms give participating agents incentives to learn their preferences and preferences of others, but limit attention to the decision of whether to learn the full ordinal or cardinal valuation for all schools. Bade (2015); Harless and Manjunath (2015) consider information acquisition in assignment problems.

The rational inattention literature that stemmed from the macroeconomic literature also looks at information acquisition by agents. This literature uses a framework introduced by Sims (2003) where the costs of signals are given by information theoretic measures of the informativeness of the signals. Matějka and McKay (2015) shows that in that framework agent's choices can be formulated as a generalized multinomial logit, and Steiner et al. (2017) give a tractable formulation for the choices of agents with endogenous information acquisition in a dynamic setting. Our approach differs in that our model uses a different cost structure, and focuses on the interaction between information acquisition and market mechanisms.

A related question is the communication complexity of transmitting known preference to a mechanism. Gonczarowski et al. (2015) consider the communication complexity of finding a stable matching and show that it requires $\Omega\left(n^2\right)$ boolean queries. Ashlagi et al. (2018a) find that the communication complexity of finding a stable matching is low under assumptions on the structure of the economy and a Bayesian prior. Their Communication-Efficient Deferred Acceptance protocol utilizes messages about both acceptances and rejections. The analysis in both papers differs from ours in that they assume agent know their full preferences (for example, can report their first choice) and only consider the cost of communicating that information to the mechanism.

Finally, our work contributes to the growing number of papers exploring the use of sequential or multi-round school choice mechanisms. Bo and Hakimov (2017) and Ashlagi et al. (2018b) propose the Iterative Deferred Acceptance mechanism (IDAM) and Communication-Efficient Deferred Acceptance mechanism respectively, which allow for multiple rounds of message-passing where students can learn the set of schools with which they are likely to be matched. Such mechanisms are also currently used in practice; Dur et al. (2016) empirically study a public school system in Wake County that implements an iterative mechanism, Gong and Liang (2016) theoretically and empirically consider a college admissions system in Inner Mongolia that implements an iterative version of Deferred Acceptance, and Bo and Hakimov (2017) propose IDAM in response to a sequential mechanism previously used for college admissions in Brazil.

2 Model

We present a model where students learn their preferences through costly information acquisition. The set of schools is denoted by $\mathcal{C} = \{1, \ldots, n\}$, and each school $i \in \mathcal{C}$ has capacity to admit $q_i > 0$ students. A student is given by a quadruple $s = (F^s, c^s, r^s, v^s)$. School priorities are publicly known, and captured by the vector $r^s \in [0,1]^{\mathcal{C}}$. School i prefers student s over student s' if and only if $r_i^s > r_i^{s'}$. We say that r_i^s is the rank of student s at school i. Student s needs to perform costly information acquisition to learn her value for attending each school. Initially student s knows that the value for attending school i is distributed according to prior F_i^s, and may pay a inspection cost of $c_i^s > 0$ to learn the realized value v_i^s. Student s privately knows F^s, c^s (importantly, the designer does not know these parameters). Students must inspect a school in order to attend it. We assume that v_i^s is independently drawn across students and schools.[1]

With slight abuse of notation, we use a student type $\theta = \theta(s) = \left(F^\theta, c^\theta, r^\theta\right)$ to denote the initially known information of a student $s = \left(F^\theta, c^\theta, r^\theta, v^s\right)$. We refer to $\theta \in \Theta$ as a student type, and refer to $s = (\theta, v^s) \in \mathcal{S}$ as the student's realized preference. Formally, $\Theta = \mathcal{F}^{\mathcal{C}} \times \mathbb{R}^{\mathcal{C}} \times [0,1]^{\mathcal{C}}$, where $\mathcal{F}$ is the set of probability distribution functions, and $\mathcal{S} = \Theta \times \mathbb{R}^{\mathcal{C}}$. We will use s and θ interchangeably to index $F^\theta, c^\theta, r^\theta$. Given a type θ the realized values are randomly distributed $v^s \sim F^\theta$, and with slight abuse of notation we write $s \sim F^\theta$.

Definition 1. An *discrete economy* is given by $E = (\mathcal{C}, S, q)$, where $S = \{s_1, \ldots, s_N\}$ is the set of students and $q = \{q_i\}_{i \in \mathcal{C}}$ is the vector of quotas at each school.

[1]This implies that preferences of other students do not provide a student any information about v^s.

We make the following assumptions. First, all students and colleges are acceptable. Second, as r_i^s carries only ordinal information, it is normalized to be equal to the percentile rank of student s in college i's preferences, i.e. $r_c^s = |\{s' \mid s \succ_c s'\}| / |S|$. Third, school have strict priorities, i.e., $r_i^s \neq r_i^{s'}$ if $s \neq s'$. Fourth, the priors F^θ are such that students have strict preferences, i.e., $\mathbb{P}(v_i^s = v_{i'}^s) = 0$ for all s and $i \neq i'$. Last, we assume there is an excess of students, that is, $\sum_{i \in \mathcal{C}} q_i < |S|$.

It will useful to consider continuum economies where there is no aggregate uncertainty. The realized preferences v^s of a single student given his type $\theta(s)$ are random. In the continuum economy there is a continuous mass of students of any given type θ, and although the realized preferences of an individual student are random, the aggregate distribution over $s = (\theta, v)$ is known from the initial information F^θ. Formally, a continuum economy is described by a measure η over $\mathcal{S}$. We require that the measure η is consistent with initial information, that is, for any $A \subset \Theta$ and sets $V_i \subset \mathbb{R}^{\mathcal{C}}$ we have that

$$\eta(\{(\theta, v) \mid \theta \in A, v_i \in V_i\}) = \int_{\theta \in A} \int_{v \in V_1 \times \cdots \times V_n} dF^\theta(v)\, d\eta(\theta).$$

Definition 2. A *continuum economy* is given by $\mathcal{E} = (\mathcal{C}, \mathcal{S}, \eta, q)$, where $q = \{q_i\}_{i \in \mathcal{C}}$ is the vector of quotas at each school, and η is a probability measure over $\mathcal{S}$ that is consistent with initial information.

We make the same assumptions about continuum economies as for finite economies: namely that all students and colleges are acceptable; r_i^s is normalized so that for any $i \in \mathcal{C}$ and $x \in [0,1]$, we have that $\eta\left(\left\{(\theta, v) \in \mathcal{S} \middle| r_i^\theta \leq x\right\}\right) = x$; school priorities are strict, i.e. for any $x \in [0,1]$ we have $\eta\left(\left\{(\theta, v) \in \mathcal{S} \middle| r_i^\theta = x\right\}\right) = 0$; student preferences are strict, i.e. for any $x \in [0,1]$ we have $\eta(\{s = (\theta, v^s) \in \mathcal{S} | v_i^s = x\}) = 0$; and there is an excess of students, $\sum_{i \in \mathcal{C}} q_i < \eta(\mathcal{S}) = 1$. In what follows, we will define concepts for both the discrete and continuum economy, and let $\eta(\cdot)$ denote the cardinality of a set in the discrete economy, and the measure in the continuum economy.

As in the standard matching models, a *matching* is a mapping $\mu : \mathcal{S} \to \mathcal{C} \cup \{\emptyset\}$ specifying the assignment of each student. Overloading notation, for school $i \in \mathcal{C}$ let $\mu(i)$ denote the set $\mu^{-1}(i) \subseteq \mathcal{S}$ of students assigned to school i. For each student $s \in \mathcal{S}$ and school $i \in \mathcal{C}$, let the *inspection indicator* χ_i^s be an indicator function that is 1 if student s has inspected school i and 0 otherwise.[2] We denote the preference information revealed from inspections χ by $v|_\chi = \{v_i^s \mid \chi_i^s = 1\}$.

A matching μ is *feasible* with respect to inspections χ if for each school $i \in \mathcal{C}$ we have that $\mu(i)$ is η-measurable and $\eta(\mu(i)) \leq q_i$, and, for each student s, if $\mu(s) \neq \emptyset$ then then $\chi_{\mu(s)}^s = 1$. This last condition is tantamount to assuming that a student must inspect a school in order to attend it. A *feasible outcome* is a matching and inspection pair (μ, χ) such that μ is feasible with respect to χ. Given (μ, χ) the utility of student s is $u^s(\mu, \chi) = v_{\mu(s)}^s - \sum_{i \in \mathcal{C}} \chi_i^s c_i^s$.

2.1 Stability with Costly Information Acquisition

Consider a feasible outcome (μ, χ). As in the complete information settings, a student-school pair (s, i) forms a blocking pair if: (i) student s has higher priority than some student who is assigned to s or school s did not fill its capacity, namely $r_i^s > \inf\left\{r_i^{s'} \mid s' \in \mu(i)\right\}$ or $\eta(\mu(i)) < q_i$; and (ii) student s inspected school i and knows she prefers school i over her assigned school $\mu(s)$, namely $\chi_i^s = 1$ and $v_i^s > v_{\mu(s)}^s$.[3] When information acquisition is costly for students there may be a student-school pair (s, i) where (i) holds and student s did not inspect school i. We extend the standard definition and say that (s, i) forms a blocking pair if (i) holds; and (ii') s has not yet inspected school i and prefers to pay the inspection cost c_i^s and be assigned to the better school of i and $\mu(s)$, namely $\chi_i^s = 0$ and $\mathbb{E}_{\tilde{v}_i^s \sim F_i^s}\left[\max\left\{v_{\mu(s)}^s, \tilde{v}_i^s\right\} - c_i^s\right] \geq v_{\mu(s)}^s$. An outcome (μ, χ) is stable if there are no pairs (s, i) that block by satisfying either (i),(ii) (i.e. the classical stability condition) or (i),(ii').

We remark that stability of (μ, χ) depends only on student's initial information $\theta(s)$ and preferences revealed by inspections $v|_\chi$. Simple examples show that if $\chi \neq \chi'$ are different

[2] We are implicitly assuming that two students with the same type and values inspect the same schools.
[3] Recall that feasibility requires that $\chi_{\mu(s)}^s = 1$.

inspections and μ is a matching, it is possible that the outcome (μ, χ) is stable but the outcome (μ, χ') is not.[4]

Given a matching μ define the *budget set* of student s by

$$B^s(\mu) = \left\{ i \in \mathcal{C} \mid r_i^s \geq r_i^{s'} \text{ for some } s' \in \mu(i) \right\} \cup \{ i \in \mathcal{C} \mid \eta(\mu(i)) < q_i \}.$$

The budget set $B^s(\mu)$ is the set of schools such that (s, i) satisfy (i). A stable outcome (μ, χ) must assign student s to a school $i \in B^s(\mu)$ if s so desires, and the student s cannot be assigned to any school in the complement set $\mathcal{C} \setminus B^s(\mu)$. We say that a school i is *available to* s if $i \in B^s(\mu)$, otherwise school i is unavailable to s. The following immediate lemma characterizes stable outcomes in terms of budget sets.

Lemma 1. *A feasible outcome* (μ, χ) *is stable if and only if for every student we have that*

$$\mu(s) = \arg\max_{i \in \mathcal{C}} \{ v_i^s \mid i \in B^s(\mu), \, \chi_i^s = 1 \},$$

and for any $i \in B^s(\mu)$ *such that* $\chi_i^s = 0$ *we have that*

$$\mathbb{E}_{\tilde{v}_i^s \sim F_i^s} \left[\max \left\{ v_{\mu(s)}^s, \tilde{v}_i^s \right\} - c_i^s \right] \leq v_{\mu(s)}^s.$$

2.2 Regret-Free Stable Outcomes

To reach an outcome, students must perform inspections to acquire information about their values. These inspections might induce regret. Sometimes this regret is unavoidable: e.g., a student will regret having inspected a school with low value. Other times, regret is avoidable: i.e., a student should carefully select her inspections based on her available information. Below we characterize the information acquisition that maximizes the student's expected payoff

[4] For example, if there are only two schools, both of which are very costly to inspect compared to the possible values they may yield, then a student who has inspected and is matched to the first but has not inspected the second (χ) might not wish to pay the inspection cost for the second school, causing the current matching to be stable. However, if she had inspected the second schoool (χ'), she may realize a high value for it and thus form a blocking pair with it.

given all potential information, including her initial information and information that could be provided by the market. We determin how market information can affect the student's information acquisition decision. This allows us to define regret-free stable matchings where agents acquired information optimally.

Consider a student s who possesses initial information $\theta(s) = (F^s, c^s, r^s)$ and needs to select which schools to inspect. Since inspections are costly, student s will want to inspect a school only if inspecting the school can lead to being assigned to that school and receiving a higher value. In particular, student s will not want to inspect a school i if she knows that school i filled its capacity with higher priority students, and therefore she will not be assigned to the school i regardless of her value v_i^s. Thus, the set of schools that student s would like to inspect depends on her potential matches and the preferences of other students.

