Information Acquisition Costs in Matching Markets

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

August 1, 2019

Abstract

How does the design of a marketplace affect the flow and acquisition of information in the market? We explore this question in a model of college admissions that formally accounts for students' information acquisition costs in forming their preferences. In this model students may rationally choose to remain partially informed, and we extend the notion of stability to this partial information setting. Our main question is whether the market can reach a stable outcome while facilitating efficient information acquisition by all students. To address this question, we define a regret-free stable outcome as a benchmark for efficient information acquisition.

We show that regret-free stable outcomes always exist, and are characterized by market-clearing cutoffs. Standard matching mechanisms can be seen as asking students to engage in price-discovery to discover these cutoffs, and this price-discovery process can be costly. Instead, we emphasize the usefulness of additional sources of information. We show that estimating cutoffs from historical market information allows for an approximately optimal outcome, and provide a procedure for bootstrapping cutoffs when historical market information is not available. Our results suggest that given sufficient information, matching markets can perform well even with simple or decentralized mechanisms.

[*]Microsoft Research, New England
[†]University of Chicago Booth School of Business
[‡]Stanford University, Department of Management Science and Engineering
[§]Microsoft Research, New England

1 Introduction

In this paper, we investigate how the design of a matching market affects how agents form their preferences. While the prevalence of incomplete information is well-studied in the context of auction markets (see e.g. Milgrom and Milgrom (2004)), it is relatively unexplored in matching market settings. In matching settings such as school choice and college admissions it is common for applicants to spend a significant amount of effort investigating potential placements before forming their preferences. This has motivated several informational interventions Corcoran et al. (2018); Dynarski et al. (2018); Grenet et al. (2019), and raises the question of how the design of the marketplace facilitates the acquisition and flow of information in the market.

To address this question, we present a model of many-to-one matching with costly information acquisition. In our model college priorities are common knowledge. Students have independent private values for each college which they can learn at a cost. We model each agent's information acquisition problem using the "Pandora's box" framework of Weitzman (1979). Students initially know their distribution of values for each college, and need to pay a cost to learn the realized value for a given college. When a student knows the set of colleges available to her, her information acquisition problem has an optimal solution in terms of a simple index policy. In the matching setting, the set of colleges available to a student depends on the decisions of other students. Thus the information acquisition decisions of different students are interlinked.

The model captures some important features of matching markets with costly information acquisition. Market outcomes consist of an assignment as well as the information acquired. This allows us to account for both the value of the match and the cost of information acquisition. In this model, students can rationally choose to remain partially informed. The model can also account for the cost of submitting a full list to DA, and rationalize a student's preference to delay decision-making, even when she has multiple offers of admission.

We extend stability to this partial information setting: an outcome for the market consists of a matching and acquired information, and is stable if there is no blocking pair, i.e. a (student, college) pair such that the student (i) has sufficiently high priority to attend the college, and the student either (ii) prefers the college to their assigned college, given their acquired information, or (ii') is

willing to pay the cost to learn whether they prefer the college over their assigned college. This definition extends the standard definition of stability, and is equivalent to the standard definition when students know their full preferences.

The model allows us to examine the role of the market in facilitating the flow and acquisition of information. Information acquisition can be inefficient since students need information about which colleges are available to them to know what information to acquire, and therefore need information about others' decisions. We formalize a benchmark for efficient information acquisition, which we term regret-free stability, under which every student has acquired information as if she had full access to all information collected from other students. We ask whether an appropriately designed market can facilitate efficient information acquisition in the sense that it achieves regret-free stable outcomes.

We first show that regret-free stable outcomes always exist and can be described using cutoffs. The set of regret-free stable outcomes forms a non-empty lattice, and there is a regret-free stable outcome that is unambiguously the best for all students. In addition, regret-free stable outcomes are characterized by market-clearing cutoffs: a college is available to a student if and only if her priority at the college meets the cutoff, each student acquires preferences optimally within the set of colleges for which she meets the cutoff, and induced aggregate demand exactly clears the supply of seats in the market. Hence the cutoffs function as prices, coordinating the matching of supply and demand. In order to facilitate efficient information acquisition a market must discover and communicate the market-clearing cutoffs to all students.

Standard matching mechanisms can be seen as asking students to engage in price-discovery processes to discover market-clearing cutoffs. In certain settings, such as when students are willing to inspect all colleges available to them, college-proposing Deferred Acceptance (DA) can learn the market-clearing cutoffs using regret-free information acquisition. Intuitively, college-proposing DA behaves like a descending-price auction that uses revealed preference information to discover the market-clearing cutoffs. This facilitates efficient information acquisition for two reasons: the side with more information (i.e. the colleges) propose first, and student preferences are such that there are always some students with sufficient information to continue optimally acquiring information.

However, we also show that in general settings standard matching mechanisms can fail to implement a regret-free stable outcome. In fact we show that many reasonable price-discovery processes can fail to implement a regret-free stable outcome. We consider canonical processes, in which a centralized platform dictates sequential information acquisition by individual students, choice-based processes, where the platform elicits only choices functions and ordinal preferences from students, and AWR processes, which induce information acquisition by informing students about whether each school accepts (A), rejects (R) or waitlists (W) them. We show that all three types of processes can result either in inefficient information acquisition or in information deadlocks, where no student has sufficient information to know how to optimally acquire information.

Instead of considering more complicated market mechanisms and processes, we instead emphasize the usefulness of having additional sources of information. For example, given oracle access to the market-clearing cutoffs, a market designer can implement a regret-free stable outcome simply by posting these cutoffs and committing to assigning students to their favorite college for which they meet the cutoff. We also provide practical suggestions for how to obtain additional market information and implement regret-free stable outcomes. If the market designer can estimate the market-clearing cutoffs, such as from historical demand, then posting the estimated cutoffs implements a regret-free stable outcome with respect to slightly perturbed capacities. We also provide a process for bootstrapping the cutoffs when historical market information is not available.

Our results suggest that given sufficient information, matching markets can perform well even with simple or decentralized mechanisms. For example, if admissions cutoffs are relatively constant from year to year, then a decentralized system where colleges post their admissions cutoffs and students apply separately can implement an approximately optimal outcome. Our results also shed light on the possible implications of costly information acquisition. In markets where students do not have access to market-clearing cutoffs, many reasonable market processes can create incentives for students to delay their decision-making processes, or cause students to acquire unnecessary information and regret their decisions to do so.

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 colleges 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); Liu (2018); Chen and Hu (2019); Bikhchandani (2017) suggest notions of stability under asymmetric information between agents that account for agents learning by observing that a given match is not blocked, and Chakraborty et al. (2010) consider agents with incomplete information who update their preferences after seeing the matching. We take an orthogonal approach in studying how unnecessary information acquisition costs can be minimized, in a setting where the lack of blocking does not impart any further information about preferences.

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. Grenet et al. (2019) provide an empirical analysis of student behavior in a university admissions context that is consistent with students undergoing costly information acquisition and accepting admissions offers in a manner that reduces regret.

Another strand of related work is the literature on costly search in single-agent search processes. Chade et al. (2014) studies the problem of deciding on an optimal simultaneous search portfolio in college admissions when values are stochastic and search is costly. Shorrer (2019) analyzes optimal simultaneous search in a setting where admissions to colleges are correlated. Our work builds on the adaptive search framework of Pandora's box consumer search as introduced by Weitzman (1979), where the agent adaptively selects a set of colleges at which they pay to learn their value, and then obtains the highest value amongst searched colleges. Doval (2018); Beyhaghi and Kleinberg (2019) explore the structure of optimal and near-optimal policies when agents can opt to attend a college without learning their value at the college. Our focus is on how to best design a marketplace where many single-agent search processes are coupled by a matching constraint.

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 by individuals and market-level information aggregation and assignment.

A related question is the communication complexity of transmitting known preferences 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.

Our work contributes to the growing number of papers exploring the use of sequential or multi-round college admissions 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 colleges 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. Echenique et al. (2016); Klijn et al. (2013) experimentally study dynamic implementations of the Deferred Acceptance mechanism. We show that such mechanisms are not always able to facilitate optimal information acquisition, and that additional market information may be needed to coordinate the interconnected search problems and mitigate the effects of costly information acquisition.

2 Model for Matching with Costly Information Acquisition

We present a model where students learn their preferences through costly information acquisition. The set of colleges is denoted by $\mathcal{C} = \{1, \ldots, n\}$, and each college $i \in \mathcal{C}$ has capacity to admit $q_i > 0$ students. A student is given by a tuple $s = (F^s, c^s, r^s, v^s)$. College priorities are publicly known,

and captured by the vector $r^s \in [0,1]^{\mathcal{C}}$. College 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 college i. Student s needs to perform costly information acquisition to learn her value for attending each college. Initially student s knows that the value for attending college 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). We assume that v_i^s is independently drawn across students and colleges.[1]

We will let a student's type $\theta = \theta^s = (F^s, c^s, r^s)$ denote the information that is initially known to student $s = (F^s, c^s, r^s, v^s)$. We refer to $\theta \in \Theta$ as a student type, and refer to $s = (\theta^s, v^s) \in \mathcal{S}$ 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 $\mathcal{S} = \Theta \times \mathbb{R}^{\mathcal{C}}$.

Let η be a measure over $\mathcal{S} = \Theta \times \mathbb{R}^{\mathcal{C}}$, and let ν be the induced measure of student types over Θ. We say that η is *consistent with initial information* if for any $A \subset \Theta$ and sets $V \subset \mathbb{R}^{\mathcal{C}}$ we have that

$$\eta\left(\{s = (\theta, v) \,|\, \theta \in A, v \in V\}\right) = \int_{\theta \in A} \int_{v \in V} dF^s\left(v\right) d\nu\left(\theta\right).$$

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 college, and η is a probability measure over $\mathcal{S}$ that is consistent with initial information.

Intuitively, there is a continuous mass of students of any given type θ, and although the realized preferences of an individual student are random, there is no aggregate uncertainty, and the aggregate distribution over $s = (\theta, v)$ is determined exactly by aggregating students' initial information F^θ.

We also define a rank distribution ζ and type distribution ν as follows. For every subset $R \subset [0,1]^{\mathcal{C}}$ of possible student ranks we let $\zeta\left(R\right) = \eta\left(\{s \mid r^s \in R\}\right)$, and for every subset $A \subset \Theta$ of possible student types we let $\nu\left(A\right) = \eta\left(\{s \mid \theta^s \in A\}\right)$.

We make the following assumptions. All students and colleges are acceptable. The rank r_i^s is normalized to be percentile rank, i.e. for any $i \in \mathcal{C}$ and $x \in [0,1]$, we have that $\eta\left(\{s \in \mathcal{S} | \, r_i^s \leq x\}\right) = x$. College priorities are strict, i.e. for any $x \in [0,1]$ and $i \in \mathcal{C}$ we have $\eta\left(\{s \in \mathcal{S} | \, r_i^s = x\}\right) = 0$; and there is an excess of students, $\sum_{i \in \mathcal{C}} q_i < \eta\left(\mathcal{S}\right) = 1$. For ease of exposition, we also assume that

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

student preferences are strict, by imposing that $\eta(\{s \in \mathcal{S} | \, \theta^s = \theta, \, v_i^s = v_j^s\}) = 0$ for any $i, j \in \mathcal{C}$ and $\theta \in \Theta$, and assuming that any further indifferences are resolved in favor of the college with lower index. To simplify notation for resolving such indifferences, if z_i and z_j are quantities related to colleges i and j respectively (e.g. v_i^s and v_j^s for some student s), we let $z_i \hat{>} z_j$ denote that either $z_i > z_j$ or $z_i = z_j$ and $i < j$.

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 college $i \in \mathcal{C}$ let $\mu(i)$ denote the set $\mu^{-1}(i) \subseteq \mathcal{S}$ of students assigned to college i. For each student $s \in \mathcal{S}$ and college $i \in \mathcal{C}$, let the *inspection indicator* χ_i^s be an indicator function that is 1 if student s has inspected college 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 college $i \in \mathcal{C}$ 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 $\chi_{\mu(s)}^s = 1$. This last condition is tantamount to assuming that a student must inspect a college in order to attend it, and is justified in settings where these costs are eventually borne should the student attend the school.[3] 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-college pair (s, i) forms a blocking pair if: (i) student s has sufficiently high priority to be assigned to college i, namely $r_i^s \geq \inf\left\{r_i^{s'} \mid s' \in \mu(i)\right\}$ or $\eta(\mu(i)) < q_i$; and (ii) student s inspected college i and knows she prefers college i over her assigned college $\mu(s)$, namely $\chi_i^s = 1$ and $v_i^s \hat{>} v_{\mu(s)}^s$.[4,5] When information acquisition is costly for students there may be a student-college pair (s, i) where (i) holds and student s did not inspect college i. We extend the standard definition and say that (s, i) also forms a blocking pair if (i) holds and (ii') s has not yet inspected college i and prefers to pay

[2] We are implicitly assuming that two students with the same type and values inspect the same colleges.

[3] For example, if a students' value is determined by the course offerings, the student must pay the cost of acquiring this information (i.e. the student must read the course catalogue) if she attends the school.

[4] Recall that feasibility requires that $\chi_{\mu(s)}^s = 1$.

[5] We assume students break ties in favor of the college with lower index; this implies (s, i) also blocks if $v_i^s = v_\mu^s(s)$ and $i < \mu(s)$.

the inspection cost c_i^s and be assigned to the better college of i and $\mu(s)$, namely $\chi_i^s = 0$ and $\mathbb{E}_{\tilde{v}_i^s \sim F_i^s}\left[\max\left\{v_{\mu(s)}^s, \tilde{v}_i^s\right\} - c_i^s\right] \hat{>} v_{\mu(s)}^s$.[6]

Definition 1. An outcome (μ, χ) is *stable* if the set of pairs (s, i) that block by satisfying either (i),(ii) (i.e. the classical stability condition) or (i),(ii') has measure 0.

