

He, Yinghua; Magnac, Thierry

Working Paper

A Pigouvian Approach to Congestion in Matching Markets

IZA Discussion Papers, No. 11967

Provided in Cooperation with:

IZA – Institute of Labor Economics

Suggested Citation: He, Yinghua; Magnac, Thierry (2018) : A Pigouvian Approach to Congestion in Matching Markets, IZA Discussion Papers, No. 11967, Institute of Labor Economics (IZA), Bonn

This Version is available at:

<https://hdl.handle.net/10419/193261>

Standard-Nutzungsbedingungen:

Die Dokumente auf EconStor dürfen zu eigenen wissenschaftlichen Zwecken und zum Privatgebrauch gespeichert und kopiert werden.

Sie dürfen die Dokumente nicht für öffentliche oder kommerzielle Zwecke vervielfältigen, öffentlich ausstellen, öffentlich zugänglich machen, vertreiben oder anderweitig nutzen.

Sofern die Verfasser die Dokumente unter Open-Content-Lizenzen (insbesondere CC-Lizenzen) zur Verfügung gestellt haben sollten, gelten abweichend von diesen Nutzungsbedingungen die in der dort genannten Lizenz gewährten Nutzungsrechte.

Terms of use:

Documents in EconStor may be saved and copied for your personal and scholarly purposes.

You are not to copy documents for public or commercial purposes, to exhibit the documents publicly, to make them publicly available on the internet, or to distribute or otherwise use the documents in public.

If the documents have been made available under an Open Content Licence (especially Creative Commons Licences), you may exercise further usage rights as specified in the indicated licence.

DISCUSSION PAPER SERIES

IZA DP No. 11967

**A Pigouvian Approach to Congestion in
Matching Markets**

Yinghua He
Thierry Magnac

NOVEMBER 2018

DISCUSSION PAPER SERIES

IZA DP No. 11967

A Pigouvian Approach to Congestion in Matching Markets

Yinghua He

Rice University and Toulouse School of Economics

Thierry Magnac

Toulouse School of Economics and IZA

NOVEMBER 2018

Any opinions expressed in this paper are those of the author(s) and not those of IZA. Research published in this series may include views on policy, but IZA takes no institutional policy positions. The IZA research network is committed to the IZA Guiding Principles of Research Integrity.

The IZA Institute of Labor Economics is an independent economic research institute that conducts research in labor economics and offers evidence-based policy advice on labor market issues. Supported by the Deutsche Post Foundation, IZA runs the world's largest network of economists, whose research aims to provide answers to the global labor market challenges of our time. Our key objective is to build bridges between academic research, policymakers and society.

IZA Discussion Papers often represent preliminary work and are circulated to encourage discussion. Citation of such a paper should account for its provisional character. A revised version may be available directly from the author.

ABSTRACT

A Pigouvian Approach to Congestion in Matching Markets*

Recruiting agents, or “programs” costly screen “applicants” in matching processes, and congestion in a market increases with the number of applicants to be screened. To combat this externality that applicants impose on programs, application costs can be used as a Pigouvian tax. Higher costs reduce congestion by discouraging applicants from applying to certain programs; however, they may harm match quality. In a multiple-elicitation experiment conducted in a real-life matching market, we implement variants of the Gale-Shapley Deferred-Acceptance mechanism with different application costs. Our experimental and structural estimates show that a (low) application cost effectively reduces congestion without harming match quality.

JEL Classification: D78, D47, D50, D61, I21

Keywords: Gale-Shapley Deferred Acceptance Mechanism, costly preference formation, screening, stable matching, congestion, matching market place

Corresponding author:

Thierry Magnac
Université de Toulouse Capitole
Manufacture des Tabacs
21 Allée de Brienne
31015 Toulouse
France
E-mail: thierry.magnac@tse-fr.eu

* We dedicate this paper to Jean-Philippe Lesne, whose enthusiasm and initiatives made this research possible. David Alary, the M2 program directors, and many others at TSE provided crucial help. For excellent research assistance, we thank Christophe Lévêque and Adrian Torchiana. We also thank the conference/ seminar participants at CEMMAP, Cowles Conference at Yale, CREST-Paris, ECARES, ESEM 2017, IAAE 2018, London, Matching in Practice (Toulouse), Nottingham, Paris School of Economics, Texas Camp Econometrics, TSE, University of Houston, Warwick, Wisconsin-Madison, WZB (Berlin), and Zhejiang University. This study was approved by the Toulouse School of Economics Research Ethics Committee (# 27052013). We gratefully acknowledge financial support from the European Research Council under the European Community's Seventh Framework Program FP7 2007-2013 grant agreement No 295298. The usual disclaimer applies.