To fix ideas, first consider the isolated information acquisition problem where student $s = (\theta, v^s)$ is given a subset of schools $C \subseteq \mathcal{C}$ to choose from, each of which guarantees her admission. Student s needs to acquire information to form her preferences and then select her assigned school from C. If χ^s is s's inspection indicator and $i^* \in C$ is her selected school her utility is $v_{i^*}^s - \sum_{i \in \mathcal{C}} \chi_i^s c_i^s$. The adaptive inspection strategy that maximizes the student's expected utility given the initial information F^θ is derived by Weitzman (1979) and is stated in the following lemma.

Lemma 2. *(Weitzman 1979) Consider a student $s = (\theta, v^s)$ with initial information F^θ and inspection costs c^θ that can adaptdively inspect schools and choose a school from $C \subset \mathcal{C}$. For each school i, define a index $\underline{v}_i^\theta$ to be the unique solution to the equation $\mathbb{E}_{\tilde{v}_i \sim F_i^\theta}\left[\max\{0, (\tilde{v}_i - \underline{v}_i^\theta)\}\right] = c_i^\theta$. Sequentially inspect schools one by one in decreasing order of their index $\underline{v}_i^\theta$.[5] Continue inspecting the following school until the score of the next school to be inspected is below the maximal realized value among inspected schools.[6]*

We denote the inspections resulting from this optimal strategy by $\chi^{opt}\left(F^\theta, c^\theta, v^s; C\right)$.

[5]In case of multiple schools with equal index $\underline{v}_i^\theta = \underline{v}_{i'}^\theta$, break the tie by first inspecting the school $\min\{i, i'\}$.

[6]That is, if the set of inspected schools is $I = \{i \mid \chi_i^s = 1\}$ then inspect $j^* = argmax_{j \in C \setminus I}\left\{\underline{v}_j^\theta\right\}$ if $\underline{v}_{j^*}^\theta > \max_{i \in I} v_i^s$ and stop otherwise.

The optimal inspection policy is an index policy, where students use the prior information F^{θ} to compute indices $\underline{v}_i^{\theta}$ for each school i and inspect schools in decreasing order of their index.[7] The set of inspected schools depends on the set of available schools C, the indices $\left\{\underline{v}_i^{\theta}\right\}_{i\in C}$, and the realized values $\{v_i^s\}_{i\in C}$. The following example illustrates this.

Example 1. Suppose that $C = \mathcal{C} = \{1, 2\}$. Let $[x; p]$ denote the probability distribution which assigns probability p to the value x and $1-p$ to 0. Consider a student with $v_1 \sim F_1 = [10; 1/2]$, and $v_2 \sim F_2 = [4; 3/4]$ and let the inspection costs be $c_1 = 3, c_2 = 1$. Then the optimal inspection strategy is to first inspect school 1, and continue to inspect school 2 only if $v_1 = 0$. If instead $C = \{2\}$ the optimal inspection strategy is to only inspect school 2.

Knowing the set of available schools C helps the student in Example 1 to conduct the adaptive information acquisition that maximizes her expected utility. If the student does not know C she her inspection strategy may be sub-optimal in two ways. First, the student may inspect school 1 when it is not available, wasting the cost c_1. Second, the student should not inspect school 2 before she inspects school 1 or learns that school 1 is not available, because it is likely that she will not choose to inspect school 2 after inspecting school 1.

When student s is part of a matching market, the set of schools that are available to her depends on the resulting matching outcome, and therefore on the preferences of other students. Suppose that student s were to delay her information acquisition until the rest of the market resolved and the matching μ is realized. Arriving last to the market, student s can learn the set of schools available to her, which is $B^s(\mu)$ by Lemma 1. The student can optimize her information acquisition by using her initial information F^{θ}, c^{θ} as well as the market information $B^s(\mu)$, and applying Lemma 2. We say that the outcome (μ, χ) is regret-free stable if every students follows the optimal inspection policy informed by all available market information, that is, every student inspected schools as if she was the last to the market.

Definition 3. An outcome (μ, χ) is *regret-free stable* if (μ, χ) is stable and every student s

[7] Such a policy can also be constructed by mapping the problem to a multi-armed bandit (MAB) problem; see e.g Olszewski and Weber (2015) for details.

inspected the optimal set of schools given her available set of schools $B^s(\mu)$, that is $\chi^s = \chi^{opt}(F^s, c^s, v^s; B^s(\mu))$ for all $s \in S$. We let $M^{RF}(E)$ denote the set of regret-free stable outcomes for the economy E.

When an outcome is not regret-free stable some students can benefit from delaying their information acquisition until the remainder of the market resolved. Note that while the definition of regret-free stability is ex post in flavor, as it is stated in terms of each student's realized preferences v^s, it only imposes the restriction that the student follows the ex ante optimal inspection strategy given θ and $B^s(\mu) = B^\theta(\mu)$ (before observing v^s). A regret-free stable outcome could be ex post suboptimal, e.g. a student s may inspect a school $i \neq \mu(s)$ with low realized value v_i^s and ex post observe that the inspection cost c_i was wasted, but student s could do no better given all available information from $\theta(s)$ and the market information.

Remark. To verify whether (μ, χ) is regret-free stable it is sufficient to know $v^s|_\chi, \chi^s$ and F^s, c^s for each s, and does not require knowledge the students' values for uninspected schools.

3 The Structure of Regret-Free Stable Outcomes

In this section we provide several results about the structure of regret-free stable outcomes. We show that the set of regret-free stable outcomes is a non-empty lattice and give a concise characterization of regret-free stable outcomes in terms of cutoffs.

We begin by exploring how the demand of student s depends on the set of available schools. Consider a student s with available schools $C \subset \mathcal{C}$. If s optimally acquires information, she inspects $\chi^s = \chi^{opt}\left(F^\theta, c^\theta, v^s; C\right)$. Denote the resulting demand of s by

$$D^s(C) = \arg\max\left\{v_i \mid i \in C,\ \chi^{opt}\left(F^\theta, c^\theta, v^s; C\right) = 1\right\} \in C,$$

which is the most preferred inspected school. Note that $D^s(C)$ depends only on information that is revealed to s. The following lemma shows that $D^s(\cdot)$ satisfies WARP, and we can construct a full preference ordering $\succeq^{\Psi(s)}$ that yields the same demand.

Proposition 1 (Reduction to demand from complete information)**.** *Let $s = (F^s, c^s, r^s, v^s)$ be a realized student. There exist an ordering $\succeq^{\Psi(s)}$ such that for all $C \subset \mathcal{C}$ we have that*

$$D^s(C) = \max_{\succeq^{\Psi(s)}} (C) .$$

Proof. Using the indices from Lemma 2, define $i \succeq^{\Psi(s)} j$ if and only if $\min \left\{ \underline{v}_i^\theta, v_i^s \right\} \geq \min \left\{ \underline{v}_j^\theta, v_j^s \right\}$. It is straightforward to verify that $D^s(C) = \max_{\succeq^{\Psi(s)}} (C)$. □

That is, if we only observe the eventual selection from a set of available schools C, the student s is indistinguishable from a student with complete preference information and preferences $\succeq^{\Psi(s)}$. Given only initial information θ, the demand of θ from a set C is uncertain, as the realized values v^s are unknown. An immediate corollary is the distribution of demand of θ from a set C is identical to the distribution of $\operatorname{argmax}_{\succeq^{\Psi(s)}} (C)$, where $\succeq^{\Psi(s)}$ is the random preference ordering induced by drawing a random student $s = \left(F^\theta, c^\theta, r^\theta, v^s \right)$ from the distribution $v^s \sim F^\theta$.

Proposition 1 also allows a characterization of regret-free stable outcomes in terms of cutoffs, as in the complete information model of (Azevedo and Leshno, 2016). Cutoffs $P = \{P_i\}_{i \in \mathcal{C}} \in \mathbb{R}^{\mathcal{C}}$ are admission thresholds for each school. Cutoffs $\boldsymbol{P}$ determine the budget set of a student s to be

$$B^s(\boldsymbol{P}) = \{ i \in \mathcal{C} \mid r_i^s \geq P_i \} ,$$

which is the set of schools where s has better rank than the cutoff at that school. Note that $B^s(\boldsymbol{P})$ depends only on r^s and can be calculated from $\boldsymbol{P}$ and the initial information $\theta(s)$.

The demand of student s given cutoffs $\boldsymbol{P}$ is defined to be equal to $D^s(C)$ for a set of available schools equal to his budget set $C = B^s(\boldsymbol{P})$; for succintness we will write this as $D^s(\boldsymbol{P})$. Note that within the definition of $D^s(\boldsymbol{P})$ we require that student acquire information optimally. Aggregate demand for school i given cutoffs $\boldsymbol{P}$ is defined to be the mass of students that demand school i,

$$D_i(\boldsymbol{P}) = D_i(\boldsymbol{P}|\eta) = \eta(\{s \in \mathcal{S} \mid D^s(\boldsymbol{P}) = i\}).$$

We define market-clearing cutoffs as in Azevedo and Leshno (2016) and show there is a one-to-one correspondence between market-clearing cutoffs and regret-free stable outcomes. Note that the effect of information acquisition is captured within the definition of $D_i(\cdot)$.

Definition 4. A vector of cutoffs $\boldsymbol{P}$ is *market-clearing* if it matches supply and demand for all schools with non-zero cutoffs:

$$D_i(\boldsymbol{P}) \leq q_i \text{ for all } i \text{ and } D_i(\boldsymbol{P}) = q_i \text{ if } P_i > 0.$$

We can now state our characterization of regret-free stable outcomes.

Theorem 1. *An outcome (μ, χ) is regret-free stable if and only if there exist market-clearing cutoffs P such that for all s*

$$\mu(s) = D^s(\boldsymbol{P})$$

and

$$\chi^s = \chi^{opt}\left(F^\theta, c^\theta, v^s; B^s(\boldsymbol{P})\right)$$

Theorem 1 shows an equivalence between market clearing cutoffs and regret-free stable outcomes. Because demand $D(\cdot)$ provides us with sufficient information to determine whether $\boldsymbol{P}$ are market clearing cutoffs, demand $D(\cdot)$ is also sufficient to determine whether a matching μ yields a regret-free stable outcome with some χ. Using Proposition 1, for any market with information acquisition $\mathcal{E} = (\mathcal{C}, \mathcal{S}, \eta, q)$ we can construct a full information economy E that has the same demand, and therefore the economy $\mathcal{E}$ has the same market clearing cutoffs as E.

Theorem 2. *For every continuum economy $\mathcal{E}$ there exists a regret-free stable outcome. More-*

over, the set of regret-free stable outcomes (μ, χ) *is a non empty lattice under the order* $\succeq$ *defined by* $(\mu, \chi) \succeq (\mu', \chi')$ *iff* $v^s_{\mu(s)}(\omega) \geq v^s_{\mu'(s)}(\omega)\ \forall s \in S$.

Proof. The theorem follows from the reduction shown in Proposition 1 to the complete information setting, and analogous results by Blair (1988) on the lattice structure of many-to-one stable matchings in the complete information setting. □

Uniqueness of the regret-free stable outcome will require that the distribution of student types is regular. As student types have probabilistic demand, we will need to expand the definition of regularity beyond that found in Azevedo and Leshno (2016).

Definition 5. We say that $\theta = \left(F^\theta, c^\theta, r^\theta\right)$ is regular if for all $i \neq j$ we have that $\underline{v}^\theta_i \neq \underline{v}^\theta_j$ and $\mathbb{P}_{\tilde{v}^s_i \sim F^\theta_i}\left(\tilde{v}^s_i = \underline{v}^\theta_j\right) = 0$.

An measure η is *regular* if $\eta\left(\{s \mid \theta(s) \text{ is not regular}\}\right) = 0$ and the image under $D(\cdot \mid \eta)$ of the closure of the set $\left\{P \in (0,1)^{\mathcal{C}} \mid D(\cdot|\eta) \text{ is not continuously differentiable at } P\right\}$ has Lebesgue measure 0.

Intuitively, a type θ is regular if there are no ties, and so there is always a unique decision for whether to continue to inspect, and if so which school to inspect. A measure η is regular if there is no positive measure of irregular students and the implied demand is sufficiently smooth.

Theorem 3. *Suppose* η *is a regular measure. Then for almost every* q *with* $\sum_i q_i < 1$ *the economy* $\mathcal{E} = (\mathcal{C}, \mathcal{S}, \eta, q)$ *has a unique regret-free stable outcome.*

Proof. If η satisfies $\eta\left(\{s \mid \theta(s) \text{ is not regular}\}\right) = 0$ then for every cutoff P demand $D(P|\eta)$ is uniquely specified, and so there is a unique reduction to the complete information setting. The theorem follows from the reduction shown in Proposition 1 to the complete information setting, and analogous results by Azevedo and Leshno (2016) in this setting. □

4 Mechanisms

To this point we have discussed properties of regret-free stable outcomes. We now turn to the process by which a market-maker might implement such outcomes. In general, the market arrives at an outcome (μ, χ) following a sequential process in which students provide information to the market, the market provides information to students, students inspect schools to obtain more information, and the process repeats. We can describe any such market procedure as a dynamic mechanism.