We remark that stability of (μ, χ) depends only on each student's initial information θ^s and the preferences $v|_\chi$ revealed by the inspections χ. Simple examples show that if $\chi \neq \chi'$ are different inspections and μ is a matching, it is possible that the outcome (μ, χ) is stable but the outcome (μ, χ') is not.[7]

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

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

The budget set $B^s(\mu)$ is the set of colleges such that (s, i) satisfy (i). A stable outcome (μ, χ) must assign student s to a college $i \in B^s(\mu)$ if s so desires, and the student s cannot be assigned to any college in the complement set $\mathcal{C} \setminus B^s(\mu)$. We say that a college i is *available to* s if $i \in B^s(\mu)$, otherwise college 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*[8]

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

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

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

[6] I.e. (s, i) blocks if $\mathbb{E}_{\tilde{v}_i^s \sim F_i^s}\left[\max\left\{v_{\mu(s)}^s, \tilde{v}_i^s\right\} - c_i^s\right] > v_{\mu(s)}^s$ or $\mathbb{E}_{\tilde{v}_i^s \sim F_i^s}\left[\max\left\{v_{\mu(s)}^s, \tilde{v}_i^s\right\} - c_i^s\right] = v_{\mu(s)}^s$ and $i < \mu(s)$.

[7] For example, if there are only two colleges, 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 college, 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.

[8] Here the argmax breaks ties in favor of the college with lower index.

2.2 Regret-Free Stability

Stable outcomes ensure that given the currently acquired information no student wishes to acquire additional information about their preferences. However a stable outcome is agnostic as to how students acquired the information in the outcome, and whether this was the correct information for them to obtain. We now refine the set of stable outcomes to capture that information acquisition is costly and hence that no student wishes to overinspect.

We first consider the isolated information acquisition problem where a student $s = (\theta, v^s)$ is given a subset of colleges $C \subseteq \mathcal{C}$ to choose from, each of which guarantees her admission. Student s needs to acquire information to form her preferences and then select her assigned college from C. The student is allowed to inspect colleges sequentially, adaptively using information from previous inspections to either choose the next college to inspect or stop and select her assigned college from the set of inspected colleges. If χ^s is s's inspection indicator and $i^* \in C$ is her selected college her utility is $v^s_{i^*} - \sum_{i \in \mathcal{C}} \chi^s_i c^s_i$. The adaptive inspection strategy that maximizes the student's expected utility given the initial information (F^s, c^s) is derived by Weitzman (1979) and is stated in the following lemma.

Lemma 2. *(Weitzman 1979) Consider a student s with initial information (F^s, c^s) and values v^s who can adaptively inspect colleges and choose a college from $C \subset \mathcal{C}$. For each college i, define the index $\underline{v}^s_i$ to be the unique solution to the equation $\mathbb{E}_{\tilde{v}_i \sim F^s_i}[\max\{0, (\tilde{v}_i - \underline{v}^s_i)\}] = c^s_i$. Then an optimal inspection and stopping strategy is to sequentially inspect colleges one by one in decreasing order of their index $\underline{v}^s_i$,[9] until the score of the next college to be inspected is strictly below the maximal realized value among inspected colleges,[10] and take the best inspected college.*

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

The optimal inspection policy is an index policy, where students use the prior information F^s to compute indices $\underline{v}^s_i$ for each college i and inspect colleges in C in decreasing order of their index.[11] We note that the indices only need to be computed at the outset and do not need to be recomputed.

[9] In case of multiple colleges with equal index $\underline{v}^\theta_i = \underline{v}^\theta_{i'}$, break the tie by first inspecting the college $\min\{i, i'\}$.

[10] That is, if the set of inspected colleges is $I = \{i \mid \chi^s_i = 1\}$ then inspect $j^* = argmax_{j \in C \setminus I} \left\{\underline{v}^\theta_j\right\}$ if $\underline{v}^\theta_{j^*} \hat{>} \max_{i \in I} v^s_i$ and stop otherwise.

[11] 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.

Now consider a student s who does not know her set of available colleges C. Then for such a student s the set of inspected colleges, as given by $\chi^{opt}\left(F^s, c^s, v^s; C\right)$, depends on the set of available colleges C, the indices $\{\underline{v}_i^s\}_{i\in C}$, and the realized values $\{v_i^s\}_{i\in C}$. The following example illustrates this.

Example 1. Suppose that $\mathcal{C} = \{1, 2\}$. Let $[x; p]$ denote the probability distribution which assigns probability p to the value x and $1-p$ to 0. Consider a student with $v_1 \sim F_1 = [10; 1/2]$, $v_2 \sim F_2 = [4; 3/4]$ and $v_3 \sim F_3 = [4; 3/4]$, and let the inspection costs be $c_1 = 3, c_2 = 1$, and $c_3 = 2$. Then $\underline{v}_1 = 4$, $\underline{v}_2 = 8/3$, $\underline{v}_3 = 4/3$. If $C = \mathcal{C}$ the optimal inspection strategy is to first inspect college 1, continue to inspect college 2 only if $v_1 = 0$, and continue to inspect college 3 only if $v_2 = v_1 = 0$. If instead $C = \{2, 3\}$ the optimal inspection strategy is to inspect first college 2, and continue to inspect college 3 only if $v_2 = 0$.

In other words, knowing the set of available colleges C helps the student in Example 1 to conduct the adaptive information acquisition that maximizes her expected utility. If the student does not know the set of available colleges C, her inspection strategy may be sub-optimal in two ways. First, the student may inspect college 1 when it is not available, wasting the cost c_1. Second, the student may inspect college 2 before she inspects college 1 or learns that college 1 is not available. This is suboptimal as her optimal strategy is to inspect college 1 first, and it is possible that after inspecting college 1 she will choose to not inspect college 2.

Consider now the optimal information inspection strategies in a full matching market. When student s is part of a matching market, the set of colleges that are available to her depends on the resulting matching outcome, and therefore on the preferences of other students. As a thought experiment, suppose that in a continuum market 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 colleges available to her, which is her budget set $B^s(\mu)$ by Lemma 1. The student can optimize her information acquisition by using her initial information F^s, c^s as well as the market information $B^s(\mu)$ and applying Lemma 2. This motivates the following definition.

Definition 2. An outcome (μ, χ) is *regret-free stable* if (μ, χ) is stable and every student s

has inspected the optimal set of colleges given her available set of colleges $B^s(\mu)$, that is $\chi^s = \chi^{opt}(F^s, c^s, v^s; B^s(\mu))$ for all $s \in S$.

In other words, an outcome is regret-free stable if every student inspects colleges as if she were the last to the market. When an outcome is not regret-free stable some students would have benefited from delaying their information acquisition until the rest of the market resolved.

We make a few technical remarks about the definition of regret-free stable outcomes. First, 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 $B^s(\mu) = B^\theta(\mu)$ (i.e. the inspection decision for school s is made before observing v^s). A regret-free stable outcome could be ex post suboptimal, e.g. a student s may inspect a college $i \neq \mu(s)$ 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 information that was available from θ^s and the rest of the market. In addition, 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. Verification does not require knowledge of the students' values for uninspected colleges.

Regret-free stability formally captures students' disutility from providing rank-ordered lists of preferences at the outset of the match. For most students such a process will require them to incur unnecessary inspection costs at schools outside of their budget set, as well as at schools within their budget set that they would not have inspected under their optimal inspection process. Regret-free stability also formalizes the idea that students who do not know their budget sets could strictly improve their expected utility by waiting to know their budget sets with certainty. For example, if the student in Example 1 were guaranteed admission to schools 2 and 3, she would still wish to wait to obtain information about her admission to school 1 before deciding whether to inspect schools 2 and 3. Hence regret-free stability rationalizes the decision to hold onto multiple admission offers and wait for more information before responding to offers.

3 Information Deadlocks

We now turn to the process by which a market-maker might implement regret-free stable outcomes. In a setting with costly information acquisition, just as an outcome consists of both an allocation μ and acquired information χ, a marketplace is both a clearinghouse that decides on the final allocation μ as well as a centralized platform where information on supply and demand is aggregated and communicated. Centralized marketplaces often induce sequential information acquisition and aggregation processes where agents provide individual information to the market, the market aggregates and disseminates information to agents in the form of prices, and the process continues until it reaches a market-clearing outcome. Similarly, we may think of a college admissions market as a *sequential communication process* in which students provide information to the market, the market provides information to students, and the process continues until it arrives at an outcome (μ, χ).

In this section, we first define two natural classes of communication processes which simplify the information that needs to be provided by students. We define *canonical processes* in which students provide full information to the platform and the platform suggests actions to students, and *choice-based processes* in which students provide ordinal preferences to the market-maker. We show that standard matching market mechanisms, such as the Gale-Shapley Deferred Acceptance (DA) mechanism, are choice-based processes. We then demonstrate that under certain conditions multi-round processes, such as dynamic implementations of DA, can implement a regret-free stable outcome.

While this bodes well for the role of mechanism design in settings with costly information acquisition costs, our main finding is a negative one. We show that for general economies both canonical processes and choice-based processes are susceptible to *information deadlocks*, where no information is acquired because every agent needs more information to know how to optimally acquire information. This occurs even in the absence of strategic considerations in how students acquire and report information, and even though (as we will show in Section 4) a regret-free stable outcome always exists. Hence even though matching mechanisms which determine how students are assigned can also serve as information aggregators and communicators, when information acquisiton

is costly they are limited in their ability to facilitate optimal information acquisition.

3.1 Communication Processes

In this section we provide informal definitions for canonical processes and choice-based processes. Formal definitions can be found in Appendix A.

Definition 3. A *communication process* $\mathcal{M} = \left(f_t^{\mathcal{M}}, g^{\mathcal{M}}, \iota^0\right)$ *with oracle* ι^0 is an iterative process via which students communicate with the market and market information is provided to the students to reach an outcome. The communication process consists of a *message function* $f_t^{\mathcal{M}}$, an *aggregation function* $g^{\mathcal{M}}$, an *oracle* ι^0. The communication process also has an associated *aggregate information space* $\mathcal{I}$ that is a key component of the domain and range of the message and aggregation functions respectively.

In each iteration (indexed by t), the market considers aggregate market information and either returns an outcome, or sends messages to the students based on this aggregate market information. This mapping from aggregate market information to outcomes or messages is specified by the message function $f_t^{\mathcal{M}}(\cdot)$ on $\mathcal{I} \times \{\iota^0\}$. Based on the messages, students unilaterally perform some part of their information acquisition process and submit individual information to the market. The process aggregates this information with past aggregate information $\iota^{\mathcal{E}} \in \mathcal{I}$ via the aggregation function $g^{\mathcal{M}}$ and proceeds to the next iteration.

Our main market design question is how to design the information space $\mathcal{I}$, message function $f^{\mathcal{M}}$ and aggregation function $g^{\mathcal{M}}$ in order to implement a regret-free stable outcome. We will be interested in settings where the communication device has a simple structure; in these settings additional information will be needed to implement a regret-free stable outcome, and we will formally let this extra information be provided via the oracle ι^0.

Definition 4. A *canonical process* $\mathcal{M} = \left(f_t^{\mathcal{M}}, g^{\mathcal{M}}, \iota^0\right)$ *with oracle* ι^0 is a communication process where: (1) the aggregate information state $\mathcal{I}$ gives each student values for some of the schools, and the messages sent by the message function $f_t^{\mathcal{M}}$ are sets of schools; and (2) the aggregation function $g^{\mathcal{M}}$ aggregates students' individual information by keeping track, for each student, of the first value

for each college that is sent to the aggregator. Intuitively, canonical processes collect information from students about their observed values at each college, and suggest which both student should next inspect and the college they should inspect.

A canonical process $\mathcal{M}$ also has a *canonical outcome*, which occurs when students follow the inspections recommended by $f_i^{\mathcal{M}}$ and submit their observed values. The canonical outcome (μ, χ) is the outcome output by $\mathcal{M}$ when it terminates, and is given by setting μ to be the student-optimal stable matching with respect to the preferences induced by the observed values and priorities r^s, and χ to be the inspections suggested by the process.

Definition 5. A *choice-based process* $\mathcal{M} = \left(f_i^{\mathcal{M}}, g^{\mathcal{M}}, \iota^0\right)$ *with oracle* ι^0 is a communication process where: (1) the aggregate information state $\mathcal{I}$ gives each student a weak ordinal preference $\succeq$ over $\mathcal{C} \cup \{\phi\}$, where ϕ denotes non-inspection; and (2) the aggregation function $g^{\mathcal{M}}$ discards all previous information and records all newly submitted individual information. We say that college i is *acceptable* if and only if $i \succ \phi$. Intuitively, choice-based processes rely only on each student's current ordinal preferences.

In order to define a canonical outcome for such processes, we will also need to specify students' information acquisition processes. An *information acquisition process* σ consists of a mapping $\sigma^s : M^s \to \{0,1\}^{\mathcal{C}}$ for each student s, specifying for each message which colleges student s will inspect.

A choice-based process $\mathcal{M}$ has a *canonical outcome* under $\sigma = (\sigma^s)_{s \in \mathcal{S}}$, which occurs when students inspect according to σ and submit the partial preferences induced by their history of inspections. The canonical outcome (μ, χ) is given by running the process until all students have strict preferences $\succ^s$, and then setting μ to be the student-optimal stable matching with respect to $\succ^s$ and χ to be the inspections induced by σ.

Canonical outcomes allow us to abstract away from strategic inspection and reporting and provide explore possible outcomes when the centralized clearinghouse has full power to dictate how costly search should proceed. We will show that information deadlock arises when implementing canonical outcomes, and so the problem of inefficient search is due to the interconnected nature of the optimal search process, and not due to misaligned incentives.

Accept-Waitlist-Reject Processes

In defining choice-based mechanisms we have been agnostic as to the types of messages that can be sent to students. We now introduce a particular class of choice-based processes, the Accept-Waitlist-Reject (AWR) processes. These communication processes considerably simplify the set of possible messages sent to the students, and inform a student whether a certain college is: certainly available to her and hence *accepts* her (A); certainly unavailable to her and hence *rejects* her (R); or either available or unavailable and hence *waitlists* her (W).