1 Introduction

In many commonly observed matching markets, agents on one side, called “programs” (e.g., employers or colleges), actively recruit agents on the other side, called “applicants” (e.g., employees or students). The matching process usually includes an application stage, in which applicants submit their candidacy to programs of their choice; it is followed by a screening stage, in which each program costly screens and ranks its candidates.

As a result of massive advancements in both information technology and market design, there has been a largely welcomed trend toward applicants applying to an increasing number of programs. For example, in the job market for new graduates of PhD programs in economics, the centralized platforms – EconJobMarket and AEA Job Opening for Economists – have dramatically reduced the application cost associated with submitting a job application. As a result, hiring departments now screen many more candidates (Nguyen, Peters and Poitevin, 2018).

While the traditional matching theory considers application cost as a market friction (see, e.g., Rogerson, Shimer and Wright, 2005) and welcomes these new developments, a practical issue has emerged: Programs must screen many more applicants and pay a high cost, which amounts to a negative externality on programs as a result of applicant behavior. Costly screening to form preferences is quite common in real-life scenarios. Recruiting programs have enough information to rank applicants only after reviewing their files or conducting interviews. When programs receive too many applications, congestion occurs, and matching outcomes can be (constrained) inefficient (Arnosti, Johari and Kanoria, 2016).

This externality motivates us to consider application cost as a Pigouvian tax. Intuitively, applicants apply to fewer programs when the application cost is higher, which mitigates market congestion. However, the total welfare effects, including those on match quality, are ambiguous because a high application cost may preclude some efficient matches.

In a setting of many-to-one matching without transfers, we provide the first empirical evidence on the performance of a set of popular market designs with different schemes of application costs. Similar to implementations of Pigouvian taxes in other settings (Pigou, 1920), we hypothesize that when designed appropriately, an application cost can enhance welfare by reducing negative externalities. Inspired by practical market designs, we consider two forms of application cost: A positive marginal cost and an application limit, both of which are used in real-life centralized and decentralized matching markets (see the examples in Section 6, especially Table 16).

Our research design is original in that it marries real-life, experimental data with a structural analysis. We implement a multiple-elicitation experiment in which application costs vary across market designs. The experiment allows the effects of application costs on congestion and match quality to be directly evaluated. Experimental variation in market design in the field is, however, necessarily limited, and we further conduct comprehensive counterfactual analyses based on the estimation of a structural model using the experimental data.

We focus on two dimensions of a matching outcome: the congestion/screening costs, measured by the number of applicants to screen, and match quality, measured by the numbers of blocking pairs, the number of unmatched applicants, and the welfare of both sides. A pair of applicant and program blocks a matching, if both would be better off by being matched together after leaving their current matches. The stability of a matching, defined as the absence of any blocking pair, is the key to the success of matching markets (Roth, 1991). Importantly, stability implies Pareto efficiency when both sides are endowed with strict preferences (Abdulkadiroglu and Sönmez, 2013).

Our experiment involves the real-life matching of 129 applicants to the 7 master's programs at the Toulouse School of Economics (TSE) and was conducted in May 2013 for admission in the 2013-14 academic year. The experimental market designs are three variants of the Gale-Shapley Deferred-Acceptance mechanism (DA) encountered in practice: The traditional DA, under which applicants can apply to all programs without any cost; DA with truncation (DA-T), under which applicants can apply to no more than four programs; and DA with cost (DA-C), under which applicants must write a motivation letter for each additional application beyond the first three applications. Under each mechanism, every applicant is required to submit a rank-ordered list of programs (ROL). As applicants are further informed that one of the mechanisms will be implemented, they have incentives to behave optimally under each mechanism.