The mechanism relies on the information it receives from students. We will be interested in two kinds of mechanisms. First, we consider *direct mechanisms* in which students report all of their private information and thereby delegate all decision-making. Second, we consider *choice mechanisms,* which are restricted in the nature of information that can be passed between the mechanism and the students. Choice mechanisms can only inform students about the availability of schools, and can only collect ordinal preference information from students.

4.1 Direct Mechanisms

In a direct mechanism the students fully delegate their decisions to the mechanism. We can think of a direct mechanism as the following iterative process. At any given state of the mechanism, we can write $\hat{\chi}$ to denote the indicator for the set of inspections the mechanism has conducted so far. Then, for each student s, the mechanism knows $v^s|_{\hat{\chi}}$ and knows F^s, c^s, r^s by assumption. Based on this information the mechanism can either decide to stop acquiring information and output the outcome $(\mu, \hat{\chi})$ for some matching μ, or to decide on the behalf of some students to inspect additional schools. We denote the information available in economy $\mathcal{E} = (\mathcal{C}, \mathcal{S}, \eta, q)$ after inspections χ by $I_{direct}(\mathcal{E}, \chi) = (\nu, q, v|_{\chi}, \chi)$, where ν is defined by $\nu(A) = \eta(\{s \mid \theta(s) \in A\})$ for $A \subset \Theta$. We denote that set of all possible inspection indicators by $\mathcal{X}$ and use $0 \in \mathcal{X}$ to denote the initial state where no student has inspected any school. Let $\mathcal{I}_{direct}$ denote the collection of all possible information sets $I_{direct}(\mathcal{E}, \chi)$.

Definition 6. A *direct mechanism* M is a mapping

$$M : \mathcal{I}_{direct} \to \left(2^{\mathcal{S}} \times \mathcal{C}\right) \cup \left(\mathcal{S}^{\mathcal{C}}, \chi\right)$$

that takes as input all the information available to the mechanism given previous inspections and returns either a next step of the inspection process, as described by tuple (S, i) of a set $S \subseteq \mathcal{S}$ of students to inspect the school $i \in \mathcal{C}$, or a final outcome (μ, χ) where χ is the current inspection indicator. To ensure termination of the mechanism, we require that iterated applications of the mechanism starting with $I_{direct}(\mathcal{E}, 0)$ will ultimately produce an outcome (μ, χ), which is the outcome of the mechanism.[8]

Imposing that the mechanism produces a regret-free stable matching ensures that the mechanism makes inspection decisions that are aligned with the optimal solution to each student's single-agent inspection problem.

Definition 7. A mechanism M is *(student-optimal) regret-free stable* if for any economy $\mathcal{E} = (\mathcal{C}, \mathcal{S}, \eta, q)$ the mechanism outputs a regret-free stable outcome (μ, χ).

4.2 Choice-Based Mechanisms

Direct mechanisms require that students directly report their initial information and all inspected values. However, students may not to communicate detailed cardinal information about their priors and costs. This may preclude the use of direct mechanisms in practice. We therefore consider mechanisms with lower communication requirements, where students provide only information about their preferred choice(s) from given sets of schools.

A choice-based mechanism is an iterative process where the mechanism provides information to students, students choose which schools to inspect, provide information back to the

[8] More formally, the mapping M induces a mapping $M' : \mathcal{I}_{direct} \to \mathcal{I}_{direct}$ defined by $M'(I_{direct}(\mathcal{E}, \chi)) = I_{direct}(\mathcal{E}, \chi')$, where: if $M(I_{direct}(\mathcal{E}, \chi)) = (S, i)$ then we let χ' be the inspections after the students in S have inspected school i, i.e. $(\chi')^s_j = 1 \Leftrightarrow \left(\chi^s_j = 1 \text{ or } s \in S, j = i\right)$; and if $M(I_{direct}(\mathcal{E}, \chi)) = (\mu, \chi)$ then we let $\chi' = \chi$. It is sufficient to require that if $I = (\nu, q, v|_\chi, \chi)$ is a fixed point of the mapping M' then $M(I) = (\mu, \chi)$ for some matching μ and the same inspections χ.

mechanism, and so on. Choice-based mechanisms do not have access to students' private information, and therefore cannot directly inform students which schools they should inspect. The mechanism can only provide information to students about which schools are available to them. Because we are interested in producing regret-free stable outcomes, which do not depend on students' beliefs about other students' preferences, we restrict our attention to mechanisms that only inform a student whether (a) it is certain that school i is available to her, (b) it is certain that school i is unavailable to her, or (c) it is uncertain whether school i is available or not. We use Accept (A), Wait-list (W), and Reject (R) to denote these three possible messages.

Students receiving an AWR message can choose which schools to inspect, and inform the mechanism of their choices. To simplify notation, we write the response of the student as a refinement of a preference ordering. Given s, χ^s define $\succcurlyeq^{s|\chi^s}$ by $i \succ^{s|\chi^s} i'$ if $\chi^s_i = \chi^s_{i'} = 1$ and $v^s_i > v^s_{i'}$, and $i \sim^{s|\chi^s} \phi$ if $\chi^s_i = 0$. Here we are using symbol ϕ to denote non-inspected schools. Let $\mathcal{L}(\mathcal{C} \cup \{\phi\})$ denote all transitive relations over $\mathcal{C}$ and the non-inspection symbol ϕ. For an economy $\mathcal{E} = (\mathcal{C}, \mathcal{S}, \eta, q)$ and χ an inspection indicator, the information available to a choice-based mechanism is $I_{choice}(\mathcal{E}, \chi) = \left(\left\{\succcurlyeq^{s|\chi^s}, r^s\right\}_{s \in S}, q, \chi\right)$. Let $\mathcal{I}_{choice}$ denote the collection of all possible information sets $\mathcal{I}_{choice}(\mathcal{E}, \chi)$.

Definition 8. An *Accept-Waitlist-Reject (AWR) mechanism M* is is defined via a mapping

$$M : \mathcal{I}_{choice} \to \left(\{A, W, R\}^{\mathcal{S} \times \mathcal{C}}\right) \cup \left(\mathcal{S}^{\mathcal{C}}, \chi\right)$$

that takes all the information available to the mechanism given previous inspections and returns either a AWR message for each student about each school, or a outcome (μ, χ) where χ is the current inspection indicator. We require that iterated applications of the mechanism starting with $I_{choice}(\mathcal{E}, 0)$ ultimately produces an outcome (μ, χ), which is the outcome of the mechanism.

We formally define general mechanisms, choice-based mechanisms and AWR mechanisms as dynamic games of incomplete information in the appendix.

5 Implementing Regret-Free Stable Matchings

We have shown that regret-free stable matchings have beautiful structural properties. They inherit the lattice structure of of stable matchings in the complete information setting, and can also be characterized using market-clearing cutoffs. In this section, we explore the mechanism design problem of implementing regret-free stable outcomes. We first show that regret-free stable matchings can be implemented by posting market-clearing cutoffs, and that information about these cutoffs is both necessary and sufficient for regret-free information acquisition.

We then show that in the incomplete information setting, the difficulties lie not in the existence of regret-free stable matchings, but in computating and verifying the stability and optimality of these matchings in a regret-free manner. While standard mechanisms popularized in the complete information setting can discover the market-clearing cutoffs, in many markets they will necessarily incur regret. This is because such mechanisms rely on students gathering and reporting information about their preferences and can result in *information deadlocks*, where no information is gathered because every student waits for others to acquire and report information first. Moreover, even when these mechanisms discover a regret-free stable matching, they will not be able to verify whether the matching is student-optimal without incurring regret.

Our conclusion is that information acquisition problems can be mitigated by posting market-clearing cutoffs. Cutoffs provide each agent with sufficient information to perform their regret-free stable inspections. The natural question, then, is how the market designer should determine which cutoffs to post. Market-clearing cutoffs can be learned and posted by the market designer without incurring regret if there is sufficient information about aggregate demand, either from historical data or from structured demand. We also show that even if market-clearing cutoffs can only be approximated, this is sufficient to implement a matching that is regret-free stable with respect to capacities that are close to the true capacities. Hence learning and posting cutoffs allows us to break the information deadlock and reach a regret-free stable outcome.

5.1 All You Need are Cutoffs

Recall from Theorem 1 that an outcome (μ, χ) is regret-free stable if and only if there exist market-clearing cutoffs P such that (a) each student follows the Weitzman optimal inspection strategy over her budget set as described by P, and (b) each student is matched to the school in her budget set that is most preferred, given the information revealed by the aforementioned optimal inspection strategy. Note, then, that if a student knows her budget set in advance, then she can optimally solve her information acquisition problem by proceeding as in a single-agent Pandora's Box problem to resolve her own incomplete information. An implication is that any matching mechanism that proceeds by committing to a collection of market-clearing acceptance cutoffs for the schools, then allowing each student to unilaterally optimize her inspection strategy and select her most-demanded school, will necessarily result in a regret-free stable match.

Mechanism 1 Acceptance with Market-Clearing Cutoffs (AwMC)

```
procedure AwMC(𝒞, 𝒮, q, P)
  Message Passing from Platform to Students
  for s ∈ 𝒮 do
      for i ∈ 𝒞 do
          if r_i^s ≥ P_i then
              Send message 'i accepts' to s
  Message Passing from Students to Platform
  for Student s in 𝒮 do
      Student s reports top choice school i^s that accepted them
      μ(s) ← i^s
  return μ
```

Theorem 4. *Let $\mathcal{E} = (\mathcal{C}, \mathcal{S}, \eta, q)$ be a continuum economy, and let $\boldsymbol{P}$ be the student-optimal market-clearing cutoffs in $\mathcal{E}$. Then Mechanism 1 is regret-free stable.*

Proof. We show that Mechanism 1 produces the student-optimal regret-free stable matching when all students report truthfully, and hence truthful reporting is a Nash equilibrium that produces the student-optimal regret-free stable matching. Indeed, the mechanism presents each student s with their budget set $B^{\mu}(s) = \{i \in \mathcal{C} \,|\, r_i^s \geq P_i(\eta)\}$, and student s is guaranteed to be matched to their reported favorite school $i^s \in B^{\mu}(s)$. Thus each student is presented precisly the single agent problem on $B^{\mu}(s)$. Solving this problem yields inspection strategy $\chi^s = \chi^{OPT}$ (by definition of χ^{OPT}), followed by truthfully reporting the students true favorite school: $i^s = D^s(P^*(\eta))$. By construction, demand exactly matches supply under this truthful reporting, so the output μ is the student-optimal regret free stable matching for $\mathcal{E} = (\mathcal{C}, \mathcal{S}, \eta, q)$. □

This result states that advance knowledge of market-clearing cutoffs are sufficient for regret-free stability. Indeed, posting cutoffs in advance of any information acquisition removes all uncertainty on the part of the agents about which schools they could match with. This, in turn, removes the possibility of regretting one's choice to explore the value of a match on the grounds that this school was ultimately unattainable. We note that this lack of regret does not depend on the posted cutoffs being market-clearing, but only that the mechanism commits to honoring the implied budget set for each student. Thus, for any economy $\mathcal{E} = (\mathcal{C}, \mathcal{S}, \eta, q)$ and any (not necessarily market-clearing) cutoffs $\boldsymbol{P}$, there exists a choice of capacities q' such that $\boldsymbol{P}$ are the student-optimal market-clearing cutoffs in $\mathcal{E}' = (\mathcal{C}, \mathcal{S}, \eta, q')$, and hence Mechanism 1 is regret-free stable with respect to $\mathcal{E}'$. We will make use of this fact when discussing notions of approximation in Section 5.3.

5.2 Regret-Free Choice-Based Mechanisms and Information Deadlock

Theorem 4 shows that knowing market-clearing cutoffs in advance of the market mechanism is sufficient for implementing a regret-free stable matching. Knowing market-clearing cutoffs in advance can be strong requirement. Indeed, aggregate uncertainty about agents' demands might make it difficult for the mechanism designer to know this information before interacting with the students. One might hope to avoid this impasse by way of a mechanism that reaches a stable matching without necessarily determining each student's full budget set. After all, each student's demand is ultimately described by a single ordering over all schools that is consistent across budget sets, which seems to suggest that it might not be necessary to fully learn every student's budget set in order to find a stable match. However, even though the realized demand is described by a single consistent ordering, the student cannot know this ordering a priori precisely because it depends on the values, which are only revealed after costly exploration. In other words, while the realized demand is ordered consistently, the order in which a student would wish to explore depends crucially on her budget set, and hence revealing the budget set can be crucial for avoiding wasteful exploration and regret.