Definition 6. Let $\mathcal{M}$ be a communication process. We say that $\mathcal{M}$ is an *Accept-Waitlist-Reject (AWR) process* if $\mathcal{M}$ is a choice-based process and the messages sent by the message functions $f_t^{\mathcal{M}}$ to any given student s are restricted to the set $M = \{A, W, R\}^{\mathcal{C}}$.

AWR processes provide a natural generalization to college-proposing Deferred Acceptance, and allow for information aggregation through iterative message-passing between the platform and individual students. We will show that under certain conditions, such as when students are willing to inspect any college that is available to them, college-proposing Deferred Acceptance implements a regret-free stable outcome. However, we will also show that in general vanilla AWR processes are insufficient for implementing regret-free stable outcomes, and additional market information is required for such processes to facilitate efficient information acquisition.

3.2 Deferred Acceptance and Regret-Free Stability

Consider the college-proposing DA mechanism, under which colleges propose to students with top priority, each student tentatively accepts at most one of these proposals and rejects the rest, and the process continues until all students are assigned or all colleges are full. When implemented in a multi-round manner as in Algorithm 1, DA can be interpreted as an AWR communication process that facilitates regret-free information acquisition through repeated interaction with students.

We show that under certain sufficient conditions on student preferences, the AWR process in Algorithm 1 is an iterative implementation of the college-proposing Deferred Acceptance that is regret-free stable. The key idea is that while students' information acquisition problems are

Algorithm 1 Iterative Deferred Acceptance (IDA)

procedure IDA($\mathcal{C}, \mathcal{S}, q, P$)
  while Some college still has students to propose to **do**

  Messages to Students

    Send a proposal ('accept')from each college i to the top q_i students who have not yet rejected them

  Information from Students

    Each student (irrevocably) rejects all but one of the colleges that have proposed to them, by listing them as inferior to the outside option ϕ.

  For each student s let i^s be the top-ranked college which has proposed to them and which they prefer to the outside option ϕ.
  return $\mu(s) = i^s$ for all s

interconnected and can create information deadlocks, college priorities also provide students with partial information, and this may be sufficient to both start and finish the information acquisition process.

Theorem 1. *Suppose that at least one of the following conditions hold:*

1. $\mathbb{E}[v_i^s] = \infty$ *for all* $s \in S$ *and* $i \in \mathcal{C}$ *and so students are willing inspect all colleges in their budget set.*

2. *For all* $m \geq \min_i q_i$ *the same students have the top* m *priority at all colleges, i.e. for all* $i, j \in C$, $\{s \,|\, r_i^s \geq 1 - m\} = \left\{s \,|\, r_j^s \geq 1 - m\right\}$.

Then the canonical outcome of Algorithm 1 is the college-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 colleges and reject some proposals. When students are willing to inspect all colleges available to them, any student with more than one uninspected school in their budget set is willing to inspect some colleges are reject all but one college in their budget set. When the same students have top priotiry at all colleges, there will always be at least one student who knows their full budget set and is able to perform their optimal inspection process. We provide full proofs in the appendix. □

Theorem 1 demonstrates that for certain priorities and preferences iterative Deferred Accep-

tance, which is a choice-based process, can discover sufficient preference information in a regret-free stable manner to implement a regret-free stable outcome. This process is iterative, and one can show that this is necessary: even under the conditions laid out in Theorem 1, no one-shot process — choice-based or otherwise — can always implement a regret-free stable outcome. We provide an example in Appendix C.1.

Remark. The matching found by the college-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. We provide in Appendix C.2 an example of a discrete economy where verifying that the student-optimal regret-free stable matching is the student-optimal one necessitates incurring regret with positive probability in the inspection process. 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).

3.3 Information Deadlock

Choice-based processes and AWR processes are naturally embedded in many existing matching markets. Early action in college admissions creates a two-round choice-based process; job markets often consist of multiple rounds of take-it-or-leave-it offers, which are AWR processes; and in its school-proposing form, the commonly used Deferred Acceptance mechanism is also an AWR process. However, we show that such choice-based processes can fail to find a regret-free stable outcome,

leading either to an information deadlock or to suboptimal information acquisition.

Theorem 2. *If $\mathcal{M}$ is a canonical process such that $\iota^0 = \{\underline{v}_i^s\}_{i\in\mathcal{C},s\in\mathcal{S}}$, then there exists an economy $\mathcal{E} = (\mathcal{C}, \mathcal{S}, \eta, q)$ such that the canonical outcome of $\mathcal{M}$ is not regret-free stable. Similarly, if $\mathcal{M}$ is a choice-based process such that $\iota^0 = \{\underline{v}_i^s\}_{i\in\mathcal{C},s\in\mathcal{S}}$ and σ is an information acquisition process, then there exists an economy $\mathcal{E} = (\mathcal{C}, \mathcal{S}, \eta, q)$ such that the canonical outcome of $\mathcal{M}$ under σ is not regret-free stable.*

The proof of Theorem 2 utilizes the following definition and technical lemma.

Definition 7. Let $\theta = \left(F^\theta, c^\theta, r^\theta\right)$ be a student type. The *induced demand* at college i given budget set C, denoted by $\delta_i^\theta(C)$, is the probability that a student of type θ wants college i,

$$\delta_i^\theta(C) = \mathbb{P}_{v\sim F}\left(\arg\max\left\{v_i \mid i \in C,\, \chi^{opt}\left(F^\theta, c^\theta, v; C\right) = 1\right\}\right).$$

We say that a function $\delta : 2^{\mathcal{C}} \to [0,1]^{\mathcal{C}}$ is a *Pandora demand function* if there is a student type θ with induced demand function $\delta^\theta = \delta$.

Lemma 3. *Let δ be a Pandora demand function, let $\underline{v}_1 > \underline{v}_2 > \cdots > \underline{v}_n$, and let $\varepsilon > 0$. Then there exists a student type $\theta = (F, c, r)$ with induced demand δ^θ such that $\mathbb{E}_{v\sim F}\left[(v_i - \underline{v}_i)^+\right] = c_i$ for all i and $\left|\delta_i^\theta(C) - \delta_i(C)\right| < \varepsilon$ for all $i \in \mathcal{C}, C \subseteq \mathcal{C}$.*

In other words, given any order of inspection, there exist priors and costs inducing that order of inspection which can also be used to approximately rationalize any Pandora demand function. This implies that when demand is given by following χ^{OPT} the optimal order of inspection is essentially completely uninformative about demand.

Sketch of Proof of Lemma 3. The intuition behind Lemma 3 is that we can construct distributions F and costs c that decouple the values in the support of F that determine the indices $\underline{v}_i$ from the values that determine the demand δ. In particular, if we pick m and sufficiently large M such that $M > \max_i \underline{v}_i > m$, and set the only possible values of v_i to be either exactly M or less than m, then there exist costs c_i for which the indices $\underline{v}_i$ are determined by $\mathbb{P}_{F_i}(M)$, and these probabilities

can be made arbitrarily small so that the demand δ is essentially determined by the distribution of $(v_i|v_i < m)$. We defer the full proof to the appendix. □

Sketch of Proof of Theorem 2. We specify optimal inspection processes $\{\underline{v}_i^s\}_{s\in\mathcal{S}}$ and a rank distribution ζ such that for all $\mathcal{M}$, σ there exists an economy $\mathcal{E} = (\mathcal{C}, \mathcal{S}, \eta, q)$ such that (1) $\mathcal{E}$ induces inspection processes $\{\underline{v}_i^s\}_{s\in\mathcal{S}}$ and rank distribution ζ, and (2) the canonical outcome (μ, χ) of $\mathcal{M}$ (under σ) with initial information $(\mathcal{C}, \mathcal{S}, \zeta, q)$ is not regret-free stable. We also show how to construct the economy $\mathcal{E}$ as a function of $\mathcal{M}$, σ.

The economy $\mathcal{E}$ has three colleges $\mathcal{C} = \{1, 2, 3\}$ with capacities $q_1 = q_2 = q_3 = \frac{1}{4}$. The optimal inspection processes are such that $\underline{v}_1^s > \underline{v}_2^s, \underline{v}_3^s$ for all students $s \in X$, $\underline{v}_2^s > \underline{v}_3^s, \underline{v}_1^s$ for all students $s \in Y$, and $\underline{v}_3^s > \underline{v}_1^s, \underline{v}_2^s$ for all students $s \in Z$. The rank distribution ζ is such that there are four non-intersecting sets of student types $X, Y, Z, D \subseteq \mathcal{S}$ satisfying $\eta(X) = \eta(Y) = \eta(Z) = \frac{1}{8}$ and $\eta(X \cup Y \cup Z \cup D) = 1$, and

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

$x \in X$, $y \in Y$, $z \in Z$, $d \in D$. Note that all students in X are guaranteed to have colleges 2 and 3 in their budget set, but if all students in Y and Z demand college 1 then no students in X (or in D) will have college 1 in their budget set. In addition, all students in X wish to inspect college 1 first if it is available to them, since $\underline{v}_1^s > \underline{v}_2^s, \underline{v}_3^s$. A similar observation can be made for Y and Z. (Students in D are dummy students who ensure that the number of students is greater than the number of seats.)

The intuition is that under any $\mathcal{M}, \sigma$ some students must inspect first. Let this set be denoted by S'. If some students in S' are in $X \cup Y \cup Z$ and inspect the first college i^* according to their optimal inspection process on $\mathcal{C}$, then since they have lower priority than at least $\frac{1}{4} = q_{i^*}$ students at this college there exists an economy $\mathcal{E}$ consistent with $\{\underline{v}_i^s\}_{s\in\mathcal{S}, i\in\mathcal{C}}$ and ζ where college i^* is not in their budget set and so they regret their inspection. Similarly if some students in S' are in D and

inspect some college i^*, since have lower priority than at least $\frac{3}{8} > q_{i^*}$ students at this college and there exists an economy $\mathcal{E}$ consistent with $\{\underline{v}_i^s\}_{s\in\mathcal{S},i\in\mathcal{C}}$ and ζ where college i^* is not in their budget set and so they regret their inspection. If some students in S' are in $X \cup Y \cup Z$ and inspect some college i that is not the first college i^* according to their optimal inspection process on $\mathcal{C}$, then since all the students with higher priority at i^*have at least two colleges in their budget set, there exists an economy $\mathcal{E}$ consistent with $\{\underline{v}_i^s\}_{s\in\mathcal{S},i\in\mathcal{C}}$ and ζ where college i^* is in the budget set of all students in S' and so they regret their inspection. Finally, if no students are suggested an inspection by $\mathcal{M}$ then the resulting outcome (μ, χ) is not stable, as no student is assigned and hence all students in X are willing to pay the cost to inspect colleges 2 and 3, and similarly all students in $Y \cup Z$ have a college for which they are willing to pay the cost to inspect.

We provide the complete proof in the appendix. □

The proof of Theorem 2 illustrates that in the presence of costly information acquisition, iterative processes 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. This occcurs even when agents are not strategic in how they acquire and provide information, and are fully aligned with the market maker's goal of finding a regret-free stable outcome. Hence Theorem 2 demonstrates that any choice-based process that converges to a stable matching in a regret-free manner (such as college-proposing DA) relies on an assumption that there always exist students willing to acquire some information. In settings where such an assumption does not hold, there will be limitations to how well matching mechanisms can guide the information acquisition processes.

Theorem 2 also suggests that in a college admissions system where it is costly to determine individual and aggregate preferences, even multi-stage implementations of Deferred Acceptance may result in many dissatisfied students. In particular, students may regret either some of the costly inspections they performed or the resulting preferences they submitted to the mechanism.

4 The Cutoff Structure of Regret-Free Stable Outcomes

In this section we provide several results about the structure of regret-free stable outcomes that allow us to shed light on how to implement regret-free stable outcomes. We show that regret-free stable outcomes always exist, and in fact form a non-empty lattice, and give a concise characterization of regret-free stable outcomes in terms of market-clearing cutoffs. The cutoffs provide a sufficient statistic for describing both components of a regret-free stable outcome, namely the matching and the optimal information acquisition process, and allow the interconnected information acquisition problems to be disaggregated across different students.

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

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

which is the most preferred inspected college. We show that $D^s\left(\cdot\right)$ satisfies the weak axiom of revealed preferences (WARP). In particular, we construct a full preference ordering $\succ^{\Psi(s)}$ such that for any budget set C student s will choose a school from C as if she had preferences $\succ^{\Psi(s)}$.

Proposition 1 (Reduction to demand from complete information)**.** *Let $s = (F^s, c^s, r^s, v^s)$ be a realized student. Let $\succ^{\Psi(s)}$ be an ordering of $\mathcal{C}$ defined by*

$$i \succ^{\Psi(s)} j \;\Leftrightarrow\; \min\left\{\underline{v}_i^s, v_i^s\right\} \hat{>} \min\left\{\underline{v}_j^s, v_j^s\right\}.$$

Then for all $C \subset \mathcal{C}$ we have that

$$D^s\left(C\right) = \max_{\succ^{\Psi(s)}}\left(C\right).$$

Proof. Using the indices from Lemma 2, define $i \succ^{\Psi(s)} j$ if and only if $\min\left\{\underline{v}_i^s, v_i^s\right\} \hat{>} \min\left\{\underline{v}_j^s, v_j^s\right\}$, or $\min\left\{\underline{v}_i^s, v_i^s\right\} = \min\left\{\underline{v}_j^s, v_j^s\right\}$ and $i < j$. We verify that $D^s\left(C\right) = \max_{\succ^{\Psi(s)}}\left(C\right)$.