First, the direct experimental evaluation of the three market designs, shows that both DA-T and DA-C reduce the screening cost of programs without significantly harming match quality. On average, relative to DA without cost which results in 75 applications per program, each program receives 13 fewer applications under DA-T and 16 fewer under DA-C. Such reductions in congestion are substantial despite the applicant preferences being rather heterogenous – 85% of the applicants always obtain their most-preferred program across the market designs.

In our second empirical approach, the structural analysis, we need information on appli-

cant and program preferences as well as applicant beliefs. Because of the strategy-proofness of DA (Dubins and Freedman, 1981; Roth, 1982) and because we announce in the experiment that it is optimal for applicants to report their true preferences under DA, ROLs submitted by applicants under DA are assumed to reveal their true ordinal preferences. We supplement this information with a survey in which each applicant reports whether a program is unacceptable to her. This leads to an extended version of the exploded logit model for the estimation of applicant cardinal preferences. Furthermore, program preferences are estimated using their observed ranking of applicants. Applicant beliefs are derived from an applicant’s (incomplete) information on other applicants’ preferences and on how each program ranks her. Assuming rational expectations and a common prior, we construct applicant beliefs from estimated parameters governing program and applicant preferences.

We then simulate the equilibrium outcomes of various configurations of DA-T and DA-C. DA-T can be of any degree K , where $K \in \{1, \dots, 7\}$ is the number of choices an applicant is allowed to submit. Constant marginal costs for every application in DA-C take various values over a wide range. In equilibrium, applicants choose an optimal strategy under each market design, while the ranking of applicants by programs is assumed to be constant.

The results show that relative to DA without cost, DA-C with a low cost reduces screening costs by half without harming match quality; a less-restrictive DA-T, e.g., DA-T-5 and DA-T-6, does not affect match quality, either, but the reduction in congestion is not as substantial. As expected, a high-cost DA-C or a highly restrictive DA-T results in lower match quality because a prohibitively high application cost prevents applicants from applying to a sufficient number of programs. This leads to a high number of blocking pairs and unmatched applicants.

These findings show that application costs have great potential for combating congestion even if the application cost is purely wasteful, and more so when it is a monetary transfer. In practice, identifying the optimal Pigouvian tax requires information on demand and supply, and our setting is no different in this respect.

Organization of the paper. We complete this section by a brief review of the related literature. Section 2 formalizes the many-to-one matching and summarizes the theoretical predictions. The experiment conducted at TSE is described in Section 3, and Section 4 summarizes the data and provides a direct evaluation. Our structural evaluation is presented in Section 5. The paper concludes in Section 6 with a discussion of our findings and their implications for practical market design.

Related literature. Congestion in matching markets has been studied extensively, but often theoretically. One of the earliest papers on the topic is Roth and Xing (1997), who show that thick markets may suffer from congestion. Search and screening costs hinder the evaluation of all potential matches by both sides.

Most related to our research, Arnosti et al. (2016) study the regulation of market congestion using application costs and provide a comprehensive survey of the literature. Their paper considers a dynamic model of one-to-one matching in which agents (applicants and employers) arrive and depart over time. All applicants and employers are homogenous. As in our paper – but in a decentralized setting – each applicant applies to programs at a cost, and each employer pays a screening cost to verify whether an applicant is compatible. If the applicant is compatible, the employer makes her an offer without knowing if she has already been hired by someone else. If she receives an offer, the applicant accepts or rejects it. These authors also consider imposing application limits that restrict the maximum number of applications allowed per applicant.

Application limits are considered by Che and Koh (2016), who model a college admissions game in which colleges manage yield by strategically targeting applicants who are less likely to receive other offers. Application costs are present in Chade, Lewis and Smith (2014). More recently, in a model of matching economics PhDs to university positions, Nguyen et al. (2018) note that a reduced application cost increases the probability of application. Furthermore, they show that this reduced cost causes an increase in the probability that offers are turned down by applicants and an increase in the probability that positions remain unmatched. As a result, some universities' welfare may decrease.

Another way educational institutions mitigate screening costs is to organize entrance exams. For example, in the 2000s, certain Brazilian universities used a low-cost first-round exam to screen out applicants who were unlikely to qualify, while a second-round exam, which was costlier and more informative, refined the selection (Carvalho, Magnac and Xiong, forthcoming). This two-tier screening device, a centralized exam followed by a costlier decentralized exam, has also been used in Japan (Hafalir, Hakimov, Kübler and Kurino, 2016).