It is perhaps useful to once again consider the school-proposing DA mechanism, which we recall can be interpreted as discovering the market-clearing cutoffs over time. Initially only the highest-ranked students are admitted to the schools of their choice, consistent with implicit cutoffs that are initially high, and these cutoffs then decrease (i.e., lower-ranked students are accepted) until the market clears. This choice-based approach does not post cutoffs in advance, but rather discovers them through repeated interaction with students. This provides hope that a mechanism that proceeds in multiple rounds can elicit enough information to find appropriate market-clearing cutoffs in a regret-free manner.

Indeed, we will show that under certain sufficient conditions on student preferences, the following iterative implementation of the school-proposing Deferred Acceptance is regret-free stable. The key idea is that while students' information acquisition problems are interconnected and can create information deadlocks, school priorities also provide students with partial information, and this may be sufficient to both start and finish the information acquisition process.

Algorithm 1 (Iterative Deferred Acceptance.)**.** *At each step:*

- *Each school i proposes to the top q_i students who have not yet rejected them.*
- *Each student (irrevocably) rejects some of the schools that have proposed to them.*

The algorithm terminates when no new proposals are performed, at which point all students are asked for their top choice school among all those that have proposed to them and which they have not rejected, and are assigned to school.

Theorem 5. *Suppose that at least one of the following conditions hold:*

1. *All students inspect all schools in their budget set and do not have indifferences, i.e.* $\mathbb{E}[v_i^s] = \infty$ *for all* $s \in S$ *and* $i \in C$*, and* $\mathbb{P}\left(v_i^s = v_j^s\right) = 0$ *for all* $s \in S$ *and* $i \neq j$*.*

2. *For all* $m \geq \min_i q_i$ *the same students occupy the top* m *ranks at all schools, i.e. for all* $i, j \in C$*,* $\left\{s \,|\, r_i^s \geq 1 - \frac{m}{|S|}\right\} = \left\{s \,|\, r_j^s \geq 1 - \frac{m}{|S|}\right\}$*. (Recall that* r_i^s *is normalized to be the percentile rank of student* s *in school* i*'s priorities.)*

Suppose all agents only perform regret-free inspections and report truthfully. Then Mechanism 1 almost surely implements the school-optimal regret-free stable matching.

Proof. The intuition is that our conditions guarantee that at all stages of proposal, there are students who have sufficient information about their budget set to both inspect some schools and reject some proposals. (This requires that students do not have indifferences in their preferences.) We show this formally for both cases.

Case (1). Suppose that all students inspect all schools in their budget set, i.e. condition (1) holds. First, assuming truthful reporting, it is regret-free for every student to inspect all schools that proposed to them. This is because if school i has proposed to student s, it is not full of students it prefers to s, so for all realized preferences v and for all outcomes $(\mu, \chi) \in M^{RF}(E)$ it holds that $i \in B^s(\mu)$. Since s is willing to inspect any school in her budget set it follows that it is optimal for student s to inspect school i.

Next, suppose students do not have indifferences in their preferences, i.e. $v_i^s \neq v_j^s$ for all s and $i \neq j$. Then it is regret-free for each student to reject all the schools that proposed to them except the one with the highest observed value. This is because if student s has inspected both i and j and $v_i^s > v_j^s$ then $\mu(s) \neq j$ for all $\mu \in M^{RF}(v)$ and it is optimal for s to reject the school with lower observed value.

When the algorithm terminates either all schools are at capacity or all students are assigned, since during each step all students reject all the schools that proposed to them except one, and so the algorithm terminates with a regret-free stable outcome.

Case (2). Suppose that for all $m \geq \min_i q_i$ the same students occupy the top m ranks at all schools, i.e. condition (2) holds. Let $\underline{q} = \min_i q_i$. It follows that there is a set $S(0)$ of students who are in the top $\underline{q}$ at all schools, and for all $m > \underline{q}$ there is a student s^m who is m-th ranked at all schools. Hence there is a unique regret-free stable matching, and Mechanism 1 essentially performs serial dictatorship and tells each student her (unique) budget set in order of her rank. When students don't have indifferences in their preferences, it follows that at each step some students reject all the schools that proposed to them except one, and so when the algorithm terminates it outputs a regret-free stable matching.

Let us show that there is a unique regret-free stable matching. Fix v and let $\mu, \mu' \in M^{RF}(v)$. Let $s^1, s^2, \ldots, s^{\underline{q}}$ be an arbitrary ordering of the students in $S(0)$ and s^m be the m-ranked student for all $m > \underline{q}$. Then for all $m > \underline{q}$ we can define $B^{s^m}(\mu|v)$ by $B^{s^m}(\mu|v) = \left\{i \mid \sum_{s'=s^{m'}, m'<m} \mathbf{1}\left\{D^{s'}\left(B^{s'}(\mu)|v\right) = i\right\} < q_i\right\}$, i.e. the budget set of s^m is the set of schools with residual capacity once all students ranked higher than s^m have chosen their school from their budget set. Hence by induction $B^s(\mu|v) = B^s(\mu'|v)$ for all s and so $\mu(s) = \mu'(s)$ for all s. □

Theorem 5 demonstrates that for certain priorities and preferences iterative Deferred Acceptance, which is a choice-based mechanism, can discover market-clearing cutoffs in a regret-free stable manner. This mechanism is iterative, and one can show that this is necessary: even under the conditions laid out in Theorem 5, no one-shot mechanism — choice-based or otherwise — can be regret-free stable. We provide an example in Appendix B.1.

Furthermore, the matching found by the school-proposing DA mechanism is not student-optimal. Can one find a student-optimal matching in a regret-free manner? The original proof of Gale and Shapley of the existence of a student-optimal stable matching was constructive: they furnished an algorithm, the Gale-Shapley Deferred Acceptance (DA) algorithm, and demonstrated that it always finds the student-optimal stable matching in polynomial time. An analogous algorithm for identifying the student-optimal regret-free stable matching could also be defined in our incomplete information setting, but would require students to provide information about both their priors and values and may induce students to acquire information in a way that incurs regret. In fact, for many economies, verifying that the student-optimal regret-free stable matching is the student-optimal one necessitates incurring regret with positive probability in the inspection process. This is because in regret-free stable matchings students cannot inspect outside of their budget set, and in many economies with positive probability the student-optimal regret-free stable matching does not provide students with their full budget set. We provide an example that admits no student-optimal regret-free stable mechanism in Appendix B.2. The intuition behind the example is that the act of verification requires some student to perform more information acquisition than is allowed under the regret-free stable inspection policy, and if the existing matching is student-optimal it is then costly for them to acquire the necessary information. This mirrors the Grossman-Stiglitz paradox, whereby under costly information acquisition equilibrium market prices cannot be stable, as this would eliminate the benefit of acquiring this information (Grossman and Stiglitz, 1980).

We next show that the conditions of Theorem 5 are necessary, in the sense that there is no (even non-choice-based) mechanism that is regret-free stable for general economies. In the more general case where students can suffer regret by inspecting the schools out of order, it may be impossible for any (even multi-round) mechanism to find a regret-free stable matching without incurring regret. Perhaps more fundamentally, this example shows that any mechanism that converges to a stable matching in a regret-free manner (such as school-proposing DA) relies heavily on an assumption that there always exist students willing to inspect some schools in their budget set.

Theorem 6. *Let M be a mechanism. Then there exists an economy $E = (\mathcal{C}, \mathcal{S}, q)$ such that, when each student reports $I^s = (F^s, c^s)$ truthfully, with positive probability mechanism M does not implement a regret-free stable matching.*

Remark. For convenience, we state and prove Theorem 6 for finite economies; we note that the result can be extended to continuum economies with minor adjustments.

Proof. Consider an economy E with three schools $\mathcal{C} = \{1, 2, 3\}$ with capacities $q_1 = q_2 = q_3 = 2$

and three students $S = \{x, y, z\}$.[9]

Suppose that school priorities are given by

$$\text{priority at } 1 : r_1^y > r_1^z > r_1^x$$
$$\text{priority at } 2 : r_2^z > r_2^x > r_2^y$$
$$\text{priority at } 3 : r_3^x > r_3^y > r_3^z,$$

and that student values at each school are $U[0,1]$ variables, i.e. with priors $F_i^s(x) = x \forall x \in [0.1]$ and student costs for inspection are given by $c_1^x = c_2^y = c_3^z = 0.1$, $c_2^x = c_3^y = c_1^z = 0.2$ and $c_3^x = c_1^y = c_2^z = 0.3$. Note that this means $\underline{v}_1^x = \underline{v}_2^y = \underline{v}_3^z = \sqrt{1-0.2} \approx 0.89$, $\underline{v}_2^x = \underline{v}_3^y = \underline{v}_1^z = \sqrt{1-0.4} \approx 0.77$ and $\underline{v}_3^x = \underline{v}_1^y = \underline{v}_2^z = \sqrt{1-0.6} \approx 0.63$ and so the order in which students $\{x, y, z\}$ wish to inspect schools is exactly the reverse of their priority at each school, e.g. student x wishes to inspect 1 then 2 then 3, and have bottom, middle and top priority out of $\{x, y, z\}$ at those schools respectively.

We will show that for all students s, there exists a school $i = \beta(s)$ such that, with positive probability, in every regret-free stable matching $\mu \in M^{RF}(E)$ student s only inspects i. Also, with positive probability, in every regret-free stable matching $\mu' \in M^{RF}(E)$ school i is not in student s's budget set $B^{\mu'}(s)$. To see why this implies the theorem, note that under Mechanism M, one of x, y, z must be the first student in $\{x, y, z\}$ to perform an inspection with positive probability. Without loss of generality we may suppose that student is x. If x first inspects $\beta(x)$ then with positive probability, in any regret-free stable matching $\mu' \in M^{RF}(E)$ student x regrets her inspection. If x first inspects some school other than $\beta(x)$ then with positive probability for any regret-free stable matching $\mu \in M^{RF}(E)$ student x regrets her inspection. Hence with positive probability there exists a student who regrets her inspection process, and so with positive probability M does not implement any regret-free stable matching

[9] Note that strictly speaking, as we assumed that there are more students than seats, the economy should have seven students $S=\{x, y, z, d_1, d_2, d_3, d_4\}$ where the d_i are four dummy students who have lower priority at every school than the students in $\{x, y, z\}$ and who have arbitrary preferences. For simplicity we omit these students in the description of the economy; however note that the proof applies as written to both economies.

We now turn to proving the claim: for all students s, there exists a school $i = \beta(s)$ such that with positive probability every regret-free stable matching involves student s inspecting only i, and with positive probability no regret-free stable matching has school i in student s's budget set. In particular, we show that $1 = \beta(x)$ satisfies the required properties. Note that since all priorities and costs are symmetric, the same arguments can be used to show that $2 = \beta(y)$ and $3 = \beta(z)$ satisfy the required properties.

Consider the event that $v_1^x, v_2^y, v_3^z \geq 0.9$, and $v_1^y, v_3^y, v_1^z, v_2^z \leq 0.5$. Note that it then holds that $D^x(\mathcal{C}) = 1$, $D^y(\mathcal{C}) = 2$ and $D^z(\mathcal{C}) = 3$. Now it is easy to check that for all $\omega \in X$ the only regret-free stable matching $\mu \in M^{RF}(E)$ is $(\mu(x), \mu(y), \mu(z)) = (1, 2, 3)$, since if any school i was not assigned to any of x, y, z (i.e. $\mu(i) \cap \{x, y, z\} = \emptyset$) then it would form a blocking pair with the student in $\{x, y, z\}$ whose top choice school is i. Hence $B^\mu(x) = \mathcal{C}$, so x inspects school 1 first, and since $v_1^x \geq 0.9 > \underline{v}_2^x, \underline{v}_3^x$ it follows that x only inspects school 1.

Next consider the event that $v_1^x, v_1^y, v_1^z \geq 0.9$ and $v_2^y, v_3^y, v_2^z, v_3^z \leq 0.5$. Note then that $D^y(\mathcal{C}) = D^z(\mathcal{C}) = 1$. Moreover, since $\underline{v}_1^y \geq \underline{v}_2^y, \underline{v}_3^y \geq 0.5 \geq v_2^y, v_3^y$ and y has top priority at 1 it follows that in any regret-free stable matching $\mu \in M^{RF}(E)$ student y inspects 1, and since $D^y(\mathcal{C}) = 1$ student y is assigned to 1, i.e. $\mu(y) = 1$. Since $q_1 = 2$ and z has second priority at 1 a similar argument shows that $\mu(z) = 1$. Hence for all regret-free stable matchings $\mu \in M^{RF}(E)$ it follows that 1 is full of students it prefers to x, and $1 \notin B^\mu(x)$.