Let $i, j \in C$ be such that $\min\left\{\underline{v}_i^s, v_i^s\right\} \hat{>} \min\left\{\underline{v}_j^s, v_j^s\right\}$. Suppose for the sake of contradiction that student s demands college j, $D^s\left(C\right) = j$. If $v_j^s \hat{<} \min\left\{\underline{v}_i^s, v_i^s\right\}$ then college i must be inspected since $\underline{v}_i^s \hat{>} v_{\max}^s$, and by assumption $v_i^s \hat{>} v_j^s$ and so since $i < j$ college i would be demanded over college j, which is a contradiction. If $\underline{v}_j^s < \min\left\{\underline{v}_i^s, v_i^s\right\}$ then college i must be inspected before college j,

since college j has lower index and is inspected. But then $v_i^s > \underline{v}_j^s$ so college j is not inspected, which gives a contradiction. In both possible cases we obtain the required contradiction. $\square$

Note that Proposition 1 implies that even though the set of colleges C in a student's budget set will affect her inspections, it does so in a manner that does not violate WARP. In particular, to determine the demanded college $D^s(C)$ it suffices to consider the relationship between the n values $\min\{\underline{v}_i^s, v_i^s\}$ without considering the effects of the set C on how information is acquired. It is worth nothing, however, that the preferences $\succ^{\Psi(s)}$ depend on the realized values v^s and hence any student s who only has initial information $\theta = \theta^s$ does not a priori know their corresponding preferences $\succ^{\Psi(s)}$. A more subtle point is that the ordering $\succ^{\Psi(s)}$ is only consistent with the demanded college $D^s(C)$; for undemanded colleges $i, j \in C \setminus \{D^s(C)\}$ it may hold that $i \succ^{\Psi(s)} j$ but given the student's information under $\chi^{opt}(F^s, c^s, v^s; C)$ student s prefers j to i.[12]

An immediate corollary of Proposition 1 is that for any set C the distribution of demand of all students of type θ from C is identical to the distribution of $\text{argmax}_{\succ^{\Psi(s)}}(C)$, where $\succ^{\Psi(s)}$ is the random preference ordering induced by drawing a random student $s = (F^s, c^s, r^s, v^s)$ from the distribution $v^s \sim F^s$.

Proposition 1 also allows us to characterize 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 college, and each student's budget set is the set of colleges where their rank meets the cutoff,

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

The demand $D^s(\boldsymbol{P})$ of student s given cutoffs $\boldsymbol{P}$ is defined in terms of her budget set $B^s(\boldsymbol{P})$ by

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

Aggregate demand for college i given cutoffs $\boldsymbol{P}$ is defined to be the mass of students that demand college i,

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

[12] For example, if $v^s_{D^s(C)} > \underline{v}_i^s > \underline{v}^s_{D^s(C)} > \underline{v}_j^s > v_i^s > v_j^s$ then under $\succ^{\Psi(s)}$ student s compares $v_i^s > v_j^s$, but under $\chi = \chi^{opt}(F^s, c^s, v^s; C)$ student s does not inspect college j and compares $v_i^s < \underline{v}_j^s$ and so given her information χ would prefer school j to school 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 effects of information acquisition are captured within the definition of $D_i(\cdot)$, as within the definition of $D^s(\boldsymbol{P})$ we require that student s acquires information according to the optimal information acquisition strategy $\chi^{opt}(F^s, c^s, v^s; B^s(\boldsymbol{P}))$.

Definition 8. A vector of cutoffs $\boldsymbol{P}$ is *market-clearing* if it matches supply and demand for all colleges 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 3. *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

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

Theorem 3 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. It follows that the regret-free stable outcomes of an economy $\mathcal{E}$ have the same attractive structural properties as the stable outcomes of a full information economy.

Theorem 4. *For every continuum economy $\mathcal{E}$ there exists a regret-free stable outcome. Moreover, the set of regret-free stable outcomes (μ, χ) is a non-empty lattice under the order $\succ$ defined by*

$$(\mu, \chi) \succ (\mu', \chi') \Leftrightarrow \mathbb{E}_{s:\theta^s=\theta}[u^s(\mu, \chi)] \geq \mathbb{E}_{s:\theta^s=\theta}[u^s(\mu', \chi')] \ \forall \theta \in \Theta,$$

where for each student type θ the expectation is taken over all students s with type $\theta^s = \theta$. In other words, the expectation is over all possible realized preferences consistent with the information initially known to a student of type θ.

Proof of Theorem 4. We show that $\mathbb{E}_{s:\theta^s=\theta}\left[u^s\left(\mu,\chi\right)\right] = \mathbb{E}_{s:\theta^s=\theta}\left[\min\left\{v^s_{\mu(s)}, \underline{v}^s_{\mu(s)}\right\}\right]$ for all $\theta \in \Theta$. The theorem then follows from the reduction shown in Proposition 1 to the complete information setting, and analogous results by Blair (1988) and Azevedo and Leshno (2016) on the lattice structure of many-to-one stable matchings in the complete information setting.

$$
\begin{aligned}
\mathbb{E}_{s:\theta^s=\theta}\left[u^s\left(\mu,\chi\right)\right] &= \mathbb{E}_{s:\theta^s=\theta}\left[v^s_{\mu(s)} - \sum_{i\in\mathcal{C}} \chi^s_i c^s_i\right] \text{ (by definition of } u^s\left(\cdot\right)\text{)} \\
&= \mathbb{E}_{s:\theta^s=\theta}\left[\max_{i\in\mathcal{C}} \chi^s_i v^s_i - \sum_{i\in\mathcal{C}} \chi^s_i c^s_i\right] \text{ (since } (\mu,\chi) \text{ is regret-free stable)} \\
&= \mathbb{E}_{s:\theta^s=\theta}\left[\max_{i\in\mathcal{C}} \chi^s_i v^s_i - \sum_{i\in\mathcal{C}} \chi^s_i \mathbb{E}\left[\left(v^s_i - \underline{v}^s_i\right)^+\right]\right] \text{ (by definition of } \underline{v}\text{)} \\
&= \mathbb{E}_{s:\theta^s=\theta}\left[\max_{i\in\mathcal{C}} \chi^s_i v^s_i - \sum_{i\in\mathcal{C}} \chi^s_i \left(v^s_i - \underline{v}^s_i\right)^+\right] \text{ (since } \chi^s_i \text{ is independent of } v^s_i\text{)} \\
&\leq \mathbb{E}_{s:\theta^s=\theta}\left[\max_{i\in\mathcal{C}} \chi^s_i \left[v^s_i - \left(v^s_i - \underline{v}^s_i\right)^+\right]\right] \\
&= \mathbb{E}_{s:\theta^s=\theta}\left[\max_{i\in\mathcal{C}} \chi^s_i \min\left\{v^s_i, \underline{v}^s_i\right\}\right] \\
&= \mathbb{E}_{s:\theta^s=\theta}\left[\min\left\{v^s_{\mu(s)}, \underline{v}^s_{\mu(s)}\right\}\right],
\end{aligned}
$$

where the last equality follows from Proposition 1, since (μ,χ) is regret-free stable and s inspects $\mu\left(s\right)$. Finally, we show that for regret-free stable outcomes if $i \neq \mu\left(s\right)$ then $\chi^s_i\left(v^s_i - \underline{v}^s_i\right)^+ = 0$, which implies that the inequality is an equality. Suppose to the contrary that $i \neq \mu\left(s\right)$, $v^s_i > \underline{v}^s_i$, and $\chi^s_i = 1$. Then student s both inspects college i and immediately stops inspecting after inspecting college i. This implies that college $\mu\left(s\right)$ was inspected before i and so $v^s_{\mu(s)} \leq \underline{v}^s_i < v^s_i$, providing the required contradiction. □

One consequence is that there is a regret-free stable matching that is ex ante optimal for all students.

Corollary 1. *For every continuum economy $\mathcal{E}$ there exists a student-optimal regret-free stable*

outcome (μ^*, χ^*)*, i.e.*

$$\mathbb{E}_{s:\theta^s=\theta}\left[u^s\left(\mu^*, \chi^*\right)\right] \geq \mathbb{E}_{s:\theta^s=\theta}\left[u^s\left(\mu, \chi\right)\right]$$

for all regret-free stable outcomes (μ, χ)*.*

Finally, we provide the following sufficient condition for there to be a unique regret-free stable outcome. The definition is in the same spirit as that found in Azevedo and Leshno (2016), and intuitively requires that the implied demand of the distribution η of students is sufficiently smooth.

Definition 9. An measure η is *regular* if $\eta\left(\{s \mid \theta^s \text{ is not regular}\}\right) = 0$ and the image under $D\left(\cdot \mid \eta\right)$ of the closure of the set $\left\{P \in (0,1)^{\mathcal{C}} \mid D\left(\cdot|\eta\right) \text{ is not continuously differentiable at } P\right\}$ has Lebesgue measure 0.

Theorem 5. *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. 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.

Remark. Theorems 4 and 5 also hold for discrete economies; moreover the proofs are exactly the same as in the continuum setting. At first glance it is not obvious why this should be the case, since in the continuum setting we assume that there is no aggregate uncertainty and in a discrete setting it is not possible to make such an assumption. Indeed some of our later results do not extend to the discrete setting, and we will make note of when this is the case. □

5 Implementing Regret-Free Stable Matchings

In this section we use the cutoff structure of regret-free stable outcomes to provide practical suggestions for market design. We first show that a market that posts market clearing cutoffs and assigns students to their favorite college for which they meet the cutoff implements a regret-free stable outcome. Hence the market-clearing cutoffs are a sufficient statistic for coordinating participants on a regret-free stable outcome. This coordination occurs even in the absence of a centralized clearinghouse.

Next we define regret-free stability in finite markets and show that in such markets estimated cutoffs can be combined with AWR processes to implement approximately regret-free stable outcomes. We provide examples of natural information oracles that can be used to estimate the market-clearing cutoffs, such as historical information and subsampled demand, and furnish cutoff mechanisms that use these oracles to implement approximately regret-free stable outcomes. The outcomes implemented by these mechanisms provably approximate the regret-free stable outcome as the accuracy of the estimation improves and as the market grows large.

Our results suggest that in practical settings with costly information acquisition, it may be useful to go beyond the standard mechanism design approach. In such settings, an equally important role of the market is facilitating the flow of information through identifying and communicating information to participants in the market. When information acquisition is costly, standard mechanisms such as Deferred Acceptance serve as costly price-discovery processes, and less costly methods for determining the market-clearing cutoffs may be beneficial for all participants in the market.

5.1 Breaking the Deadlock with Posted Cutoffs

Recall from Theorem 3 that an outcome (μ, χ) is regret-free stable if and only if there exist market-clearing cutoffs P such that (1) each student follows the Weitzman optimal inspection strategy over her budget set as described by P, and (2) each student is matched to the college 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 process that posts and commits to a collection of market-clearing acceptance cutoffs for the colleges, then allows each student to unilaterally optimize her inspection strategy and select her most-demanded college, will necessarily result in a regret-free stable match.

Consider the following single-round AWR communication process with oracle $\iota^0 = P$, where P are the market-clearing cutoffs. Intuitively, for each student it sends an acceptance message for each college i for which the student's rank at i meets the cutoff P_i.

Algorithm 2 Acceptance with Market-Clearing Cutoffs (AwMC)

procedure AwMC($\mathcal{C}, \mathcal{S}, q, P$)

Messages to Students

  for $s \in \mathcal{S}$ **do**

    for $i \in \mathcal{C}$ **do**

      if $r_i^s \geq P_i$ **then**

        Send message 'i accepts' to s

Information from Students

  for Student s in $\mathcal{S}$ **do**

    Student s submits ordinal preferences with top college i^s

  return $\mu(s) = i^s$ for all s

Theorem 6. *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 Algorithm 2 with oracle access to $\iota^0 = P$ implements the (student-optimal) regret-free stable outcome.*

Proof. It is easy to see that Algorithm 2 produces the student-optimal regret-free stable matching when all students perform the strategy $\sigma^* = \left(\chi^{OPT}, D^s\left(P\left(\eta\right)\right)\right)$ given by optimally inspecting the colleges which accept them and submit their favorite inspected college. Indeed, each student s is presented with their budget set $B^\mu\left(s\right) = \{i \in \mathcal{C} \,|\, r_i^s \geq P_i\left(\eta\right)\}$. Hence student s will be matched to their reported top choice that is in $B^\mu\left(s\right)$. Thus each student is presented with precisely the single agent problem on $B^\mu\left(s\right)$. Solving this problem yields an optimal strategy σ^* with inspection strategy $\chi^s = \chi^{OPT}$, followed by submitting preferences which list the student's favorite college: $i^s = D^s\left(P\left(\eta\right)\right)$ at the top. The assumption that P is (student-optimal) market-clearing guarantees that the resulting outcome (μ, χ) is (student-optimal) regret-free stable. □

Theorem 6 states that advance knowledge of market-clearing cutoffs are sufficient for regret-free stability. Posting cutoffs in advance of any information acquisition removes all uncertainty on the part of the agents about which colleges 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 college was ultimately unattainable. We note that this lack of regret does not depend on the posted cutoffs being market-clearing, but only that the market-maker 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 Algorithm 2 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 Stability in Finite Economies

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, AWR processes that learn and post estimated cutoffs implement outcomes that are regret-free stable with respect to slightly perturbed college capacities. This suggests that centralized processes like Deferred Acceptance and decentralized procedures based on posted cutoffs may perform well even when information acquisition is costly, as long as there is sufficient initial information about aggregate preferences.

5.2.1 Finite Model

In this section, we consider a model with a finite set of students. The main departure from the continuum setting is that there is aggregate uncertainty; students are atomic, so uncertainty about individual preferences leads to uncertainty about aggregate demand.

A *discrete economy* is given by $E = (\mathcal{C}, S, q)$, where S is a finite set of students $s = (F^s, c^s, r^s, v^s)$, r_i^s is the percentile rank of student s at college i, and $q = \{q_i\}_{i \in \mathcal{C}} \in \mathbb{N}^{\mathcal{C}}$ is the vector of quotas at each college. We may think of E as inducing a continuum economy $\mathcal{E} = \left(\mathcal{C}, \mathcal{S}, \eta, \frac{q}{|S|}\right)$ with $|S|$ atoms, where for each subset of students $S' \subseteq S$ we let $\eta(S') = \frac{|S'|}{|S|}$ and $\eta(\mathcal{S} \setminus S) = 0$; however this continuum economy has aggregate uncertainty and outcomes will depend on the specific realization of student values v^s.