Our paper also sheds light on the design of online platforms, an extensively studied topic (Fradkin, 2014; Halaburda, Piskorski and Yildirim, 2015; Horton, 2015). That literature often shows in different contexts that limiting the number of potential matches might have desirable properties in terms of aggregate welfare because of decreased screening costs. For example, in a dynamic search model with costly discovery of match quality, Kanoria and

Saban (2017) find that a platform can mitigate wasteful competition in match-partner search via restricting what agents can see/do.

Signaling in matching markets also relates to the issues we study (Coles, Kushnir and Niederle, 2013; Lee and Schwarz, 2017). Programs can target applicants by letting applicants signal their interests in specific programs. To avoid cheap talk, signals are made costly or are limited in terms of the total number of signals permitted. This step may improve welfare in equilibrium because offers are made to applicants who have a higher probability of accepting the offer (Coles et al., 2013). In decentralized settings, pre-match interviews can be organized between applicants and programs (Lee and Schwarz, 2017), which are costly for both sides in terms of application costs and screening costs. Focusing on the information about their preferences agents need to learn and communicate with others to form the final matching, Ashlagi, Braverman, Kanoria and Shi (2018) explore signaling and recommendation protocols with which a market reaches a desirable equilibrium outcome with a small communication cost.

Our study is also related to the literature on labor markets with costly search (for a survey, see Rogerson et al., 2005). However, the externality between applicants and programs goes in the opposite direction in that setting, as decreasing individual search costs leads to better prospects for programs, while in our setting, decreasing these costs can be detrimental to program welfare because of screening costs. An exception in this literature is Seabright and Sen (2015). In their theoretical model, a reduction in application costs can lower a firm's payoff, because it attracts lower quality applicants.

The use of Pigouvian taxes to correct negative externalities is originally attributable to Pigou (1920) and has been extensively studied theoretically (see, for example, Baumol, 1972; Sandmo, 1975). Recently, it has received increased attention, and in particular, the effectiveness of Pigouvian taxes in regulating pollution has been empirically evaluated; see for example Muller and Mendelsohn (2009) and Knittel and Sandler (2013).

In terms of theoretical tools, the two-sided matching framework we study is well analyzed and summarized in Roth and Sotomayor (1990). More specifically, our model builds on the theoretical results derived by Haeringer and Klijn (2009) and Fack, Grenet and He (Forthcoming) in which the properties of the DA mechanism with truncation and/or application costs are investigated.

The estimation of applicant preferences from applicants' submitted ROLs is an important building block of our empirical analysis. In a strand of the recent empirical school choice literature, applicants are assumed to report their true preferences, as in our experiment under

DA. Standard discrete choice methods are extended – e.g., as an exploded logit – to utilize the identifying information contained in ROLs (Hastings, Kane and Staiger, 2008; Abdulkadiroğlu, Agarwal and Pathak, 2017; Pathak and Shi, 2017). Another strand of this literature explicitly considers possible strategic applicant behavior, for example, using data from the Boston immediate-acceptance mechanism (Agarwal and Somaini, Forthcoming; Cal-samiglia, Fu and Güell, 2014; He, 2015; Hwang, 2017) or from DA-T (Ajayi, 2013; Carvalho et al., forthcoming; Fack et al., Forthcoming). In contrast to our study of both sides’ welfare, this literature often focuses on student welfare only.

More generally, this paper builds on a growing body of literature in which structural methods are applied to real-life, experimental data and thus enlarge the span of counterfactual policy analyses (for a survey, see Blundell, 2017).