Note that if no student has performed any inspections then we are unable to discern whether either these events is true, and for any student any inspection they perform will incur regret in either one event or the other, i.e. any inspections incurs regret with positive probability.

This example shows that there does not exist any mechanism that always finds a regret-free stable matching in a regret-free manner. This makes it even more surprising that, for any realization of preferences, the set of regret-free stable matchings $M^{RF}(E)$ not only has a student-optimal member, but also inherits the lattice structure induced by the deterministic economy with students $\Psi(s)$. □

To build additional intuition for Theorem 6, let us briefly demonstrate why Mechanism 1 might not be regret-free in a general economy. In the first round of Mechanism 1, every student s is proposed to by all schools except $\beta(s)$. When students were willing to inspect all schools in their budget set this was enough to induce some inspections and rejections. However, in general, there is an inspection *order* that maximizes the resulting expected payoff, and with positive probability students do not inspect school some schools in their budget set. Hence even though every student knows some of the schools in their budget set, no student s wants to start inspecting schools until she knows for sure whether $\beta(s)$ is in her budget set. In other words, it is *strictly optimal* for each student to wait until the mechanism forces them to perform inspections.

This intuition illustrates a more general principal: in the presence of costly information acquisition, iterative mechanisms without an activity rule may result in an *information deadlock*, where no actions are taken because every agent can achieve higher utility if another agent acts first.

5.3 Regret-Free Learning

While it may not always be possible to discover market-clearing cutoffs through observed choice, the structure of regret-free stable matchings gives us hope that they can still be learned and implemented in an approximate manner. We show that when we have sufficient initial information or market structure, cutoff mechanisms that use estimated cutoffs implement outcomes that are regret-free stable with respect to slightly perturbed school capacities.

Before formalizing these ideas, we first turn to the following question: How do we estimate population demand? In matching markets with one-sided incomplete information, the preferences of the side with full information are a key source of information. For example, in Theorem 5, condition (2) guarantees that at any point in iterative Deferred Acceptance there are students who have full information about their budget set. For more general priority structures, there will be students who have such information at the outset of the mechanism.

Definition 9. Let $E = (\mathcal{C}, S, q)$ be an economy. A student s has **free market infor-**

mation if for all schools i student s is either in the top q_i percentile of students or the bottom $1 - \sum_j q_j$ percentile of students, i.e. $\forall i \; r_i^s \notin \left[1 - \sum_j q_j, 1 - q_i\right)$. We let $S^f(E) = \left\{s \,|\, \forall i \; r_i^s \notin \left[1 - \sum_j q_j, 1 - q_i\right)\right\}$ denote the set of students with free market information in E.

Knowing only their priors and school priorities, students with free market information can determine both their budget sets and their preferences in a regret-free manner.

Hence we may estimate market demand as follows. In some markets historical demand is sufficient for estimating population demand. For example, in college admissions in many countries aggregate student demand for different university courses do not vary much from year to year and historical demand can be used to estimate current demand. Even when such prior information is not available, as long as there are students with free information we can start learning about student preferences. For example, if running iterative Deferred Acceptance assigns some students before reaching a deadlock, the demand of the assigned students could be used to estimate the demand of the remaining students.

Estimated Cutoffs are Robust

We now formalize the claim that outcomes of cutoff mechanisms are robust to errors in estimated demand. The intuition behind these results is that demand with costly information acquisition satisfies WARP (Proposition 1), and so all questions about cutoff mechanisms under costly information acquisition reduce to analogous questions about cutoff mechanisms in markets without costly information acquisition.

We first show that the outcome $\left(\mu^P, \chi^P\right)$ from posting cutoffs P when using Mechanism 1 is regret-free stable for capacities q^P that are slightly perturbed from the true capacities, and differs from the regret-free stable assignment under q for only a small number of students.

Theorem 7. *Let $\mathcal{E}$ be a continuum economy, let μ be a regret-free stable matching for $\mathcal{E}$ corresponding to market-clearing cutoffs P^*, and let $\overline{q} = D(P^*)$ be the measures of seats assigned under μ. Let $\left(\mu^P, \chi^P\right)$ be the outcome of running Mechanism 1 on $\mathcal{E}$ with cutoffs P,* *and for all i let $q_i^P = \left|\left\{s \mid \mu^P(s) = i\right\}\right|$. Then $\left(\mu^P, \chi^P\right)$ is regret-free stable with respect to q^P, $\left\|q^P - \overline{q}\right\|_2 \leq \|P - P^*\|_2$, and $\left|\left\{s \mid \mu(s) \neq \mu^P(s)\right\}\right| \leq \|P - P^*\|_2$.*

Proof. Note that by definition $q^P \equiv D(P)$. Now it is easy to see that $\left(\mu^P, \chi^P\right)$ is regret-free stable with respect to $D(P)$. Moreover, each student's assignment $\mu(s)$ is equal to their demand $D^s(P)$, which is determined by their budget set $B^s(P)$. Finally, in moving from P to P^* only $\|P - P^*\|_2$ students receive different budget sets. The result follows. □

When the error in the estimated cutoffs is due to sampling error, the outcome is regret-free stable for capacities that are normally distributed around the market-clearing demand.

Definition 10. For a capacity vector $q' = (q'_1, \ldots, q'_n)^T$ we let $\Sigma^{q'}$ denote the matrix with entries

$$\Sigma_{ij}^{q} = \begin{cases} -q_i q_j & \text{if } i \neq j \\ q_i(1 - q_i) & \text{if } i = j. \end{cases}$$

Proposition 2 (Distribution of approximately feasible capacities)**.** *Suppose the continuum economy $\mathcal{E}$ admits a unique stable matching μ with cutoffs P^*, and let $\overline{q} = D(P^*|\eta)$.*[10] *Let $E^k = \left(\eta^k, q^k\right)$ be a randomly drawn finite economy, with k students drawn independently according to η, ω drawn independently, and where $q^k = D\left(P^*|\eta^k\right)$ is defined so that P^* is a market-clearing cutoff for E^k. Then*

$$\sqrt{k} \cdot \left(q^k - q\right) \xrightarrow{d} \mathcal{N}\left(0, \Sigma^{\overline{q}}\right),$$

where $\mathcal{N}(\cdot|\cdot)$ denotes a C-dimensional normal distribution with given mean and covariance.

Proof. The result follows from the central limit theorem, as $D\left(P^*|\eta^k\right) = \frac{1}{k}\sum_{a=1}^k X_a$, where $X_a = D^\theta(P^*)$ is a random variable with $\theta \sim \eta$ capturing the demand of a single student drawn randomly from η and the X_a are independently drawn. □

Hence the mapping from cutoffs to demand is continuous, so approximate cutoffs yield regret-free stable outcomes for approximately feasible capacities. We similarly show in Ap-

[10] Note that $\overline{q}_i = q_i$ for all overdemanded schools i, i.e. those such that $P_i^* > 0$.

pendix C.1 that the mapping from demand to market-clearing cutoffs is continuous, and estimated cutoffs are robust to errors due to sampling. Thus in order to obtain a desirable estimate of the market-clearing cutoffs P it suffices to furnish an accurate estimate of demand D.

Examples

These results suggest that if we use cutoff mechanisms based on estimated population demand, the resulting outcomes will be robust to small biases or noise due to sampling error. We illustrate this intuition in the following examples.

Example 2. In this example, we show how to implement an approximately regret-free stable matching in a setting with historical demand data. Suppose that this year's economy $E = \left(\mathcal{C}, S^k, \lfloor qk \rfloor\right)$ is given by drawing k students independently from a distribution η, and last year's economy $E^{hist} = \left(\mathcal{C}, S^{hist}, \lfloor \alpha qk \rfloor\right)$ is given by drawing αk students independently also from η for some fixed $\alpha > 0$. Then the student-optimal market-clearing cutoffs $\hat{\boldsymbol{P}}$ for the economy E^{hist} give an unbiased estimator for the student-optimal market-clearing cutoffs both for E and for $\mathcal{E} = (\mathcal{C}, \mathcal{S}, \eta, q)$, and we can show that in this year's economy E posting $\hat{\boldsymbol{P}}$ implements a regret-free stable matching with respect to capacities $\hat{q}^k k$ that are close to qk. Specifically, if we let P^* be the market-clearing cutoffs of $\mathcal{E} = (\mathcal{C}, \mathcal{S}, \eta, q)$ and $\overline{q} = D\left(P^* | \eta\right)$ then we can use classic results about the convergence of two-step estimators[11] to show

$$\sqrt{k}\left(\hat{q}^k - \overline{q}\right) \xrightarrow{d} \mathcal{N}\left(0, \frac{(\alpha+1)}{\alpha}\Sigma^{\overline{q}}\right),$$

for $\Sigma^{\overline{q}}$ defined as in Definition 10. The full proof can be found in Appendix C.2.

For large economies the capacities that make the outcome regret-free stable converge to the true capacities, and the variance depends only on $\overline{q}$ and α. Moreover, in the absence of historical information ($\alpha = 0$) the cutoff mechanism can perform arbitrarily poorly, whereas more accurate historical information ($\alpha \to \infty$) leads to smaller perturbations in the capacities.

[11] See, e.g. Newey and McFadden (1994) for details.

Example 3. Suppose that this year's economy $E = \left(\mathcal{C}, S^k, qk\right)$ is given by drawing k students independently from a distribution $\eta(\Gamma^*)$, where student demand $D^s(P|\eta(\Gamma))$ is parametrized by $\Gamma = (\gamma_1, \gamma_2, \ldots, \gamma_n)$. Suppose also that some positive fraction α of students have free market information. Then by first obtaining preferences from the students with free market information, we can estimate Γ^* and provide an estimate $\hat{\boldsymbol{P}}$ for the student-optimal market-clearing cutoffs for E. We can also show that posting $\hat{\boldsymbol{P}}$ implements a regret-free stable matching with respect to capacities $\hat{q}^k k$ that are close to qk.

Formally, consider the mechanism M^F that runs in two rounds. In the first round it proposes to all students $s \in S^f(E)$, assigns them each to their chosen school and obtains their aggregate demand $\hat{q}_i^f k$ for each school i, and uses this demand to provide an estimate $\hat{\Gamma}$ for Γ^*. In the second round it runs Mechanism 1 with cutoffs $\hat{P}$ on a residual economy E^r computed as follows. The cutoffs $\hat{P}$ are the market-clearing cutoffs for an estimated residual economy $\hat{E}^r = (\mathcal{C}, S^r, \hat{q}^r k)$, where $\hat{E}^r$ is given by drawing $k - \left|S^f(E)\right|$ students without free market information in E independently from the distribution $\eta\left(\hat{\Gamma}\right)$ and the residual capacity for each school i is $\hat{q}_i^r = q_i - \hat{q}_i^f$. The residual economy $E^r = \left(\mathcal{C}, S^k \setminus S^f(E), \hat{q}^r k\right)$ is given by removing the students in $S^f(E)$ from E and reducing capacities accordingly at their assigned schools.

Let α denote the measure of students in E who have free market information, let $D_i^f(\Gamma)$ denote the proportion of students in $S^f(E)$ who demand school i as a function of Γ, and let $q_i^f = \alpha D_i^f(\Gamma^*)$ be the target first-round capacities. Define target capacities $\overline{q} = D(P^*|\eta(\Gamma^*))$ in terms of the market-clearing cutoffs P^* of $\mathcal{E} = (\mathcal{C}, \mathcal{S}, \eta(\Gamma^*), q)$. We can show that Mechanism M^F implements a regret-free stable matching μ with respect to perturbed capacity $\hat{q}^k k$, where

$$\sqrt{k}\left(\hat{q}^k - \overline{q}\right) \xrightarrow{d} \mathcal{N}\left(0, \Sigma^{\overline{q}} + 2\left(\frac{1}{\alpha}A + I\right)\Sigma^{q^f} A^T\right)$$

as $k \to \infty$ for $A = \nabla_\Gamma D(\Gamma^*)\left(\nabla_\Gamma D^f(\Gamma^*)\right)^{-1}$ and $\Sigma^{\overline{q}}, \Sigma^{q^f}$ defined as in Definition 10. (Note that $\alpha \leq \min_i q_i + 1 - \sum_i q_i \overset{def}{=} \alpha^*$, and $\alpha = \alpha^*$ is achieved when schools have aligned preferences, i.e. condition (2) in Theorem 5.) The idea is that since in the first round we assign only students in $S^f(E)$, the budget set and demand of these students is that same whether we assign them in the first round, or in the second round after posting the cutoffs $\hat{P}$. Hence the outcome after both rounds is regret-free stable with respect to realized demand. Convergence and variance expressions can be derived using two-step GMM. The full proof can be found in Appendix C.2.