Stability and regret-free stability are defined as in the continuum model, replacing the measure η as necessary with the proportion of involved students. Specifically, given a matching μ, we define the budget set of student s by $B^s(\mu) = \left\{i \in \mathcal{C} \mid r_i^s \geq \min\left\{r_i^{s'} \mid s' \in \mu(i)\right\}\right\} \cup \{i \in \mathcal{C} \mid |\mu(i)| < q_i\}$ .We say that (s, i) forms a blocking pair if (i) $i \in B^s(\mu)$ and either (ii) y $\chi_i^s = 1$ and $v_i^s \hat{>} v_{\mu(s)}^s$, or (ii')

$\chi_i^s = 0$ and $\mathbb{E}_{\tilde{v}_i^s \sim F_i^s}\left[\max\left\{v_{\mu(s)}^s, \tilde{v}_i^s\right\} - c_i^s\right] \hat{>} v_{\mu(s)}^s$.

Definition 10. An outcome (μ, χ) is *stable* if there are no pairs (s, i) that block by satisfying either (i),(ii) or (i),(ii'). A stable outcome (μ, χ) is *regret-free stable* if $\chi^s = \chi^{opt}\left(F^s, c^s, v^s; B^s(\mu)\right)$ for all $s \in S$.

These definitions directly derive from Definitions 1 and 2 in the continuum setting. However, their interpretation changes now that students are atomic. In particular, the intuition that the budget set $B^s(\mu)$ is the set of colleges that student s could be admitted to when arriving 'last to market' is no longer correct, as the student's preferences could affect the set of colleges in their budget set.[13] In addition under regret-free stable outcomes students are not necessarily performing optimal search given their knowledge of all other students' preferences, as such outcomes may require students to inspect colleges with which they cannot credibly block.[14] This mirrors the issue with stability identified in Ehlers and Morrill (2017), whereby a student in a blocking pair (s, i) may be unjustifiably envious of another student assigned to i, as there are no possible outcomes where s is assigned to i. When information acquisition is costly this may also lead to unnecessary costly inspection.

Remark. The presence of aggregate uncertainty presents additional difficulties in implementing regret-free stable outcomes without allowing students to interleave their information acquisition processes, and also in verifying that a given regret-free stable outcome is student-optimal. We illustrate these difficulties in Appendix C.

[13] Consider the following modification of Example X in [Erdil and Ergin Y]. $E = (\mathcal{C} = \{1, 2, 3\}, S = \{x, y, z\}, q = \mathbf{1})$, where college priorities satisfy $r_1^y > r_1^z > r_1^x$ and $r_2^x > r_2^y > r_2^z$, student x prefers $1 \succ 2 \succ 3$ and has cost 0 at each college, student y prefers $2 \succ 1 \succ 3$ and has cost 0 at each college, and student z has information acquisition problem (F^z, c^z, r^z) such that $\underline{v}_1^z > v_2^z > v_3^z$. There is a unique matching that corresponds to any stable outcome of $E \setminus z$, given by $\mu(x, y) = (1, 2)$, so z's budget set arriving 'last to market' is $\{1, 3\}$. If, given a budget set of $\{1, 3\}$, z demands 3, then she is assigned to 3 and her budget set remains $\{1, 3\}$. However, if z demands 1, this changes the unique stable matching in any stable outcome to be $\mu(x, y, z) = (2, 1, 3)$.

[14] For example, in the example in footnote 13, since $\underline{v}_1^z > \underline{v}_3^z$ if v^z is such that z demands 3 then in the unique regret-free stable outcome z must inspect college 1. However it is easy to see that z cannot be assigned to 1 in any stable outcome, and so z cannot credibly block with 1 and hence can improve their expected utility by not inspecting college 1.

5.2.2 Approximate Cutoffs and Approximate Regret-Free Stability

In this section, we show that estimated market-clearing cutoffs implement outcomes that are regret-free stable with respect to perturbed capacities, and provide convergence results for large markets. The intuition behind these results is that by Proposition 1 the costly information acquisition setting reduces to the full information setting, and so all questions about cutoff-based processes under costly information acquisition reduce to analogous questions about cutoff-based processes in markets without costly information acquisition.

We first show that the outcome $\left(\mu^P, \chi^P\right)$ from posting cutoffs P when using Algorithm 2 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 Algorithm 2 on $\mathcal{E}$ with cutoffs P, and for all i let $q_i^P = \left|\left\{s \,|\, \mu^P(s) = i\right\}\right|$. 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 $\left|\left\{s \,|\, \mu(s) \neq \mu^P(s)\right\}\right| \leq \|P - P^*\|_2$.*

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

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

Proposition 2 (Distribution of approximately feasible capacities)**.** *Suppose the continuum economy $\mathcal{E}$ admits a unique stable matching μ with cutoffs P^*, and let $\overline{q} = D(P^*|\eta)$.*[15] *Let $E^k = \left(\mathcal{C}, S^k, q^k k\right)$ be a randomly drawn finite economy, where S^k is a set of k students drawn independently according to η, and where $q^k = \frac{1}{k} D\left(P^*|S^k\right)$ 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 $D\left(P^*|S^k\right) = \sum_{a=1}^{k} X_a$, 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 Appendix D.1 that the mapping from demand to market-clearing cutoffs is continuous, so 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.

5.3 Regret-Free Learning

In this section, we provide illustrative examples of how to estimate population demand and use the estimated demand to implement approximately regret-free stable outcomes. We explore two simple sources of information a market designer could use to estimate market demand in college admissions. The first is historical demand. 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. The second is through subsampling student demand. Subsampled student demand can be used to estimate or provide bounds on the market-clearing cutoffs and bootstrap the information acquisition process. For example, if running iterative Deferred Acceptance assigns some students before reaching a deadlock, the demand of the assigned students

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

could be used to estimate market-clearing cutoffs.

We also exploit the fact that 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 1, the alignment of college priorities in 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 information communication process.

Definition 12. Let $\mathcal{E} = (\mathcal{C}, S, \eta, q)$ be an economy. A student s has **free market information** if for all colleges i student s is either in the top q_i students or the bottom $|S| - \sum_j q_j$ percentile of students. Formally, we let $S^f(\mathcal{E}) = \left\{ s \,|\, \forall i \, r_i^s \notin \left[1 - \frac{1}{|S|} \sum_j q_j, 1 - \frac{1}{|S|} q_i \right) \right\}$ denote the set of students with free market information in $\mathcal{E}$.

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

In the following examples, we estimate population demand using historical information and subsampling of demand from students with free market information, and illustrate how to use cutoff-based processes to implement approximately regret-free stable outcomes. Moreover we also quantify the variance in the required violations of capacity. Our results suggest that the outcomes are reasonably robust to small biases or noise due to sampling error.

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 assume that $\mathcal{E} = (\mathcal{C}, \mathcal{S}, \eta, q)$ is an economy with a unique regret-free stable outcome and let P^* be the market-clearing cutoffs

and $\overline{q} = D\left(P^*|\eta\right)$ then we can use classic results about the convergence of two-step estimators[16] 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 11. The full proof can be found in Appendix D.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-based process can perform arbitrarily poorly, whereas more accurate historical information ($\alpha \to \infty$) leads to smaller perturbations in the capacities.

Example 3. 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 $\eta\left(\Gamma^*\right)$, where student demand $D^s\left(P|\eta\left(\Gamma\right)\right)$ is parametrized by an n-dimensional parameter $\Gamma = (\gamma_1, \gamma_2, \ldots, \gamma_n)$ (recall that $n = |\mathcal{C}|$ is the number of colleges). Let α denote the fraction of students with free market information. (Note that if we define $\alpha^* = \min_i q_i + 1 - \sum_i q_i$ then for all economies $\mathcal{E}$ it holds that $\alpha \leq \alpha^*$, and $\alpha = \alpha^*$ is achieved when colleges have aligned preferences, i.e. condition (2) in Theorem 1 holds.) If $\alpha > 0$, 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 two-round AWR process $\mathcal{M}^F$ defined as follows. In the first round it reveals all acceptances to all students $s \in S^f\left(E\right)$, assigns them each to their chosen college and obtains their aggregate demand $\hat{q}_i^f k$ for each college i. It then uses this demand to provide an estimate $\hat{\Gamma}$ for Γ^*. In the second round $\mathcal{M}^F$ runs Algorithm 2 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 $k - \left|S^f\left(E\right)\right|$ students without free market information in E independently from the distribution $\eta\left(\hat{\Gamma}\right)$ and the residual capacity for each college i is $\hat{q}_i^r = q_i - \hat{q}_i^f$. The residual economy $E^r = \left(\mathcal{C}, S^k \setminus S^f\left(E\right), \hat{q}^r k\right)$ is given by removing the students in $S^f\left(E\right)$ from E and reducing capacities accordingly at their assigned colleges.

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

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

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

as $k \to \infty$ for $A = \nabla_\Gamma D(P^*(\Gamma)|\eta(\Gamma^*))|_{\Gamma=\Gamma^*}\left(\nabla_\Gamma D^f(\Gamma^*)\right)^{-1}$ and $\Sigma^{\overline{q}}, \Sigma^{q^f}$ defined as in Definition 11. 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 results about the distribution of two-step estimators. The full proof can be found in Appendix D.2.

The matrix $A = \nabla_\Gamma D(P^*(\Gamma)|\eta(\Gamma^*))|_{\Gamma=\Gamma^*}\left(\nabla_\Gamma D^f(\Gamma^*)\right)^{-1}$ consists of two multiplicative components: the sensitivity of the true demand $D(\cdot|\eta(\Gamma^*))$ to errors in the cutoffs $P^*(\Gamma)$ from misspecifying Γ, and the inverse sensitivity of the demand $D^f(\cdot)$ of students with free market information. 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 and A. In the absence of free market information ($D^f(\Gamma^*) \to 0$) the cutoff-based process can perform arbitrarily poorly. Suppose, however, that priorities are such that more students have free market information ($\alpha \to \alpha^*$), or that demand is insensitive to misspecifications in cutoffs ($\nabla_\Gamma D \to 0$) while the demand of students with free information provides an accurate estimate of Γ^* ($\nabla_\Gamma D^f \to \infty$). Then there are smaller perturbations in the capacities. Finally, the first round in this process corresponds to the first round of iterative college-proposing Deferred Acceptance. If we allow for further rounds of proposals, we can further reduce the variance of the perturbed capacities.

6 Discussion

We have provided a model for matching with costly information acquisition, and proposed regret-free stability as a suitable solution concept that rationalizes agents' decisions to delay preference formation or withhold decisions about preferences. We have also identified the phenomenon of *information deadlocks* in markets with costly information acquisition, and demonstrated the insufficiency of mechanism design for overcoming these deadlocks. Hence, in contrast with the predictions of stability and student-optimality found in models without costly information acquisition, our results suggest that AWR processes like Deferred Acceptance may cause students to incur significant regret when they do not have access to *oracles* which provide accurate information about aggregate demand.

To break the information deadlock, we have shown that *market-clearing cutoffs* form a sufficient statistic for characterization regret-free stable outcomes. We have demonstrated how oracles for these cutoffs can be constructed using information about historical demand or by subsampling demand in an existing economy. Our findings suggest that in stable matching markets with costly information acquisition, recurring markets where aggregate information does not fluctuate much over time should perform better than new marketplaces or in marketplaces where aggregate demand varies widely. Moreover given sufficient information decentralized markets can be just as efficient as centralized clearinghouses. Hence our results underscore the importance of analyzing the role of markets in facilitating information acquisition and aggregation, and also help explain why long-running markets using Deferred Acceptance perform well despite informational frictions.

We conclude with some directions for future work.

Other models of incomplete information. We have considered a setting with one-sided incomplete information where students know their priorities exactly. We have also assumed that a student must inspect a college in order to be assigned to it. Both assumptions were required for analytical tractability, as the optimal information acquisition process is not known even for a single agent when students have non-trivial priors over their admissions probabilities, or when students can be assigned to a college and obtain the expected value without inspection. Shorrer (2019); Doval (2018); Beyhaghi and Kleinberg (2019) provide additional insight to the single-agent

setting, and an investigation of the impact of matching constraints on these problems is an exciting direction for future work.

Approximation algorithms. Our results demonstrate that in general there is a tradeoff between the regret of a communication process, the feasibility of the solution, and the amount of information provided to the market-maker. We have provided one class of algorithms that relax the feasibility constraint in order to achieve optimal regret. It may also be possible to relax the regret of the algorithm 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 processes 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 processes 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 processes 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 additional implications do our results have for practical applications? Colleges in many countries, such as China, India, Ireland and Australia post historical cutoffs for admission into college programs. Our results on processes 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 processes 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.

H. Beyhaghi and R. Kleinberg. Pandora's problem with nonobligatory inspection. *arXiv preprint arXiv:1905.01428*, 2019.

S. Bikhchandani. Stability with one-sided incomplete information. *Journal of Economic Theory*, 168:372–399, 2017.

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.

H. Chade, G. Lewis, and L. Smith. Student portfolios and the college admissions problem. *Review of Economic Studies*, 81 (3):971–1002, 2014.

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.

Y.-C. Chen and G. Hu. Learning by matching. *Theoretical Economics*, 2019.

S. P. Corcoran, J. L. Jennings, S. R. Cohodes, and C. Sattin-Bajaj. Leveling the playing field for high school choice: Results from a field experiment of informational interventions. Technical report, National Bureau of Economic Research, 2018.

L. Doval. Whether or not to open pandora's box. *Journal of Economic Theory*, 175:127–158, 2018.

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.

S. Dynarski, C. Libassi, K. Michelmore, and S. Owen. Closing the gap: The effect of a targeted, tuition-free promise on college choices of high-achieving, low-income students. Technical report, National Bureau of Economic Research, 2018.

F. Echenique, A. J. Wilson, and L. Yariv. Clearinghouses for two-sided matching: An experimental study. *Quantitative Economics*, 7(2):449–482, 2016.

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

L. Ehlers and T. Morrill. (il) legal assignments in school choice. 2017.

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.

J. Grenet, Y. He, and D. Kübler. Decentralizing centralized matching markets: Implications from early offers in university admissions. 2019.

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.

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.

F. Klijn, J. Pais, and M. Vorsatz. Preference intensities and risk aversion in school choice: A laboratory experiment. *Experimental Economics*, 16(1):1–22, 2013.