2 Many-to-one Matching: Set-up

A many-to-one matching market in which applicants are to be matched with programs is denoted by

$$\left\{ [v_{i,j}]_{i \in \mathcal{I}, j \in \mathcal{J} \cup \{0\}}, [s_{i,j}]_{i \in \mathcal{I} \cup \{0\}, j \in \mathcal{J}}, [q_j]_{j \in \mathcal{J}}, C_A(\cdot), C_P(\cdot) \right\},$$

where $\mathcal{I} \equiv \{1, \dots, I\}$ is the set of applicants and $\mathcal{J} \equiv \{1, \dots, J\}$ is the set of programs, with the addition of 0 as an outside option or the option of being unmatched. Applicant i receives a von Neumann-Morgenstern utility $v_{i,j} \in \mathbb{R}$ if matched with j ; and $v_i \equiv (v_{i,0}, v_{i,1}, \dots, v_{i,J})$. Programs have responsive preferences, and more specifically, program j values its match with i at $s_{i,j} \in \mathbb{R}$ regardless of who else is matched with j . Let $s_i = (s_{i,0}, \dots, s_{i,J})$. Each applicant can be matched with at most one program, and each program has a capacity of q_j identical positions. There is no indifference in preferences on either side. Program j is acceptable to i if $v_{i,j} \geq v_{i,0}$, and otherwise, it is unacceptable. Program j finds applicant to be *qualified* if $s_{i,j} \geq s_{0,j}$; otherwise, i is unqualified for j . Moreover, applicant i is qualified for program j if and only if i meets j ’s prerequisites.

Deviating from the previous literature, we assume that it is costly for an applicant to “apply” to programs, $C_A(\cdot) \geq 0$, while every program has to conduct a costly screening of its candidates to form its preferences, $C_P(\cdot) \geq 0$. Both cost functions, which are defined below, are homogeneous across i and j , respectively.

We consider a centralized market. Programs first announce their capacities and prerequisites, and every applicant then submits a rank-ordered list (ROL) of $K_i \leq J$ programs, denoted by $L_i = (l_i^1, \dots, l_i^{K_i})$, where $l_i^k \in \mathcal{J}$ is i ’s k^{th} choice. An ROL defines a relationship

$>_{L_i}$ such that $j >_{L_i} j'$ if and only if j is ranked above j' in L_i . The set of all possible ROLs, \mathcal{L} , includes all ROLs that rank at least one program. We define program j 's *candidates* as the set of applicants who include program j in their submitted ROLs.

Application cost and screening cost. When submitting ROL L , an applicant incurs a cost that is assumed to depend on the number of programs being ranked in L , denoted by $|L|$, but not on how they are ranked, $C_A(|L|) : \{0, \dots, J\} \rightarrow \overline{\mathbb{R}}^+ \equiv [0, +\infty]$. Similarly, upon receiving an application, a program pays a screening cost to learn its value of being matched with the applicant. The total screening cost of forming preferences over the (subset of) candidates \mathcal{I}_j is $C_P(|\mathcal{I}_j|) : \{0, \dots, I\} \rightarrow \overline{\mathbb{R}}^+$. For simplicity, we assume that a program always pays the cost of forming preferences over all its candidates. In other words, programs do not strategically choose to remain uninformed about their preferences over (a subset of) its candidates. Moreover, as we shall see in Section 4.1.4, programs in centralized markets are often prohibited from using information on candidates' ROLs to screen them, while in decentralized markets, applicants do not disclose their ROLs to any party.

Both the application cost and the screening cost monotonically increase such that for all $n \geq 0$, $C_A(n+1) \geq C_A(n)$ and $C_P(n+1) \geq C_P(n)$. These specifications are rather flexible, and in particular, such application cost captures many common practices of matching markets, as we show shortly.

2.1 Matching and Matching Mechanisms

We define a matching $\mu : \mathcal{I} \rightarrow \mathcal{J} \cup \{0\}$ such that (i) $\mu(i) = j$ if applicant i is matched with j ; (ii) $\mu(i) = 0$ if applicant i is unmatched; and (iii) $\mu^{-1}(j)$ is the set of applicants matched with j , while $|\mu^{-1}(j)|$ is the number of applicants matched with j and does not exceed j 's capacity.

A matching μ is **individually rational** if each applicant prefers her current match to remaining unmatched and if each program prefers each of its currently matched applicants to having a vacant position. Given matching μ , (i, j) form a **blocking pair** if i prefers j over her matched program $\mu(i)$, while i is preferred by j to either having a vacant position or keeping the least-preferred applicant of the currently matched ones. μ is **stable** if there is no blocking pair and if it is individually rational. Stability implies Pareto efficiency when both sides have strict preferences,¹ and stability is essential to the success of many matching

¹Abdulkadiroglu and Sönmez (2013) show this result for one-to-one matching, and we can extend it to many-to-one matching with responsive preferences. Suppose that μ' Pareto dominates a stable matching

markets (Roth, 1991).