For large economies the capacities that make the outcome regret-free stable converge to the true capacities, and the variance depends on $\overline{q}$, q^f, $A = \nabla_\Gamma D(\Gamma^*)\left(\nabla_\Gamma D^f(\Gamma^*)\right)^{-1}$ and α. Moreover, in the absence of free market information ($\alpha = 0$) the cutoff mechanism can perform arbitrarily poorly, whereas priorities that yield more students with free market information ($\alpha \to \alpha^*$) or more accurate estimates of Γ^* ($A \to 0$) lead to smaller perturbations in the capacities. Finally, the first round in this mechanism corresponds to the first round of iterative school-proposing Deferred Acceptance. If we allow for further rounds of proposals, we can further reduce the noise in the perturbed capacities.

6 Discussion

Summary of findings. We have proposed regret-free stability as a suitable solution concept in matching markets with costly information acquisition. We have also shown that, surprisingly, regret-free stable matchings always exist and the set of regret-free stable matchings has a lattice structure. However, we have also shown that the effect of costly information acquisition is that it may be impossible to compute a regret-free stable matching in a regret-free manner, and that standard matching market mechanisms can result in information deadlocks. We have also provided some mechanisms for when we are willing to relax feasibility and provide varying amounts of information in order to achieve a regret-free outcome, and shown that for large economies they can be implemented by perturbing the capacities by $O\left(\sqrt{k}\right)$ students, where k is the number of students in the market.

Approximation algorithms. Our results demonstrate that in general there is a tradeoff between the regret of a mechanism, the feasibility of the solution, and the amount of information provided to the mechanism. We have provided one class of mechanisms that relax the feasibility constraint in order to achieve optimal regret. It may also be possible to relax the regret of the mechanism in order to achieve exact feasibility, or to increase the number of rounds of communication in order to better approximate both. We leave these questions open for future work.

Activity rules. We demonstrated that in the presence of costly information acquisition standard matching mechanisms can create situations where it is strictly optimal for every agent to wait for other agents to move first. This illustrates a more general principal, that in the presence of costly information acquisition iterative mechanisms will need an activity rule to converge. Another relevant question is what the appropriate design of activity rules would be for such situations.

Stable matchings. We have concentrated our efforts on mechanisms that implement regret-free stable matchings. However, we have also provided a more general notion of stability in incomplete information settings. Is this more general space of outcomes predictive and does it have attractive structural properties? We selected the class of regret-free stable matchings as they compare each agent's utility only with her own utility under other information acquisition strategies. However it may be possible to improve social welfare by moving to a stable matching that transfers utility from one student to another. Are there stable matchings that are more desirable than the student-optimal regret-free stable matching? Our notion of stability under incomplete information can also be naturally extended to settings with two-sided incomplete information, as well as to settings with more general models for costly information acquisition, such as rational inattention models, or other models where agents may refine their priors for a cost. All of these questions become much more interesting in these general settings. We leave them open for future investigation.

Practical market design. Finally, what implications do our results have for practical applications? Colleges in many countries, such as China, India and Australia post historical cutoffs for admission into college programs. Our results on mechanisms with historical cutoffs suggests that if colleges capacities are flexible this can eliminate unnecessary preference formation by applicants. In Israel colleges post a pair of cutoffs for each program; students above the higher cutoff are guaranteed admission, students below the lower cutoff are advised to consider other options, students between the cutoffs are advised to wait for further information on enrolment for that year, and the cutoffs are updated as students register for programs. This very closely mirrors our Accept-Waitlist-Reject mechanisms and suggests that they can be of practical use. Our result on information deadlock also brings to mind the behavior of participants reacting to activity rules such as deadlines and exploding offers in other markets. In markets such as job markets and Ph.D. admissions, participants often wait until the last minute before expressing their preferences. Clearly costly information acquisition is an important issue in many other markets, and we leave further investigation of the empirical and practical consequences for future work.

Bibliography

References

I. Ashlagi, M. Braverman, Y. Kanoria, and P. Shi. Communication requirements and informative signaling in matching markets. Technical report, Working Paper, 2018a.

I. Ashlagi, M. Braverman, Y. Kanoria, and P. Shi. Clearing matching markets efficiently: Informative signals and match recommendations. 2018b.

E. M. Azevedo and J. D. Leshno. A supply and demand framework for two-sided matching markets. *Journal of Political Economy*, 124(5):1235–1268, 2016.

H. Aziz, P. Biró, S. Gaspers, R. de Haan, N. Mattei, and B. Rastegari. Stable matching with uncertain linear preferences. In *International Symposium on Algorithmic Game Theory*, pages 195–206. Springer, 2016.

S. Bade. Serial dictatorship: The unique optimal allocation rule when information is endogenous. *Theoretical Economics*, 10(2):385–410, 2015.

C. Blair. The lattice structure of the set of stable matchings with multiple partners. *Mathematics of Operations Research*, 13(4):619–628, 1988.

I. Bo and R. Hakimov. The iterative deferred acceptance mechanism. 2017.

A. Chakraborty, A. Citanna, and M. Ostrovsky. Two-sided matching with interdependent values. *Journal of Economic Theory*, 145(1):85–105, 2010.

Y. Chen and Y. He. Information acquisition and provision in school choice. Technical report, mimeo, 2015.

J. Drummond and C. Boutilier. Elicitation and approximately stable matching with partial preferences. In *IJCAI*, 2013.

J. Drummond and C. Boutilier. Preference elicitation and interview minimization in stable matchings. In *AAAI*, pages 645–653, 2014.

U. Dur, R. G. Hammond, and T. Morrill. Identifying the harm of manipulable school-choice mechanisms. Technical report, 2016.

L. Ehlers and J. Massó. Matching markets under (in) complete information. *Journal of Economic Theory*, 157:295–314, 2015.

P. Eso and B. Szentes. Optimal information disclosure in auctions and the handicap auction. *Review of Economic Studies*, 74:705–731, 2007.

Y. A. Gonczarowski, N. Nisan, R. Ostrovsky, and W. Rosenbaum. A stable marriage requires communication. In *Proceedings of the Twenty-Sixth Annual ACM-SIAM Symposium on Discrete Algorithms*, pages 1003–1017. Society for Industrial and Applied Mathematics, 2015.

B. Gong and Y. Liang. A dynamic college admission mechanism in Inner Mongolia: Theory and experiment. Technical report, Working paper, 2016. 2, 10, 2016.

S. J. Grossman and J. E. Stiglitz. On the impossibility of informationally efficient markets. *The American economic review*, 70(3):393–408, 1980.

P. Harless and V. Manjunath. The importance of learning in market design. 2015.

A. Hassidim, D. Marciano-Romm, A. Romm, and R. I. Shorrer. "Strategic" behavior in a strategy-proof environment. Technical report, working paper, 2015.

S. Kadam. Interviewing in matching markets. Technical report, Working Paper, 2015.

A. Kapor, C. Neilson, and S. Zimmerman. Heterogeneous beliefs and school choice. 2016.

B. Kleinberg, B. Waggoner, and E. G. Weyl. Descending price optimally coordinates search. 2016.

D. Knuth. *Mariages Stables Et Leurs Relations Avec D'Autres Problèmes Combinatoires: Introduction À L'Analyse Mathématique Des Algorithmes*. Presses de l'Université de Montréal, 1976.

R. S. Lee and M. Schwarz. Interviewing in two-sided matching markets. Technical report, National Bureau of Economic Research, 2009.

Q. Liu, G. J. Mailath, A. Postlewaite, and L. Samuelson. Stable matching with incomplete information. *Econometrica*, 82(2):541–587, 2014.

F. Matějka and A. McKay. Rational inattention to discrete choices: A new foundation for the multinomial logit model. *American Economic Review*, 105(1):272–98, 2015.

P. Milgrom and R. Weber. A theory of auctions and competitive bidding. *Econometrica*, 50 (5):1089–1122, September 1982.

W. K. Newey and D. McFadden. Large sample estimation and hypothesis testing. *Handbook of Econometrics*, 4:2111–2245, 1994.

W. Olszewski and R. Weber. A more general pandora rule? *Journal of Economic Theory*, 160: 429–437, 2015.

B. Rastegari, A. Condon, N. Immorlica, and K. Leyton-Brown. Two-sided matching with partial information. In *Proceedings of the fourteenth ACM conference on Electronic commerce*, pages 733–750. ACM, 2013.

B. Rastegari, A. Condon, N. Immorlica, R. Irving, and K. Leyton-Brown. Reasoning about optimal stable matchings under partial information. In *Proceedings of the fifteenth ACM conference on Economics and computation*, pages 431–448. ACM, 2014.

A. E. Roth. *Who Gets Whatâand Why: The New Economics of Matchmaking and Market Design*. Houghton Mifflin Harcourt, 2015.

A. E. Roth and M. Sotomayor. Two-sided matching. *Handbook of game theory with economic applications*, 1:485–541, 1992.

C. A. Sims. Implications of rational inattention. *Journal of Monetary Economics*, 50(3): 665–690, 2003.

J. Steiner, C. Stewart, and F. Matějka. Rational inattention dynamics: Inertia and delay in decision-making. *Econometrica*, 85(2):521–553, 2017.

M. L. Weitzman. Optimal search for the best alternative. *Econometrica: Journal of the Econometric Society*, pages 641–654, 1979.

Appendices

A Regret-Free Mechanisms in Extensive Form

In this section, we provide three extensive-form game descriptions of regret-free mechanisms. We first define general mechanisms as extensive-form dynamic games of imperfect information. We note that as regret-free mechanisms are incentive-compatible, we may restrict our attention to direct mechanisms, where students need only report either their type, or given a set of schools to inspect need only inspect that set of schools and truthfully report their inspected values. We then formally define choice-based messages, where we restrict the messages from students to be only choice-based information about their own preferences, such as choice functions and partial orders, and we restrict the actions and messages of the mechanism designer to use only the choice-based information. Finally, we formally define Accept-Waitlist-Reject (AWR) mechanisms, which restrict the mechanism designer to only tell students which schools will definitely be in their budget set, which schools definitely will not, and which schools are uncertain.

A.1 General Mechanisms

We formally define general mechanisms as dynamic games of incomplete information. There is a set of players: $s \in \mathcal{S}$, possible actions: $a \in \mathcal{A}$, and possible messages $m \in \mathcal{M}$. There is a set of nodes $\mathcal{Z}$, with initial node z^0 and terminal nodes T. At each node $z \in \mathcal{Z}$ there is a message history $\mathcal{M}_z = \{\mathcal{M}_z^s\}_{s \in \mathcal{S}}$; at each non-terminal node there is a set of active students $S(z)$ who are sent new messages $\{\iota_z^s = \iota^s(\mathcal{M}_z)\}_{s \in S(z)}$. At each node z every student has private information about their history $h_z^s = (\{a_{z'}^s, \mathcal{M}_{z'}^s, \iota_{z'}^s\})$ of their actions. These actions result in inspections $\chi = \chi_z^s$ and values $v^s|_\chi$, which are also privately known to the student.

There is a partition $\mathcal{H}$=$\{H_1, H_2, \ldots\}$ of the nodes into information sets, which represent the information available to the students. Each student s has a partition $\mathcal{H}^s = \{H_1^s, \ldots\}$ such that any two nodes z, z' are in the same information set (i.e. $z, z' \in H_i^s$) for some i if and only if, up to relabeling of other students, they have the same information state $\mathcal{M}_z^s = \mathcal{M}_{z'}^s$ and history $h_z^s = h_{z'}^s$ Note that for a given student s the mechanism designer only knows r^s as well as $\mathcal{M}^s$, and so if students s, s' satisfy $(r^s, \mathcal{M}_z^s) = \left(r^{s'}, \mathcal{M}_z^{s'}\right)$ then they are *indistinguishable to the mechanism designer at node* z. Each information set H_i has a set of active students $S = S(H_i)$ of positive measure such that $s \in S \Rightarrow H_i \in \mathcal{H}^s$ and if two students s, s' are indistinguishable to the mechanism designer at any node $z \in H_i$ then $s \in S \Leftrightarrow s' \in S$.

The available actions $A^s(H_i) \subseteq \mathcal{A}$ for the set of active students $s \in S(H_i)$ at an information set H_i are as follows. Students first inspect some subset of uninspected schools $\{i : \chi_z^s = 0\}$, where the subset can be adaptively chosen based on the observed values of other schools inspected at that node. Students then report a message m to the mechanism.

At terminal nodes $z \in T$ the mechanism outputs a matching μ_z and inspections χ_z, where χ_z^s is consistent with history h_z^s.