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

Q. Liu. Rational expectations, stable beliefs, and stable matching. Technical report, Working Paper, 2018.

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 P. R. Milgrom. *Putting auction theory to work*. Cambridge University Press, 2004.

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.

R. I. Shorrer. Simultaneous search: Beyond independent successes. In *Proceedings of the 2019 ACM Conference on Economics and Computation*, pages 347–348. ACM, 2019.

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 Communication Processes

A *communication process* is an iterative process via which students communicate with the market and market information is provided to the students to reach an outcome.

Definition 13. A *communication process* $\mathcal{M} = \left(f_t^{\mathcal{M}}, g^{\mathcal{M}}, \iota^0, T\right)$ with *communication device* $d = (\mathcal{I}, I, M)$ and oracle ι^0 is defined via the mappings

$$f_t^{\mathcal{M}} : \mathcal{I} \times \left\{\iota^0\right\} \to M \cup \left(\mathcal{S}^{\mathcal{C}}, \chi\right) \ \forall t \leq T,$$
$$g^{\mathcal{M}} : \mathcal{I} \times \left\{\iota^0\right\} \times I \to \mathcal{I},$$

where $f_t^{\mathcal{M}}$ is the *message function* at time t, and $g^{\mathcal{M}}$ is the *aggregation function.*

The process has an associated *aggregate information space* $\mathcal{I}$, *message space* M, *individual information space* I and *number of rounds* T. The mapping $f^{\mathcal{M}}$ takes as input the current information state $\iota^{\mathcal{E}}$ and returns either a message $m = \{m^s\}_{s\in\mathcal{S}}$ to send to the students, or an outcome (μ, χ) where χ is the current inspection indicator. The mapping $g^{\mathcal{M}}$ aggregates students' individual information $\iota = \{\iota^s\}_{s\in\mathcal{S}} \in I$ with the past information state $\iota^{\mathcal{E}} \in \mathcal{I}$ to produce a new information state $g^{\mathcal{M}}\left(\iota^{\mathcal{E}}, \iota^0, \iota\right) \in \mathcal{I}$.

If the number of rounds is finite, i.e. $T < \infty$, we require that the process always outputs an outcome by the last round, $f_T^{\mathcal{M}}\left(\iota^{\mathcal{E}}, \iota^0\right) \in \left(\mathcal{S}^{\mathcal{C}}, \chi\right)$ for all $\iota^{\mathcal{E}} \in \mathcal{I}$. If T is not finite, to ensure termination of the process, we require that iterated applications of the process with any initial information state $\iota^{\mathcal{E}} \in \mathcal{I}$ and individual information $\iota \in I$ will ultimately produce an outcome (μ, χ). [17]

[17] More formally, fix any mapping $\sigma : M \to I$; this mapping will represent what individual information students acquire as a function of the messages they are given. We assume that the information state $\mathcal{I}$ is endowed with a metric. Define a mapping $h : \mathcal{I} \to \mathcal{I}$ by $h\left(\iota^{\mathcal{E}}\right) = g^{\mathcal{M}}\left(\iota^{\mathcal{E}}, \iota^0, \left(\sigma \circ f^{\mathcal{M}}\right)\left(\iota^{\mathcal{E}}, \iota^0\right)\right)$ if $f^{\mathcal{M}}\left(\iota^{\mathcal{E}}\right) \in M$ and $h\left(\iota^{\mathcal{E}}\right) = \iota^{\mathcal{E}}$ otherwise. The requirement is twofold. First, we require that for all mappings σ and all initial information states $\iota^{\mathcal{E}} \in \mathcal{I}$ iterated applications of h starting with $\iota^{\mathcal{E}}$ will converge to some $\iota_1^{\mathcal{E}} \in \mathcal{I}$. In addition, we require if for each $k \geq 1$ we let $\iota_{k+1}^{\mathcal{E}} \in \mathcal{I}$ denote the limit of the process starting with $\iota_k^{\mathcal{E}}$, and for each ordinal ω' let the union of all the information from the $\iota_k^{\mathcal{E}}$ for smaller ordinals $k < \omega'$ be some $\iota_{\omega'}^{\mathcal{E}} \in \mathcal{I}$, then there is some ordinal ω^n for which $\iota_{\omega^n}^{\mathcal{E}}$ is a fixed point of h and produces an outcome $f^{\mathcal{M}}\left(\iota_{\omega^n}^{\mathcal{E}}\right) \in \mathcal{S}^{\mathcal{C}} \times \chi$.

A.1 Canonical Processes

Definition 14. Let $\mathcal{M}$ be a communication process with oracle ι^0. We say that $\mathcal{M}$ is *canonical* if: (1) the communication device $d = (\mathcal{I}, I, M)$ is given by $\mathcal{I} = (\zeta, q) \times \{\mathbb{R} \cup \{\phi\}\}^{\mathcal{C}\times\mathcal{S}}$, $I^s = \{\mathbb{R} \cup \{\phi\}\}^{\mathcal{C}}$ and $M^s = \{0,1\}^{\mathcal{C}}$; and (2) the mapping $g^{\mathcal{M}}$ is defined by $g^{\mathcal{M}}\left(\iota^{\mathcal{E}}, \iota^0, \iota\right) = \tilde{\iota}^{\mathcal{E}} \in \mathcal{I}$ satisfying $\left(\tilde{\iota}^{\mathcal{E}}\right)_i^s = \left(\iota^{\mathcal{E}}\right)_i^s$ if $\left(\iota^{\mathcal{E}}\right)_i^s \neq \phi$, and $\left(\tilde{\iota}^{\mathcal{E}}\right)_i^s = \iota_i^s$ otherwise. (Here we have used the notation $\iota^{\mathcal{E}} = (\zeta, q) \times \left(\left(\iota^{\mathcal{E}}\right)_i^s\right)_{i\in\mathcal{C},\, s\in\mathcal{S}} \in \mathcal{I}$, $\iota = (\iota_i^s)_{i\in\mathcal{C},\, s\in\mathcal{S}} \in \prod_{s\in\mathcal{S}} I^s$ and $m = (m_i^s)_{i\in\mathcal{C},\, f\in\mathcal{S}} \in \prod_{s\in\mathcal{S}} M^s$.)

The first condition can be thought of as setting the information space to be student ranks and college capacities augmented by a value for each student at each college, and setting the message and individual information spaces so that students can be suggested sets of colleges to inspect and the individual information of each student is their values for the colleges. The second condition can be thought of as aggregating students' individual information by keeping track of the first value for each college that is sent to the aggregator.

Definition 15. Let $\mathcal{M}$ be a canonical process. The *canonical outcome* (μ, o) of $\mathcal{M}$ is defined as follows. Students follow the actions suggested by $\mathcal{M}$ and update their individual information based on their observed values, i.e. $(\iota^s)_i = v_i^s$ if $(m^s)_i = 1$ and $(\iota^s)_i = \phi$ otherwise. Each student's information history is aggregated to give a preference ordering $\succ^s$ where colleges are acceptable if and only if they are inspected, and for inspected colleges i, j the student prefers the one with higher value. The outcome (μ, o) is given by setting μ to be the student-optimal stable matching with respect to the induced preferences $\succ^s$ and priorities r^s, and χ to be the inspections induced by the actions suggested by the messages m.

A.2 Choice-Based Processes

Let $\mathcal{L}(\mathcal{C} \cup \{\phi\})$ denote the set of all weak ordinal preferences over $\mathcal{C} \cup \{\phi\}$, where ϕ denotes non-inspection. Recall that college i is *acceptable* if and only if $i \succ \phi$.

Definition 16. Let $\mathcal{M}$ be a communication process with oracle ι^0 . We say that $\mathcal{M}$ is *choice-based* if: (1) the communication device $d = (\mathcal{I}, I, M)$ is given by $\mathcal{I} = (\zeta, q) \times \mathcal{L}(\mathcal{C} \cup \{\phi\})$ and

$I^s = \mathcal{L}(\mathcal{C} \cup \{\phi\})$; and (2) the mapping $g^{\mathcal{M}}$ is given by $g^{\mathcal{M}}\left(\iota^{\mathcal{E}}, \iota^0, \iota\right) = \iota$, where we write $\iota^{\mathcal{E}} = (\zeta, q) \times \left(\left(\iota^{\mathcal{E}}\right)_i^s\right)_{i \in \mathcal{C},\, s \in \mathcal{S}} \in \mathcal{I}$, $\iota = (\iota_i^s)_{i \in \mathcal{C},\, s \in \mathcal{S}} \in \prod_{s \in \mathcal{S}} I^s$ and $m = (m_i^s)_{i \in \mathcal{C},\, s \in \mathcal{S}} \in \prod_{s \in \mathcal{S}} M^s$.

In other words, the communication process $\mathcal{M}$ is choice-based if communication from students and aggregate market information is based on student choices. The first condition restricts the aggregate and individual message spaces to be weak ordinal preferences, and the second condition says that indvidual information is re-reported every round.

Definition 17. An *information acquisition process* σ consists of a mapping $\sigma^s : M^s \to \{0,1\}^{\mathcal{C}}$ for each student s, specifying for each message which colleges student s will inspect.

Definition 18. Let $\mathcal{M}$ be a choice-based process, and let σ be an information acquisition process. The *canonical outcome* (μ, o) of $\mathcal{M}$ under $\sigma = (\sigma^s)_{s \in \mathcal{S}}$ is defined as follows. Students' individual information is given by the partial preferences $\succeq^s$ induced by inspecting according to σ^s, i.e. colleges are acceptable if and only if they are inspected, and for inspected colleges i, j the student prefers the one with higher value. The outcome (μ, o) is given by running the process until all students have strict preferences $\succ^s$, and then setting μ to be the student-optimal stable matching with respect to $\succ^s$ and χ to be the inspections induced by σ.

B Proofs for Section 3.3

Proof of Lemma 3. We construct priors F and costs c that induce demand close to δ and such that $\mathbb{E}\left[(v_i - \underline{v}_i)^+\right] = c_i$ for all i. Let M and m be reals such that $M > \max_i \underline{v}_i$ and $0 < m < \min_i \underline{v}_i$, and let $\rho > 0$ be sufficiently small so that $\rho + 2\left(1 - (1-\rho)^n\right) < \varepsilon$. Let $\tilde{v} \sim \tilde{F}$ and $\tilde{c}$ be random variables, priors and costs such that δ is induced by $\tilde{F}$ and $\tilde{c}$. Without loss of generality, we may assume that $\tilde{F}$ and $\tilde{c}$ are scaled and shifted so that $\tilde{F}_i(0) = 0$ and $\tilde{F}_i(m) \geq 1 - \rho$. For each i let $\underline{\tilde{v}}_i$ satisfy $\mathbb{E}\left[(\tilde{v}_i - \underline{\tilde{v}}_i)^+\right] = \tilde{c}_i$, and let $\tilde{\kappa}_i = \min\{\tilde{v}_i, \underline{\tilde{v}}_i\}$. For each $C \subseteq \mathcal{C}$ let $\tilde{\kappa}_C = \max_{i \in C} \tilde{\kappa}_i$. It follows that $\delta_i(C) = \mathbb{P}\left(i = \min\left\{j : \tilde{\kappa}_j = \tilde{\kappa}_C\right\}\right)$ (where the choice of the minimum follows our tie-breaking convention). Let $\hat{F}_i$ be the cumulative distribution function of κ_i. It follows that $\hat{F}_i(0) = 0$, $\hat{F}_i(m) \geq 1 - \rho$.

For each i, we define F_i and c_i as follows. Define

$$F_i(x) = \begin{cases} \frac{1-\rho}{\hat{F}_i(m)} \hat{F}_i(x) & \text{for all } x \in [0, m] \\ 1-\rho & \text{for all } x \in [m, M), \qquad c_i = \rho(M - \underline{v}_i) > 0 \\ 1 & \text{for } x \geq M, \end{cases}$$

so that v_i is M with probability ρ and less than m otherwise. Then $\mathbb{E}\left[(v_i - \underline{v}_i)^+\right] = \rho(M - \underline{v}_i) = c_i$ as required. Morever, it follows that $\mathbb{P}\left(v_i > \underline{v}_j \, \forall j\right) = \mathbb{P}(v_i = M) = \rho$, and that $\mathbb{P}\left(v_i < \underline{v}_j \, \forall j\right) = \mathbb{P}(0 \leq v_i \leq m) = 1 - \rho$.

Let $\hat{\delta}$ be induced by F, c. We show that for all $C \subseteq \mathcal{C}$ and $i \in C$ it holds that $\left|\hat{\delta}_i(C) - \delta_i(C)\right| < \varepsilon$, from which the theorem follows. Now

$$\begin{aligned} \hat{\delta}_i(C) - \delta_i(C) &= \mathbb{P}\left(\min(v_i, \underline{v}_i) \geq \min(v_k, \underline{v}_k) \ \forall k \in C\right) - \mathbb{P}\left(\tilde{\kappa}_i \geq \tilde{\kappa}_C\right) \\ &= \rho(1-\rho)^{|C\setminus\{i\}|} + (1-\rho)^{|C|} \mathbb{P}\left(\tilde{\kappa}_i \geq \tilde{\kappa}_C \,|\, \tilde{\kappa}_C \leq m\right) - \mathbb{P}\left(\tilde{\kappa}_i \geq \tilde{\kappa}_C\right) \end{aligned}$$

where

$$\begin{aligned} &\left|\rho(1-\rho)^{|C\setminus\{i\}|} + (1-\rho)^{|C|} \mathbb{P}\left(\tilde{\kappa}_i \geq \tilde{\kappa}_C \,|\, \tilde{\kappa}_C \leq m\right)\right| \\ =& \left|\mathbb{P}\left(\tilde{\kappa}_i \geq \tilde{\kappa}_C\right)\left(\frac{(1-\rho)^{|C|}}{\mathbb{P}\left(\tilde{\kappa}_C \leq m\right)} - 1\right) - (1-\rho)^{|C|} \frac{\mathbb{P}\left(\tilde{\kappa}_i \geq \tilde{\kappa}_C \,|\, \tilde{\kappa}_C \leq m\right)\mathbb{P}\left(\tilde{\kappa}_C > m\right)}{\mathbb{P}\left(\tilde{\kappa}_C \leq m\right)}\right| \\ \leq& \left|\frac{(1-\rho)^{|C|}}{\mathbb{P}\left(\tilde{\kappa}_C \leq m\right)} - 1\right| + \left|(1-\rho)^{|C|} \frac{\mathbb{P}\left(\tilde{\kappa}_C > m\right)}{\mathbb{P}\left(\tilde{\kappa}_C \leq m\right)}\right| \\ \leq& \left(1 - (1-\rho)^{|C|}\right) + \left(1 - (1-\rho)^{|C|}\right). \end{aligned}$$

Hence

$$\left|\hat{\delta}_i(C) - \delta_i(C)\right| \leq \rho + 2\left(1 - (1-\rho)^n\right) < \varepsilon$$

as required. □

Proof of Theorem 1. We remark that in both cases the result holds even for finite economies, as defined in Section 5.2.