Deferred Acceptance and its variants. The applicant-program match is solved by a matching mechanism that is a function of applicants' submitted ROLs and programs' preferences over applicants. As a computerized algorithm, the applicant-proposing DA works as follows:

Round 1. Every applicant applies to the first choice listed in her ROL. Each program rejects unqualified applicants and the least-preferred applicants in excess of its capacity and tentatively holds the other applicants.

Generally, later rounds are described as follows:

Round k . Every applicant who is rejected in Round $(k - 1)$ applies to the next choice on her ROL. Each program, pooling new applicants and those who were held in Round $(k - 1)$, rejects unqualified applicants and the least preferred applicants in excess of its capacity. Those who are not rejected are tentatively held by the programs.

The process ends after the round in which rejections are no longer issued. Each program is then matched with the applicants it is currently holding.

The three variants of DA that we consider differ only in their application costs. In the traditional DA mechanism, denoted by DA, the application cost is always zero: $C(|L|) = 0$ for all $L \in \mathcal{L}$. The DA with truncation, denoted by DA-T, does not allow applicants to apply to more than $K \in \{1, \dots, J\}$ programs, which is defined as DA-T of degree K and denoted by DA-T- K . Therefore, $C(|L|) = 0$ whenever $|L| \leq K$; otherwise, $C(|L|) = +\infty$. The last is DA with costs, denoted by DA-C, under which applicants must pay a cost for each program beyond their top K choices. In this case, $C(|L|) = c \times (|L| - K) \times \mathbb{1}(|L| > K)$ if the marginal cost of application beyond K is c , a positive constant.

2.2 Timeline, Information, Strategy, and Equilibrium Concept

We consider a matching game with incomplete information and a timeline depicted in Figure 1. First, at the “announcement” stage, a mechanism is chosen and made public to both programs and applicants; in addition, programs reveal their capacities and prerequisites.

μ . There must exist a program or an applicant who is strictly better off in μ' . Begin by assuming that applicant i strictly prefers $\mu'(i)$ to $\mu(i)$. Because μ is stable and because program preferences are strict, program $\mu'(i)$ must be strictly worse off in matching μ' than in matching μ ; otherwise, i and $\mu'(i)$ could form a blocking pair in μ . This situation contradicts the Pareto domination presumption. We would reach the same contradiction had we begun by assuming that there exists program j that strictly prefers $\mu'^{-1}(j)$ to $\mu^{-1}(j)$. Therefore, the stable matching μ is Pareto efficient.

Prerequisites define the necessary and sufficient conditions for qualification.

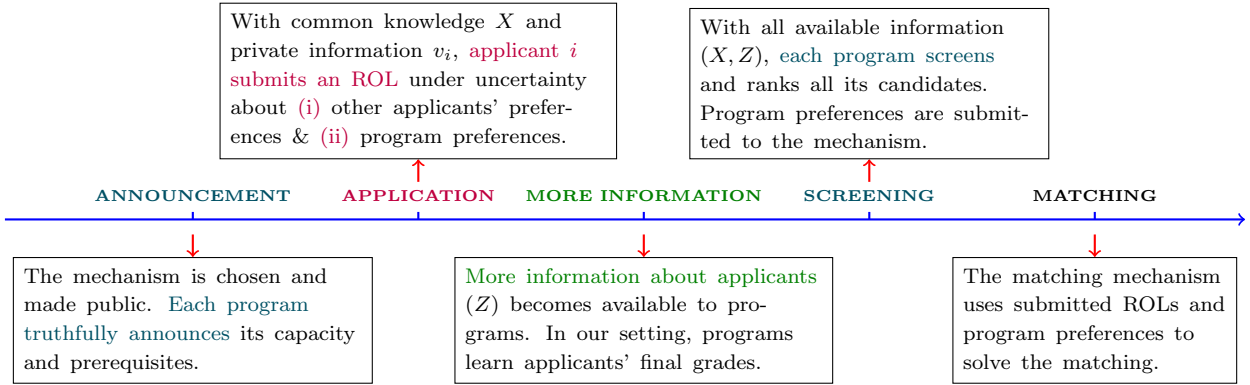


Figure 1: Timeline of the Matching Game under a Given Mechanism

Second, at the “application” stage, applicant preferences (v_i) are private information, while the distribution of applicant preferences – conditional on some common-knowledge applicant characteristics (denoted by $X \equiv [x_i]_{i \in \mathcal{I}}$) – is common knowledge. For example, X includes information previously announced by programs.