We let Σ^s denote the set of strategies for student s, i.e. an action $a \in A^s(H_i)$ for each history H_i such that $s \in S(H_i)$.

Definition 11. We say that a general mechanism is **regret-free stable** if at all terminal nodes $z = (\mu_z, \chi_z) \in T$ the matching μ_z is regret-free stable with any underlying economy consistent with the mechanism designer's current information state $\mathcal{M}_z$, and $\chi_z = \chi_z^{RF}(\mu|\cdot)$.

A.2 Choice-Based Mechanisms

We formally define choice-based mechanisms as dynamic games of incomplete information as follows.

There is a set of players $s \in \mathcal{S}$, possible actions, $a \in \mathcal{A}$, and possible messages $m \in \mathcal{M}$. The main restriction of a choice-based mechanism is that $\mathcal{M}$ is restricted to be the set of *choice-based information states* $\mathcal{I} = \{\succeq^s\}_{s \in \mathcal{S}}$. There is a set of nodes $\mathcal{Z}$, with initial node z^0, and terminal nodes T. At each node $z \in \mathcal{Z}$ there is a *partial message history, given by a single choice-based information state* $\mathcal{I}_z = \{\succeq_z^s\}_{s \in \mathcal{S}}$; at each non-terminal node there is a set of active students $S(z)$ who are sent new messages $\{\iota_z^s = \iota^s(\mathcal{I}_z)\}_{s \in S(z)}$ based on the information state

$\mathcal{I}_z$. At each node each student has private information about their history $h_z^s = (\{a_{z'}^s, \mathcal{I}_{z'}^s, \iota_{z'}^s\})$ of actions. These actions result in inspections $\chi = \chi_z^s$ and values $v^s|_\chi$, which are also privately known to the student.

There is a $\mathcal{H}$=$\{H_1, H_2, \ldots\}$ of the nodes into information sets. Each student s has a partition $\mathcal{H}^s = \{H_1^s, \ldots\}$ such that any two nodes z, z' are in the same information set $z, z' \in H_i^s$ for some i if and only if, up to relabeling of other students, they have the same choice-based information state $\mathcal{I}_z^s = \mathcal{I}_{z'}^s$ and history $h_z^s = h_{z'}^s$. Note that for a given student s the mechanism designer only knows r^s as well as $\mathcal{I}^s$, and so if students s, s' satisfy $(r^s, \mathcal{I}_z^s) = \left(r^{s'}, \mathcal{I}_z^{s'}\right)$ then they are indistinguishable to the mechanism designer at node z. Each information set H_i has: (1) a set of active students $S = S(H_i)$ of positive measure such that $s \in S \Rightarrow H_i \in \mathcal{H}^s$ and if two students s, s' are indistinguishable to the mechanism designer at any node $z \in H_i$ then $s \in S \Leftrightarrow s' \in S$; and (2) a set of schools $c(H_i)$.

The available actions $A^s(H_i) \subseteq \mathcal{A}$ for the set of active students $s \in S(H_i)$ at an information set H_i are as follows. Students first inspect some subset of uninspected schools $\{i : \chi_z^s = 0\}$ in $c(H_i)$, where the subset can be adaptively chosen based on the observed values of other schools inspected at that node. Students then report a refinement of $\succeq_z^s$, which encodes their choice of schools in that set.

At terminal nodes $z \in T$ the mechanism outputs a matching μ_z and inspections χ_z, where χ_z^s is consistent with history h_z^s

We let Σ^s denote the set of strategies for student s, i.e. an action $a \in A^s(H_i)$ for each history H_i such that $s \in S(H_i)$.

Definition 12. We say that a choice-based mechanism is **regret-free stable** if at all terminal nodes $z = (\mu_z, \chi_z) \in T$ the matching μ_z is regret-free stable with any underlying economy consistent with the mechanism designer's current information state $\mathcal{I}_z$, and $\chi_z = \chi_z^{RF}(\mu|\cdot)$

A.3 Accept-Watilist-Reject Mechanisms

We formally define Accept-Waitlist-Reject (AWR) mechanisms as dynamic games of incomplete information as follows.

There is a set of players $s \in \mathcal{S}$, possible actions, $a \in \mathcal{A}$, and possible messages $m \in \mathcal{M}$. The main restriction of an AWR mechanism is that $\mathcal{M}$ is restricted to be the set messages of the form $\mathcal{C}^{\{A,W,R\}} \cup \mathcal{C}$, where each element of $\mathcal{C}^{\{A,W,R\}}$ oncodes a possible message from the mechanism to a student about the set of schools that accept them (A) as the school is definitely in their budget set, waitlist them (L) as the school may or may not be in their budget set, or reject them (R) as the school is definitely not in their budget set. There is a set of nodes $\mathcal{Z}$, with initial node z^0, and terminal nodes T. At each node $z \in \mathcal{Z}$ there is a *partial message history* m_z*, given by a single element* $m_z \in \mathcal{M}^{\mathcal{S}}$*,* with m_z^s representing the last message sent between the mechanism and each student. At each non-terminal node there is a set of active students $S(z)$ who are sent new messages $\left\{\iota_z^s \in \mathcal{C}^{\{A,W,R\}}\right\}_{s \in S(z)}$ based on the message history m_z. At each node each student has private information about their history $h_z^s = (\{a_{z'}^s, m_{z'}^s, \iota_{z'}^s\})$ of actions. These actions result in inspections $\chi = \chi_z^s$ and values $v^s|_{\chi}$, which are also privately known to the student.

There is a $\mathcal{H}$=$\{H_1, H_2, \ldots\}$ of the nodes into information sets. Each student s has a partition $\mathcal{H}^s = \{H_1^s, \ldots\}$ such that any two nodes z, z' are in the same information set $z, z' \in H_i^s$ for some i if and only if they have the same history $h_z^s = h_{z'}^s$. Note that for a given student s the mechanism designer only knows r^s as well as $\mathcal{I}^s$, and so if students s, s' satisfy $(r^s, m_z^s) = \left(r^{s'}, m_z^{s'}\right)$ then they are indistinguishable to the mechanism designer at node z. Each information set H_i has: (1) a set of active students $S = S(H_i)$ of positive measure such that $s \in S \Rightarrow H_i \in \mathcal{H}^s$ and $m_z^s \in \mathcal{C}^{\{A,W,R\}}$, and if two students s, s' are indistinguishable to the mechanism designer at any node $z \in H_i$ then $s \in S \Leftrightarrow s' \in S$; and (2) a message in $m_z \in \mathcal{C}^{\{A,W,R\}}$ for those students.

The available actions $A^s(H_i) \subseteq \mathcal{A}$ for the set of active students $s \in S(H_i)$ at an information set H_i are as follows. Students first inspect some subset of uninspected schools $\{i : \chi_z^s = 0\}$

such that i accepts them (i.e. the element of m_z^s corresponding to school i is A), where the subset can be adaptively chosen based on the observed values of other schools inspected at that node. Students then report their favorite school $i \in A$.

At terminal nodes $z \in T$ the mechanism outputs a matching μ_z and inspections χ_z, where χ_z^s is consistent with history h_z^s

We let Σ^s denote the set of strategies for student s, i.e. an action $a \in A^s(H_i)$ for each history H_i such that $s \in S(H_i)$. Note that in an AWR mechanism each student will be an active student at a node exactly once.

Definition 13. We say that an AWR mechanism is **regret-free stable** if at all terminal nodes $z = (\mu_z, \chi_z) \in T$ the matching μ_z is regret-free stable with any underlying economy consistent with the mechanism designer's current information state $\mathcal{I}_z$, and $\chi_z = \chi_z^{RF}(\mu|\cdot)$

B Examples demonstrating Impossibility of Regret-Free Stable Mechanisms

We now demonstrate that standard mechanisms can fail spectacularly in learing market-clearing cutoffs and alleviating the costs associated with information acquisition. Intuitively, in choice-based mechanisms students need to know other students' choices in order to determine their optimal inspection strategy, and so in general the student who performs the 'first' inspection will incur additional inspections costs. Standard Deferred Acceptance mechanisms, which are played as one-shot games where students submit their full preference lists, perform especially poorly, as students are given almost no information about their choices before deciding on their inspection strategy. While in some settings regret can be eliminated by allowing for multi-round mechanisms, we prove the stronger result that for general economies even multiple-round mechanisms must either incur regret, or create an information deadlock, where every student waits for others to acquire information first.

B.1 Direct One-Shot Mechanisms

To demonstrate the issues in computing regret-free stable matchings, let us first consider the case where all students are willing to inspect any school as long as it is in their budget set. We may view this as a setting where the costs affect *which schools* students are willing to inspect, but not the *order* in which they are willing to inspect them. It is clear that the standard implementation of Deferred Acceptance as a one-shot game will not be regret-free even for such students, as students' budget sets will depend on the preferences of other students, and so students who have low priority at the schools they prefer are likely to incur regret. We illustrate this in the following example.

Example 4. Consider a discrete economy $E = (\mathcal{C}, \mathcal{S}, q)$ with n students and n schools each with capacity $q_i = 1$. Suppose that school priorities are perfectly aligned, i.e. $r^s_i = r^s_j$ for all $s \in \mathcal{S}$, $i, j \in \mathcal{C}$, and students have random preferences and are willing to incur the cost to attend any school. Such demand can be rationalized e.g. by the priors $F^s_i(x) = 0$ for all $x \in [0, 1)$, $F^s_i(x) = \frac{1}{4}$ for all $x \in [1, 2)$, $F^s_i(x) = 1 - \frac{1}{2^k}$ for all $k \geq 1$ and $x \in \left[2^k, 2^{k+1}\right)$ and costs $c^s_i = 1$ for all $s \in \mathcal{S}$.

In any one-shot choice-based mechanism, a student s will have no regret only if she chooses to examine precisely the set of all schools not selected by higher-ranked students. This is because a student is willing to incur the cost to examine any school if and only if it is in her budget set. As student preferences are random, the probability that every student other than the highest-ranked student regrets her inspections is at least $\prod_i \left(1 - \frac{1}{\binom{n}{i}}\right) \geq \left(\frac{n-1}{n}\right)^{n-1} \to \frac{1}{e}$ all $n - 1$. The example can also be modified so that with probability $\to \frac{1}{e}$ a proportion $\to 1$ of students incur unbounded regret.[12]

This example demonstrates that single-shot choice-based mechanisms cannot hope to find regret-free stable matchings, even in settings where students are willing to incur the costs of searching any number of schools, due to their inability to coordinate the students' search.

[12]For each bound K the example can be modified so that with probability $\to \frac{1}{e}$ all $n - 1$ students other than the top priority student incur regret at least K times their utility.

B.2 Impossibility of Student-Optimal Regret-Free Stable Mechanisms

In this section we provide an example demonstrating that even in settings where it is possible to implement a regret-free stable choice-based mechanism, it may be impossible to verify that a matching is student-optimal without incurring regret.

Example 5. Consider an economy E with two schools $\mathcal{C} = \{1, 2\}$ with capacities $q_1 = q_2 = 1$ and 2 students $S = \{x, y\}$.[13]

Suppose that school priorities are given by

$$\text{priority at } 1 : r_1^y > r_1^x$$
$$\text{priority at } 2 : r_2^x > r_2^y.$$

Suppose also that student values at each school have discrete distribution $\mathbb{P}(v_i^s = 1) = \mathbb{P}(v_i^s = 2) = \frac{1}{4}$, $\mathbb{P}\left(v_i^s = 2^k\right) = \frac{1}{2^k}$ for all $k > 1$ and $\mathbb{P}(v_i^s = x) = 0$ for all $x \notin \left\{\frac{1}{2^k}\right\}_{k\in\mathbb{N}}$, i.e. with priors $F_i^s(x) = 0$ for all $x \in [0, 1)$, $F_i^s(x) = \frac{1}{4}$ for all $x \in [1, 2)$, $F_i^s(x) = 1 - \frac{1}{2^k}$ for all $k \geq 1$ and $x \in \left[2^k, 2^{k+1}\right)$, and that student costs for inspection are given by $c_1^x = c_2^y = 1$ and $c_2^x = c_1^y = 2$. As $\mathbb{E}[(v_i^s - \underline{v})] = \infty$ for all s, i and $\underline{v} \in \mathbb{R}$ it follows that both students' optimal strategies are to inspect all the schools that are available to them.[14]

Note that the matching $\mu = \mu^{school}$ defined by $(\mu(x), \mu(y)) = (2, 1)$ is always regret-free stable, and is the school-optimal regret-free stable matching. Let $\mu' = \mu^{student}$ be defined by $(\mu'(x), \mu'(y)) = (1, 2)$. We will consider two separate events. Let X denote the event that $v_1^x = v_2^y = 2$ and $v_2^x = v_1^y = 4$. Let X' denote the event that $v_1^x, v_2^y > 4$ and $v_2^x = v_1^y(\omega) = 4$. Note that μ^{school} is the student-optimal regret-free stable matching subject to event X, as both x and y obtain their highest valued schools, and that $\mu^{student}$ is the student-optimal regret-free stable matching subject to event X', as both x and y again obtain their highest valued schools.