Case (1). Suppose that all students inspect all colleges in their budget set, i.e. condition (1) holds. First, assuming truthful reporting, it is regret-free for every student to inspect all colleges that proposed to them. This is because if college 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 regret-free stable outcomes (μ, χ) it holds that $i \in B^s(\mu)$. Since s is willing to inspect any college in her budget set it follows that it is optimal for student s to inspect college i.

Next, it is regret-free for each student to reject all the colleges that proposed to them except the one (up to tiebreaking by indices) with the highest observed value. Moreover, monotonicity of the number of rejections implies that $\mathcal{M}$ converges. The algorithm terminates at a fixed point, which implies that either all colleges are at capacity or all students are assigned, 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 colleges, 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 colleges, and for all $m > \underline{q}$ if a student has rank $1 - m$ at one college then their rank at all colleges is $1 - m$. Intuitively Algorithm 1 performs serial dictatorship and tells each student her (unique) budget set in order of her rank, which results in a unique regret-free stable outcome. When students don't have indifferences in their preferences, it follows that at each step some students reject all the colleges that proposed to them except one, and so again by a monotonicity argument the algorithm terminates, and when it terminates it outputs a regret-free stable outcome.

Let us show that there is a unique regret-free stable outcome. Let $(\mu, \chi), (\mu', \chi')$ be regret-free stable outcomes. Let s be an arbitrary student, and note that either s has rank $\geq 1 - \underline{q}$ at all colleges, or there exists $m > \underline{q}$ such that s has rank $1 - m$ at all colleges. It suffices to show that $B^s(\mu) = B^s(\mu')$ for all s.

Suppose that s has rank $\geq 1 - \underline{q}$ at all colleges. Then $B^s(\mu) = B^s(\mu') = \mathcal{C}$.

Let $X = \{s : B^s(\mu) \neq B^s(\mu')\}$. Note that since students with rank $\geq 1 - \underline{q}$ at any college have the same budget sets under μ and μ', students in X are ordered by rank, i.e. for all $s, s' \in X$ and $i, j \in \mathcal{C}$ if $r_i^s \geq r_i^{s'}$ then $r_j^s \geq r_j^{s'}$.

If the set of students were finite, it would follow by induction that for any rank $\tilde{r} < 1 - \underline{q}$ the student with rank $r_i^s = \tilde{r}$ has the same budget set under μ and μ' and so $X = \emptyset$. Since we have a continuum of students, we have to select

Let $r^* = \inf \{\tilde{r} : \eta (X \cap \{s : r_i^s \leq \tilde{r}\}) = 0\}$, and for each college i let

$$\tilde{q}_i = q_i - \eta \left(\left\{ s' : r_1^{s'} \geq r^*,\, D^{s'} \left(B^s (\mu) \right) = i \right\} \right)$$

be the residual capacity at college i after all students with rank better than r^* have selected their college. Note that the $\tilde{q}_i$ would remain the same if we defined them instead in terms of μ'.

Let $\tilde{q}_{min} = \min \{\tilde{q}_i : \tilde{q}_i > 0\}$, and let $s \in X$ be such that $r_i^s \in [r^* - \tilde{q}_{min}, r^*]$. By the definition of r^* the set of such s has postive η-measure. But such a student s knows her budget set exactly: under both μ and μ' it is given by the set of colleges $\{i : \tilde{q}_i > 0\}$ that have positive residual capacity after all students with rank better than r^* have selected their college, since there are not enough unassigned students with rank in $[r_i^s, r^*]$ to exhaust the residual capacity of these colleges. Hence $B^s (\mu) = B^s (\mu')$ and $s \notin X$, which is a contradiction. □

Proof of Theorem 2. We specify optimal inspection processes $\{\underline{v}_i^s\}_{s \in \mathcal{S}}$ and a rank distribution ζ such that for all $\mathcal{M}$, σ there exists an economy $\mathcal{E} = (\mathcal{C}, \mathcal{S}, \eta, q)$ such that (1) $\mathcal{E}$ induces inspection processes $\{\underline{v}_i^s\}_{s \in \mathcal{S}}$ and rank distribution ζ, and (2) the canonical outcome (μ, χ) of $\mathcal{M}$ (under σ) with initial information $(\mathcal{C}, \mathcal{S}, \zeta, q)$ is not regret-free stable. We also show how to construct the economy $\mathcal{E}$ as a function of $\mathcal{M}$, σ.

The economy $\mathcal{E}$ has three colleges $\mathcal{C} = \{1, 2, 3\}$ with capacities $q_1 = q_2 = q_3 = \frac{1}{4}$. The optimal inspection processes are such that $\underline{v}_1^s > \underline{v}_2^s, \underline{v}_3^s$ for all students $s \in X$, $\underline{v}_2^s > \underline{v}_3^s, \underline{v}_1^s$ for all students $s \in Y$, and $\underline{v}_3^s > \underline{v}_1^s, \underline{v}_2^s$ for all students $s \in Z$. The rank distribution ζ is such that there are four non-intersecting sets of student types $X, Y, Z, D \subseteq \mathcal{S}$ satisfying $\eta (X) = \eta (Y) = \eta (Z) = \frac{1}{8}$ and $\eta (X \cup Y \cup Z \cup D) = 1$, and

$$\text{priority at } 1 : r_1^y > r_1^z > r_1^x > r_1^d$$

$$\text{priority at } 2 : r_2^z > r_2^x > r_2^y > r_2^d$$

$$\text{priority at } 3 : r_3^x > r_3^y > r_3^z > r_3^d$$

$x \in X$, $y \in Y$, $z \in Z$, $d \in D$. This is possible by distributing sets of students of mass $\frac{1}{8}$, $\frac{1}{8}$, $\frac{1}{8}$ and $\frac{5}{8}$ uniformly on $[5/8, 6/8] \times [6/8, 7/8] \times [7/8, 1]$, $[7/8, 1] \times [5/8, 6/8] \times [6/8, 7/8]$, $[6/8, 7/8] \times [7/8, 1] \times [5/8, 6/8]$ and $[0, 5/8] \times [0, 5/8] \times [0, 5/8]$ respectively (these regions give us the students in X, Y, Z and D respectively). Note that all students in X are guaranteed to have colleges 2 and 3 in their budget set, but if all students in Y and Z demand college 1 then no students in X (or in D) will have college 1 in their budget set. In addition, all students in X wish to inspect college 1 first if it is available to them, since $\underline{v}_1^s > \underline{v}_2^s, \underline{v}_3^s$. A similar observation can be made for Y and Z. (Students in D are dummy students who ensure that the number of students is greater than the number of seats.)

We now specify the priors and costs for the students as a function of $\mathcal{M}$ and σ. These will be such that with probability 1 all students prefer being assigned to being unassigned.

Consider the first non-zero measure set of students S' who are suggested an inspection by $\mathcal{M}$ (or, if $\mathcal{M}$ is a choice-based process, for whom the set of colleges to inspect is non-empty, $\sigma^s(m^s) \neq \mathbf{0}$). As $\mathcal{M}$ does not have access to oracle information this set of students will not depend on the priors and costs. Without loss of generality, let S'' be a subset of S' of maximum size that satisfies that: (1) either $S'' \subseteq X$ or $S'' \subseteq D$; and (2) all students in S'' are suggested to inspect the same college, and let $\varepsilon := \eta(S'') > 0$ denote the measure of this set.

Case 1. Suppose $S'' \subseteq X \cup D$ and the suggested action is to inspect college 1. We first define priors and costs for students not in S'' so that almost all students in S'' will not have 1 in their budget set. For all $s \in (Y \cup Z) \setminus S''$, let $\tilde{F}_2^s, \tilde{F}_3^s \sim \mathit{Unif}[0, 1]$ (i.e. $\tilde{F}_2^s(x) = \tilde{F}_3^s(x) = x$ for $x \in [0, 1]$) and let $\tilde{F}_1^s(x) \sim \mathit{Bernoulli}\left(1 - \frac{1}{4}\varepsilon\right)$ (i.e. $\tilde{F}_1^s(x) = \frac{1}{4}\varepsilon$ for $x < 1$ and $\tilde{F}_1^s(1) = 1$), and let $\tilde{c}_i = 0$ for all i. Notice that for a student of type $\theta = \theta^s$ this induces demand $\tilde{\delta}^\theta$ such that if $1 \in C$ and

$|C| > 1$ then $\tilde{\delta}_1^\theta(C) = 1 - \frac{1}{4}\varepsilon$, i.e. a $1 - \frac{1}{4}\varepsilon$ fraction of students in $(Y \cup Z) \setminus S'$ always demand college 1, for all budget sets consistent with the rank distribution ζ. For all $s \in Y \setminus S''$, by Lemma 3 there exists $\left(F^\theta, c^\theta\right)$ such that (1) $\mathbb{E}_{F_i^\theta}\left[\left(v_i^\theta - \underline{v}_i^\theta\right)^+\right] = c_i^\theta$ for $\theta = \theta^s$ and all i; and (2) induces demand δ^θ such that if $1 \in C$ and $|C| > 1$ then $\tilde{\delta}_1^\theta(C) \geq 1 - \frac{1}{2}\varepsilon$. In other words, all students in $Y \setminus S''$ have the specified optimal inspection policy, and for all budget sets consistent with ζ least a $1 - \frac{1}{2}\varepsilon$ fraction of students in $Y \setminus S''$ demand college 1. We similarly define priors and costs for all students in $Z \setminus S''$ such that all students in $Z \setminus S''$ have the specified optimal inspection policy, and for all budget sets consistent with ζ at least a $1 - \frac{1}{2}\varepsilon$ fraction of students in $Z \setminus S''$ demand college 1. It follows that in the canonical outcome at least $2(1/8)\left(1 - \frac{1}{2}\varepsilon\right) = \frac{1}{4} - \frac{1}{8}\varepsilon$ students in $(Y \cup Z) \setminus S'$ demand 1, and so at least $\frac{7}{8}\varepsilon$ students in S'' regret inspecting college 1 under their canonical outcome under $\mathcal{M}, \sigma$ as it is not available to them.

Case 2. Suppose $S'' \subset X$ and the suggested action is to inspect college 2. We first define priors and costs for students not in S'' so that almost all students in S'' will have 1 in their budget set. For all $s \in Y \setminus S''$, let $\tilde{F}_1^s, \tilde{F}_2^s \sim \mathit{Unif}[0,1]$ (i.e. $\tilde{F}_1^s(x) = \tilde{F}_2^s(x) = x$ for $x \in [0,1]$) and let $\tilde{F}_3^s(x) \sim \mathit{Bernoulli}\left(1 - \frac{1}{4}\varepsilon\right)$ (i.e. $\tilde{F}_3^s(x) = \frac{1}{4}\varepsilon$ for $x < 1$ and $\tilde{F}_3^s(1) = 1$), and let $\tilde{c}_i = 0$ for all i. Notice that for a student of type $\theta = \theta^s$ this induces demand $\tilde{\delta}^\theta$ such that if $3 \in C$ and $|C| > 1$ then $\tilde{\delta}_3^\theta(C) = 1 - \frac{1}{4}\varepsilon$, i.e. a $1 - \frac{1}{4}\varepsilon$ fraction of students in $(Y \cup Z) \setminus S'$ always demand college 3, for all budget sets consistent with the rank distribution ζ. Hence by Lemma 3 there exists $\left(F^\theta, c^\theta\right)$ such that all students in $Y \setminus S''$ have the specified optimal inspection policy, and for all budget sets consistent with ζ least a $1 - \frac{1}{2}\varepsilon$ fraction of students in $Y \setminus S''$ demand college 3. We similarly define priors and costs for all students in $Z \setminus S''$ such that all students in $Z \setminus S''$ have the specified optimal inspection policy, and for all budget sets consistent with ζ at least a $1 - \frac{1}{2}\varepsilon$ fraction of students in $Z \setminus S''$ demand college 2. It follows that in the canonical outcome at most $2(1/8)\left(\frac{1}{2}\varepsilon\right) = \frac{1}{8}\varepsilon$ students in $(Y \cup Z) \setminus S'$ demand 1.

We now define priors and costs for students in S'' such that if college 1 is available then the unique optimal outcome for almost all these students is to inspect and take 1. For example, let $\tilde{F}_2^s, \tilde{F}_3^s \sim \mathit{Unif}[0,1]$, let $\tilde{F}_1^s(x) = 0$ for $x < 1$ and $\tilde{F}_1^s(1) = 1$, and let $\tilde{c}_i = 0.1$ for all i. By Lemma 3 there exists $\left(F^\theta, c^\theta\right)$ such that all students in S'' have the specified optimal inspection policy, and

if college 1 is available then at least a $\frac{6}{7}$ fraction of the students in S'' demand college 1. Hence at least $\frac{7}{8}\varepsilon \times \frac{6}{7} = \frac{3}{4}\varepsilon$ students in S'' regret inspecting college 2 under their canonical outcome under $\mathcal{M}, \sigma$ as their optimal outcome is to inspect 1 and stop.

Case 3. Suppose no students are suggested an inspection by $\mathcal{M}$. Then the resulting outcome (μ, χ) is not stable. This is because no student has inspected, so no student is assigned, but all students in X are willing to pay the cost to inspect colleges 2 and 3, and similarly all students in $Y \cup Z$ have a college for which they are willing to pay the cost to inspect.