After applications are submitted in the form of ROLs, more information about applicants, denoted by Z , can become available to programs. In our experiment, applicants’ final grades are realized and learned by programs after applicants have submitted their ROLs.² With information on (X, Z) , a program always behaves truthfully at the “screening” stage, in which it screens all its candidates who, as defined above, are the applicants including it in their ROLs. Each program then submits its preferences over its candidates.

Finally, at the “matching” stage, the mechanism finds a matching by considering submissions from both sides.

Strategies and equilibrium. The following assumption is imposed on programs’ behavior. As we shall see in Section 3, this assumption is plausible in our experiment.

Assumption 1 *Programs do not behave strategically in the following sense: (i) Each program truthfully reports its capacity and prerequisites, (ii) screens all its candidates, and (iii) reports to the mechanism its true preference ranking over all its candidates.*

To mitigate the issue of multiple equilibria, we also impose several assumptions on applicant strategy as an “equilibrium refinement.”

Assumption 2 *Applicants do not apply to any unacceptable program or to any program for which they are not qualified.*

²It is important however to ensure that when they screen candidates, programs are not informed of actual ROLs or any signal thereof.

Because applicants always have access to their outside options, the marginal benefit of applying to unacceptable programs is zero; similarly, applying to a program for which an applicant does not qualify or meet the prerequisites brings zero payoff because the program will never accept an unqualified applicant. Assumption 2 therefore regulates how applicants behave when they have multiple best responses and is binding only when the marginal application cost is zero. With this assumption, we will possibly underestimate congestion in markets without application cost, because applicants may apply to unacceptable programs in real life (see, e.g., He, 2015). As a result, our estimated effect of application costs on congestion is likely to be a lower bound.

We focus on symmetric pure strategies such that $\sigma(v_i, x_i, X) : \mathbb{R}^{J+1} \times \mathbb{R}^{|x_i|} \times \mathbb{R}^{|x_i|} \rightarrow \mathcal{L}_i$, where x_i is i 's characteristics/information that are common knowledge and $\mathcal{L}_i \subseteq \mathcal{L}$ are all the ROLs consistent with Assumption 2. When playing the game, an applicant's beliefs are equated with her probability of acceptance by each program and depend on her information set. The probability that applicant i is accepted by program j when submitting list L_i is denoted by $\pi_j(L_i|X, \sigma)$, which is conditional on common knowledge X and on others' strategy σ . A Bayesian Nash equilibrium is σ^* such that for all (v_i, x_i, X) ,

$$\sigma^*(v_i, x_i, X) \in \arg \max_{\sigma(v_i, x_i, X) \in \mathcal{L}_i} \sum_{j \in \mathcal{J}} \pi_j(\sigma(v_i, x_i, X)|X, \sigma^*) \max(v_{i,j}, v_{i,0}) - C_A(|\sigma(v_i, x_i, X)|),$$

where $\pi_j(\cdot|X, \sigma^*)$ is the probability of acceptance by program j consistent with σ^* . In other words, given that everyone plays strategy σ^* , when i submits a list L , the probability that i is accepted by j is exactly $\pi_j(L|X, \sigma^*)$. Furthermore, only $\max(v_{i,j}, v_{i,0})$ matters in utility terms because i can always take the outside option whenever matched with an unacceptable program. The existence of a Bayesian Nash equilibrium in pure strategies can be established by applying Theorem 4 (Purification Theorem) in Milgrom and Weber (1985). Moreover, we can solve for equilibrium, and thus $\pi_j(\cdot|X, \sigma^*)$, following the steps in Appendix D.1.

Before discussing equilibrium properties, we introduce the following assumption.