[13] Strictly speaking, as we assumed that there are more students than seats, the economy should have three students S=$\{x, y, d\}$ where d is a dummy student who has lower priority at every school than the students in $\{x, y\}$ and who has arbitrary preferences. For simplicity we omit these students in the description of the economy; however note that the proof applies as written to both economies.

[14] It is simple to extend this example so that $v_i^s(\cdot)$ is continuous random variable with continuous density by smoothing the density for 2^k over the interval $\left[2^{k-1}, 2^k\right]$.

Notice that events X and X' are mutually exclusive, and that $\mathbb{P}(X) = \mathbb{P}(X') = \left(\frac{1}{4}\right)^4 > 0$. Furthermore, $v_2^x = v_1^y(\omega) = 4$ in either event. Thus, conditional on one of the events X or X' occurring, the only way to distinguish which event occurred is for student x to inspect school 1 or student y to inspect school 2.

We now first demonstrate why the existence of such X and X' shows that we cannot verify student-optimality in a regret-free manner. Note that if $\mu = \mu^{school}$ is the student-optimal regret-free stable matching then each school is assigned their top choice student, and so based on school preferences alone there are no blocking pairs and the corresponding student budget sets are $B^{\mu}(x) = \{2\}, B^{\mu}(y) = \{1\}$. Hence under $\chi^{RF}(\mu|\cdot)$ student x only inspects school 2, and student y only inspects school 1. However, if $\mu' = \mu^{student}$ is the student-optimal regret-free stable matching then under $\chi^{RF}(\mu'|\cdot)$ both students inspect both schools. Thus, since it is impossible to distinguish between events X and X' without requiring either student x to inspect school 1 or student y to inspect school 2, one of these inspections must occur in the event $X \vee X'$ in order to determine the student-optimal regret-free stable matching, which incurs regret under event X. In other words, it is impossible to verify that μ^{school} is the student-optimal regret-free stable matching without incurring regret. Since X has positive probability, we conclude that it is impossible to verify that the student-optimal regret-free stable matching is student-optimal without incurring regret with positive probability.

C Estimating Regret-Free Stable Cutoffs

C.1 Continuity and Convergence of Market-Clearing Cutoffs

We first define a metric on the space of economies and on the space of stable matchings. Fix a set of schools $\mathcal{C}$ and a set of students $\mathcal{S}$. We say that a sequence of continuum economies $\mathcal{E}^k = \left(\eta^k, q^k\right)$ converges to the continuum economy $\mathcal{E} = (\eta, q)$ if η^k converges in the weak sense to η, and $q^k \to q$. We define the distance between stable matchings to be the distance between their associated cutoffs, $d(\mu,\mu') = \max_{P,P':\mathcal{M}(P)=\mu,\mathcal{M}(P')=\mu'} \|P-P'\|$.[15] Given a finite economy $E=(\mathcal{C},S,q)$ define the continuum economy $\Phi(E)=(\mathcal{C},\mathcal{S},\eta,q)$ by taking the distribution η defined by

$$\eta\left(\left\{s\in\mathcal{S}\,|\,\theta^s=\theta^t,\ v_i^s\in\left\{v_i^t(\omega)\ |\ \omega\in X\right\}\right\}\right)=\frac{1}{|S|}p(X)\ \forall t\in S,\ X\subseteq\Omega.$$

We may think of this as first taking the empirical distribution $\sum_{t\in S}\frac{1}{|S|}\delta_t$ and then changing the point distribution δ_t for student t to mirror the possible distribution of values v^t. We say that a sequence of finite economies E^k converges to the continuum economy $\mathcal{E}$ if the embeddings $\Phi\left(E^k\right)$ converge to $\mathcal{E}$.

Theorem 8. *Suppose the continuum economy $\mathcal{E}$ admits a unique regret-free stable matching μ. Then the regret-free stable matching correspondence mapping economies to regret-free stable matchings is continuous at $\mathcal{E}$ within the set of continuum economies.*

Proof. The theorem follows from the analogous result in Azevedo and Leshno (2016) as well as observing that the set of regular measures is open. □

Theorem 9. *Suppose the continuum economy $\mathcal{E}$ admits a unique regret-free stable matching μ, and has a C^1 demand function that is non-singular at the market-clearing cutoffs (i.e. $\partial D(P^*)$ non-singular). Let $E^k=\left(\eta^k,qk\right)$ be a randomly drawn finite economy, with k students drawn independently according to η, and let P^k be a market-clearing cutoff of E^k. Then*

$$\sqrt{k}\cdot\left(P^k-P^*\right)\xrightarrow{d}\mathcal{N}\left(0,\partial D(P^*)^{-1}\cdot\Sigma^q\cdot\partial D(P^*)^{-T}\right),$$

where $\mathcal{N}(\cdot|\cdot)$ denotes a C-dimensional normal distribution with given mean and covariance matrix, and

$$\Sigma_{ij}^q=\begin{cases}-q_iq_j & \text{if } i\neq j\\ q_i(1-q_i) & \text{if } i=j.\end{cases}$$

[15] Note that if $\mathcal{E}$ is the embedding of a finite economy then there are many cutoffs that give the same matching μ.

Theorem 9 shows that the estimated cutoffs P^k are normally distributed around P^*, and follows directly from the analogous result in Azevedo and Leshno (2016). Another interpretation is that given and underlying population η and cutoffs P^*, if demand is given by sampling k students from η then the resulting market-clearing cutoffs P^k will be normally distributed around P^*.

C.2 Omitted Proofs for Section 5

C.2.1 Example 2

We show that

$$\sqrt{k}\left(\hat{q}^k - \overline{q}\right) \xrightarrow{d} \mathcal{N}\left(0, \frac{(\alpha+1)}{\alpha}\Sigma^{\overline{q}}\right).$$

Let X be a random variable that gives a student randomly drawn according to η with corresponding demand D^X, and with probability $\frac{\alpha}{\alpha+1}$ and $\frac{1}{\alpha+1}$ assigns them to be a 'past' student and 'present' student respectively. Let $m(X,P) = \mathbf{1}\{X \text{ is 'past'}\}\left(D^X(P) - \overline{q}\right)$ and let $g(X,q,P) = \mathbf{1}\{X \text{ is 'present'}\}\left(D^X(P) - q\right)$. Note that as $\lfloor qk \rfloor - qk \leq \frac{1}{k}$ it follows that $\sqrt{k}\hat{P}$ converges in distribution to $\sqrt{k}\hat{P}'$ where $\hat{P}'$ satisfies $\hat{m}\left(X,\hat{P}'\right) \stackrel{def}{=} \sum_{i=1}^{(\alpha+1)k} \frac{\alpha}{\alpha+1}\left(\frac{D^{(X_i)}(\hat{P}')}{k} - \overline{q}\right) = 0$. Note also that similarly $\sqrt{k}\hat{q}^k$ converges in distribution to $\sqrt{k}\hat{q}^k$ satisfying $\hat{g}\left(X,\hat{q}^k,P^*\right) \stackrel{def}{=} \sum_{i=1}^{(\alpha+1)k} \frac{1}{\alpha+1}\left(\frac{D^{(X_i)}(P^*)}{k} - \hat{q}^k\right) = 0$. Hence $\sqrt{k}\left(\hat{q}^k - \overline{q}\right) \xrightarrow{d} \mathcal{N}(0,V)$, where

$$V = (1+\alpha)\text{var}\left(g(X,\overline{q},P^*) - \frac{1}{\alpha}m(X,P^*)\right).$$

Since $\text{var}(g(\cdot)) = \frac{1}{1+\alpha}\text{var}\left(D^X(P^*)\right)$, $\text{var}(m(\cdot)) = \frac{\alpha}{1+\alpha}\text{var}\left(D^X(P^*)\right)$ and $\text{cov}(g(\cdot), m(\cdot)) = 0$, this is equal to $\text{var}\left(D^X(P^*)\right) + \frac{1}{\alpha}\text{var}\left(D^X(P^*)\right) = \frac{(1+\alpha)}{\alpha}\Sigma^{\overline{q}}$ as required.

C.2.2 Example 3

We first show that the outcome (μ,χ) after both rounds is regret-free stable with respect to realized demand $\hat{q}^k$. It suffices to show that for all students $s \in S^f(E)$ with free market information it follows that $\mu(s) = D^s\left(\hat{P}\right)$. Now since $s \in S^f(E)$ it holds that $\forall i\ r_i^s \notin$

$\left[1-\sum_j q_j, 1-q_i\right)$, and her first-round budget set is $B^s = \{i \mid r_i^s \geq 1-q_i\}$. Moreover, if $i \in B^s$ then if $\mu(i) = q_i$ it follows that $\hat{P}_i < r_i^s$, and so $i \in B^s\left(\hat{P}\right)$. Finally, if $i \in B^s\left(\hat{P}\right)$ then $1 - r_i^s \leq \sum_j q_j$, as all students find all schools acceptable and so if $r_i^{s'} \geq r_i^s$ then student s' is assigned to some school. In other words, $B^s \subseteq B^s\left(\hat{P}\right) \subseteq B^s$, and so $\mu(s) = D^s(B^s) = D^s\left(B^s\left(\hat{P}\right)\right) = D^s\left(\hat{P}\right)$.

We now show that

$$\sqrt{k}\left(\hat{q}^k - \overline{q}\right) \xrightarrow{d} \mathcal{N}\left(0, \Sigma^{\overline{q}} + 2\left(\frac{1}{\alpha}A + I\right)\Sigma^{q^f}A^T\right).$$

Let X be a random variable that gives a student randomly drawn according to η with corresponding demand D^X. Let the first-round cutoffs be $P_i^f = 1 - q_i$, let $m(X,\Gamma) = \mathbf{1}_{\{X \in S^F(E)\}} D^X\left(P^f | \eta(\Gamma)\right) - q^f$ and let $g(X,q,\Gamma) = m(X,\Gamma) + \mathbf{1}_{\{X \notin S^F(E)\}} D^X\left(P^*(\Gamma) | \eta(\Gamma^*)\right) - \left(q - q^f\right)$. Note that the estimated $\hat{\Gamma}$ satisfies $\hat{m}\left(X, \hat{\Gamma}\right) \stackrel{def}{=} \sum_{i=1}^k \alpha \frac{\left(D^{X_i}\left(P^f | \eta\left(\hat{\Gamma}\right)\right) | X \in S^F(E)\right)}{k} - q^f = 0$, and that the estimate demand $\hat{q}^k$ of all students satisfies $\hat{g}\left(X, \hat{q}^k, \hat{\Gamma}\right) \stackrel{def}{=} \sum_{i=1}^k \alpha \frac{\left(D^{X_i}\left(P^f | \eta\left(\hat{\Gamma}\right)\right) | X \in S^F(E)\right)}{k} + (1-\alpha)\frac{\left(D^{X_i}\left(P^*\left(\hat{\Gamma}\right) | \eta(\Gamma^*)\right) | X \notin S^F(E)\right)}{k} - \hat{q}^k = 0$. Hence $\sqrt{k}\left(\hat{q}^k - \overline{q}\right) \xrightarrow{d} \mathcal{N}(0, V)$, where

$$V = \text{var}\left(g(X, \overline{q}, \Gamma^*) - \left(I + \frac{1}{\alpha}A\right) m(X, \Gamma^*)\right)$$

and $A = \mathbb{E}\left[\nabla_\Gamma\left(g(X,\overline{q},\Gamma^*) - m(X,\Gamma^*)\right)\right] \mathbb{E}\left[\nabla_\Gamma m(X,\Gamma^*)\right]^{-1}$.

Note that $A = \nabla_\Gamma D\left(P^*(\Gamma) | \eta(\Gamma^*)\right)|_{\Gamma=\Gamma^*}\left(\nabla_\Gamma D^f(\Gamma^*)\right)^{-1}$. Moreover cov $(g(\cdot), m(\cdot))$=var $(m(\cdot))$ and so

$$V = \text{var}(g(\cdot)) + \frac{1}{\alpha}\text{var}(m(\cdot))\left(I + \frac{1}{\alpha}A\right)^T.$$

Since var $(g(\cdot))$=var $\left(D^X(\Gamma^*)\right)$=$\Sigma^{\overline{q}}$, var $(m(\cdot))$=2αvar $\left(D^f(\Gamma^*)\right)$=$2\alpha\Sigma^{q^f}$, this is equal to $\Sigma^{\overline{q}} + 2A\Sigma^{q^f}\left(1 + \frac{1}{\alpha}A\right)^T$ as required.