By symmetry, this covers all possible cases and there always exists an economy $\mathcal{E}$ consistent with $\{\underline{v}_i^s\}_{s\in\mathcal{S},i\in\mathcal{C}}$ and ζ where the canonical outcome of $\mathcal{M}$ under σ is not regret-free stable. □

C Additional Impossibility Results for Regret-Free Stable Processes

We now demonstrate that standard processes can fail rather spectacularly in learning market-clearing cutoffs and alleviating the costs associated with information acquisition. Intuitively, in choice-based processes 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 processes, 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 processes, we prove the stronger result that for general economies even multiple-round processes must either incur regret, or create an information deadlock, where every student waits for others to acquire information first.

C.1 Direct One-Shot Processes

To demonstrate the issues in computing regret-free stable matchings, let us first consider the case where all students are willing to inspect any college as long as it is in their budget set. We may

view this as a setting where the costs affect *which* college*s* students are willing to inspect, but not the *order* in which they are willing to inspect them. If each student is uncertain about aggregate demand, then the standard implementation of Deferred Acceptance as a one-shot game will not be regret-free even if they are willing to inspect all colleges in their budget set. This is because students' budget sets will depend on the preferences of other students, and so students who have low priority at the colleges 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 colleges each with capacity $q_i = 1$. Suppose that college priorities are perfectly aligned, i.e. $r_i^s = r_j^s$ 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 college. Such demand can be rationalized e.g. by the 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 costs $c_i^s = 1$ for all $s \in \mathcal{S}$.

In any one-shot choice-based process, a student s will have no regret only if she chooses to examine precisely the set of all colleges not selected by higher-ranked students. This is because a student is willing to incur the cost to examine any college 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.[18]

This example demonstrates that in settings with aggregate uncertainty single-shot choice-based processes cannot hope to find regret-free stable matchings, even in settings where students are willing to incur the costs of searching any number of colleges, due to their inability to coordinate the search of worse-ranked students with the preferences of the better-ranked students.

[18] 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.

C.2 Student-Optimal Regret-Free Stable Processes

In this section we provide an example demonstrating that even in settings where it is possible to implement a regret-free stable choice-based process, if there is aggregate uncertainty or if the market designer does not know aggregate demand it may be impossible to verify that a matching is student-optimal without incurring regret.

Example 5. Consider a discrete economy E with two colleges $\mathcal{C} = \{1, 2\}$ with capacities $q_1 = q_2 = 1$ and 2 students $S = \{x, y\}$.[19]

Suppose that college 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 college have discrete distribution $\mathbb{P}\left(v_i^s = 1\right) = \mathbb{P}\left(v_i^s = 2\right) = \frac{1}{4}$, $\mathbb{P}\left(v_i^s = 2^k\right) = \frac{1}{2^k}$ for all $k > 1$ and $\mathbb{P}\left(v_i^s = x\right) = 0$ 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 $\mathbb{E}\left[(v_i^s - \underline{v})\right] = \infty$ for all s, i and $\underline{v} \in \mathbb{R}$ it follows that both students' optimal strategies are to inspect all the colleges that are available to them.[20]

Note that the matching $\mu = \mu^{college}$ defined by $\left(\mu(x), \mu(y)\right) = (2, 1)$ is always regret-free stable, and is the college-optimal regret-free stable matching. Let $\mu' = \mu^{student}$ be defined by $\left(\mu'(x), \mu'(y)\right) = (1, 2)$. 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 $\mu^{college}$ is the student-optimal regret-free stable matching subject to event X, as both x and y obtain their highest valued colleges, 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 colleges.

[19] 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 college 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.

[20] 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 college 1 or student y to inspect college 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 $\mu = \mu^{college}$ is the student-optimal regret-free stable matching then each college is assigned their top choice student, and so based on college 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 college 2, and student y only inspects college 1. However, if $\mu' = \mu^{student}$ is the student-optimal regret-free stable matching then under $\chi^{RF}(\mu'|\cdot)$ both students inspect both colleges. Thus, since it is impossible to distinguish between events X and X' without requiring either student x to inspect college 1 or student y to inspect college 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 $\mu^{college}$ 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.

D Estimating Regret-Free Stable Cutoffs

D.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 colleges $\mathcal{C}$ 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, $d(\mu, \mu') = \max_{P,P':\mathcal{M}(P)=\mu,\mathcal{M}(P')=\mu'} \|P - P'\|$.[21] Given a finite economy $E = (\mathcal{C}, S, q)$ define the

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

continuum economy $\Phi(E) = \left(\mathcal{C}, \mathcal{S}, \eta, \frac{q}{|S|}\right)$ 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 $\mathcal{E}$.

The following two theorems follow directly from the analogous results in Azevedo and Leshno (2016), as well as observing that the set of regular measures is open.

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.*

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 $E^k = \left(\mathcal{C}, S^k, qk\right)$ be a randomly drawn finite economy, where S^k is given by drawing k students 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

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

In other words, Theorem 8 states that given an 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^*. Another interpretation is that if we estimate cutoffs by randomly sampling k students, then the estimated cutoffs P^k are normally distributed around P^*.

D.2 Proofs for Section 5.3

D.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 $m(X,P) = \mathbf{1}\{X \text{ is 'past'}\}\left(D^X(P) - \overline{q}\right)$ be the excess demand of a past student, and let $g(X,q,P) = \mathbf{1}\{X \text{ is 'present'}\}\left(D^X(P) - q\right)$ be the excess demand of a present student if the capacities are q. Recall that $\hat{P}$ are the market-clearing cutoffs for the economy E^{hist} and hence we may invoke Theorem 7 for discrete economies converging to a continuum economy to show that $\hat{P}$ satisfies $\frac{1}{k}\sum_{i=1}^{k} m\left(X_i, \hat{P}\right) \to 0$. Similarly $\hat{q}^k k$ is the demand resulting from posting $\hat{P}$ for the economy E and hence $\frac{1}{k}\sum_{i=1}^{k} g\left(X_i, \hat{q}^k, \hat{P}\right) \to 0$. Since the capacities $\overline{q}$ are such that there is a unique regret-free stable outcome, it follows that $\mathbb{E}\left[\sup \|\nabla_q g(X, \overline{q}, P^*)\|\right] < \infty$.

It follows $\sqrt{k}\left(\hat{q}' - \overline{q}\right) \xrightarrow{d} \mathcal{N}(0, V)$, where

$$V = \frac{1}{1+\alpha} G_q^{-1} \mathbb{E}\left[AA^T\right] G_q^{-T},$$

for $G_q = \mathbb{E}\left[\nabla_q g(X, \overline{q}, P^*)\right] = -\frac{1}{\alpha+1}$, $G_P = \mathbb{E}\left[\nabla_P g(X, \overline{q}, P^*)\right] = \frac{1}{\alpha+1}\mathbb{E}\left[\nabla_P D^X(P^*)\right]$, $M = \mathbb{E}\left[\nabla_P m(X, P^*)\right] = \frac{\alpha}{\alpha+1}\mathbb{E}\left[\nabla_P D^X(P^*)\right]$, and

$$\begin{aligned} A &= g(X, \overline{q}, P^*) - G_P M^{-1} m(X, P^*) \\ &= g(X, \overline{q}, P^*) - \frac{1}{\alpha} m(X, P^*), \end{aligned}$$

i.e.

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

Now $\text{var}(m(X, P^*)) = \text{var}\Big(Y\left(D^X(P^*) - \overline{q}\right)\Big)$ where $Y \sim Bernoulli\left(\frac{\alpha}{\alpha+1}\right)$ is independent

from $D^X(P^*) - \overline{q}$, and hence

$$\begin{aligned}
\text{var}(m(X, P^*)) &= \mathbb{E}\left[Y^2\right]\mathbb{E}\left[\left(D^X(P^*) - \overline{q}\right)^2\right] - \mathbb{E}[Y]^2\,\mathbb{E}\left[D^X(P^*) - \overline{q}\right]^2 \\
&= \frac{\alpha}{\alpha+1}\mathbb{E}\left[\left(D^X(P^*) - \overline{q}\right)^2\right] \\
&= \frac{\alpha}{\alpha+1}\text{var}\left(D^X(P^*) - \overline{q}\right) \\
&= \frac{\alpha}{\alpha+1}\Sigma^{\overline{q}}.
\end{aligned}$$

Similarly $\text{var}(g(X, \overline{q}, P^*)) = \text{var}\Big(Z\left(D^X(P^*) - \overline{q}\right)\Big)$ where $Z \sim Bernoulli\left(\frac{1}{\alpha+1}\right)$ is independent from $D^X(P^*) - \overline{q}$, and hence

$$\text{var}(g(X, \overline{q}, P^*)) = \frac{1}{\alpha+1}\Sigma^{\overline{q}}.$$

Finally

$$\text{cov}(g(\cdot), m(\cdot)) = \mathbb{E}\left[g(X)\, m(X)^T\right] - \mathbb{E}[g(X)]\,\mathbb{E}[m(X)]^T = 0$$

since $g(X)\, m(X)^T = 0$ (since X cannot be both 'past' and 'present') and $\mathbb{E}[g(X)] = \mathbb{E}[m(X)] = 0$.

Hence

$$\begin{aligned}
V &= (1+\alpha)\left(\text{var}(g(\cdot)) + \frac{1}{\alpha^2}\text{var}(m(\cdot))\right) \\
&= \Sigma^{\overline{q}} + \frac{1}{\alpha}\Sigma^{\overline{q}} \\
&= \frac{1+\alpha}{\alpha}\Sigma^{\overline{q}}
\end{aligned}$$

as required.

D.2.2 Example 3

We first show that the outcome (μ, χ) after both rounds is regret-free stable with respect to realized demand $\hat{q}^k k$. It suffices to show that for all students $s \in S^f(E)$ with free market information it

holds 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 $B^s = \{i \,|\, r_i^s \geq 1 - q_i\}$. 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 colleges acceptable and so if $r_i^{s'} \geq r_i^s$ then student s' is assigned to some college. 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

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

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 $m(X, \Gamma) = \mathbf{1}_{\left\{X \in S^f(\mathcal{E})\right\}} D^X\left(P^f | \eta(\Gamma)\right) - q^f$ and let $g(X, q, \Gamma) = m(X, \Gamma) + \mathbf{1}_{\left\{X \notin S^f(\mathcal{E})\right\}} D^X\left(P^*(\Gamma) | \eta(\Gamma^*)\right) - \left(q - q^f\right)$. Note that by Theorem 7 it holds that $\frac{1}{k}\sum_{i=1}^k m\left(X_i, \hat{\Gamma}\right) \to 0$ and $\frac{1}{k}\sum_{i=1}^k g\left(X_i, \hat{q}^k, \hat{\Gamma}\right) \to 0$.

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

$$\begin{aligned} B &= g(X, \overline{q}, \Gamma^*) - G_\Gamma M^{-1} m(X, \Gamma^*) \\ &= g(X, \overline{q}, \Gamma^*) - (I + A)\, m(X, \Gamma^*). \end{aligned}$$

Then $\sqrt{k}\left(\hat{q}^k - \overline{q}\right) \xrightarrow{d} \mathcal{N}(0, V)$, where

$$V = G_q^{-1}\mathbb{E}\left[BB^T\right] G_q^{-T},$$

for $G_q = \mathbb{E}\left[\nabla_q g(X, \overline{q}, \Gamma^*)\right] = -1$, i.e.

$$V = \mathrm{var}(B) = \mathrm{var}\left(g(X, \overline{q}, \Gamma^*) - (I + A)\, m(X, \Gamma^*)\right).$$

Note that $\mathrm{var}\left(g(X, \overline{q}, \Gamma^*)\right) = \mathrm{var}\left(D^X(\Gamma^*)\right) = \Sigma^{\overline{q}}$, $\mathrm{var}\left(m(X, \Gamma^*)\right) = \mathrm{var}\left(\mathbf{1}_{\left\{X \in S^f(\mathcal{E})\right\}} D^X\left(P^f | \eta(\Gamma^*)\right)\right) = \Sigma^{q^f}$, and $\mathrm{cov}(g(\cdot), m(\cdot)) = \mathrm{var}\left(m(\cdot)\right)$, and so

$$V = \Sigma^{\overline{q}} + (A - I)\,\Sigma^{q^f}\,(A + I)^T\,.$$

Finally, we may rewrite M and A in terms of demand functions to obtain that

$$\begin{aligned}
M &= \mathbb{P}\left(X \in S^f(\mathcal{E})\right)\mathbb{E}\left[\nabla_\Gamma D^X\left(P^f|\eta\left(\Gamma, X \in S^f(\mathcal{E})\right)\right)\right] \\
&= \nabla_\Gamma D^f(\Gamma^*)\,, \\
A &= \mathbb{P}\left(X \notin S^f(\mathcal{E})\right)\mathbb{E}\left[\nabla_\Gamma D^X\left(P^*(\Gamma)\,|\eta(\Gamma^*), X \notin S^f(\mathcal{E})\right)\right] M^{-1} \\
&= \nabla_\Gamma D^r\left(P^*(\Gamma)\,|\eta(\Gamma^*)\right)|_{\Gamma=\Gamma^*}\left(\nabla_\Gamma D^f(\Gamma^*)\right)^{-1},
\end{aligned}$$

so A is the product of the sensitivity of the remaining demand $D^r(\cdot|\eta(\Gamma^*)) = D(\cdot|\eta(\Gamma^*)) - D^f(\Gamma^*)$ to errors in the cutoffs $P^*(\Gamma)$ from misspecifying Γ, with the inverse sensitivity of the demand $D^f(\cdot)$ of students with free market information. Finally, since $D^f(\Gamma^*)$ is constant, it follows that $\nabla_\Gamma D^r\left(P^*(\Gamma)\,|\eta(\Gamma^*)\right)|_{\Gamma=\Gamma^*} = \nabla_\Gamma D\left(P^*(\Gamma)\,|\eta(\Gamma^*)\right)|_{\Gamma=\Gamma^*}$.