Assumption 3 *Equilibrium acceptance probabilities are non-degenerate for qualified applicants: $\pi_j(L|X, \sigma^*) \in (0, 1)$ for every $j \in L$ and for all i and $L \in \mathcal{L}_i$.*

In other words, there is sufficient uncertainty in other applicants' preferences and in program preferences such that every applicant has some chance of being accepted by any program to which she applies as long as she meets the prerequisites. This assumption is satisfied in our econometric model because we assume that both applicant and program preferences (v_i, s_i) have full support on the real line (see Section 5).

Given Assumptions 1, 2, and 3, the literature provides the following results:

- (i) *Under DA*, i.e., $C_A(|L|) = 0$ for all L , there is a unique Bayesian Nash equilibrium in which every applicant truthfully ranks all of the acceptable programs for which she qualifies (Dubins and Freedman, 1981; Roth, 1982; Fack et al., Forthcoming). The equilibrium outcome is the applicant-optimal stable matching that Pareto dominates all other stable matchings in terms of applicant welfare (Gale and Shapley, 1962).
- (ii) *Under DA-T-K for $K \in \{1, \dots, J\}$* , it is a dominated strategy if one submits ROL L that does not truthfully rank programs included in L (Haeringer and Klijn, 2009). An equilibrium matching outcome is not necessarily stable; for any applicant-program pair (i, j) , the probability of (i, j) forming a blocking pair is bounded above by a function that is decreasing in K and equal to zero when $K = J$ (Fack et al., Forthcoming).
- (iii) *Under DA-C*, it is a dominated strategy if one submits ROL L that does not truthfully rank programs included in L . An equilibrium matching outcome is not necessarily stable; in the case of a constant marginal cost (c), for any applicant-program pair (i, j) , the probability that (i, j) forms a blocking pair is bounded above by a function that is increasing in c and equal to zero when $c = 0$ (Fack et al., Forthcoming).

We use these results to guide our experimental and research design and, in particular, the mechanisms considered in counterfactual analysis.

2.3 Evaluation Criteria and Pigouvian Tax

Our analysis evaluates the performance of a market design with five criteria. The first is about congestion or screening cost, while the last four are about match quality.

- (i) **Congestion/screening costs:** Recall that by assumption, every program screens all its candidates. In the absence of direct information on screening technology, we measure a program's screen costs by the number of applications it receives (and thus screens).
- (ii) **Stability (number of blocking pairs):** In two-sided matching, both sides' welfare matter. Stability is therefore a natural welfare criterion.³ Although finding all blocking pairs unrealistically requires that every applicant and program know everyone's preferences, we use this concept in the sense that no blocking pair implies Pareto efficiency when both sides have strict preferences. Moreover, blocking pairs may affect

³In contrast, the recent literature on school choice is mostly interested in the welfare of one side, students (see, for a survey, Pathak, 2011). In that setting, schools are passive and only endowed with priorities instead of preferences.

the survival of a market design in the future, because such pairs can sometimes be identified and thus reveal some drawbacks of the existing design.

- (iii) **Number of unmatched applicants:** An extreme form of match inefficiency is some applicants being unmatched, especially when everyone can be matched. Indeed, our setting is such a matching market, as we shall see shortly. The number of unmatched applicants therefore provides a measure of match quality.
- (iv) **Applicant welfare:** We evaluate applicant welfare both ordinally and cardinally. The ordinal measure uses the applicants' true ordinal preferences revealed under DA, while the cardinal measure relies on an estimation of cardinal preferences. An advantage of the cardinal measure is that we can explicitly consider application costs and measure applicant welfare net of costs.
- (v) **Program welfare:** Measuring program welfare over their assigned applicants is more demanding. In the data, we observe program ordinal preferences over individual applicants; even with this information and assuming responsive program preferences, we may still fail to construct program preferences over sets of applicants. We report program ordinal preferences whenever the comparison is conclusive. In addition, we construct two approximations to cardinal preferences – one based on applicants' grades, the other using program ordinal preferences.

Externalities and Pigouvian Tax The above welfare criteria are inputs to the social welfare function which shows externalities and possible Pigouvian taxes.

Suppose that applicant i submits L_i in equilibrium. As a simple illustration, let us consider an additive social welfare function as follows:

$$\begin{aligned}
 & Total_Social_Welfare(L_1, \dots, L_I) \\
 &= \sum_{i=1}^I [Applicant_i's_Welfare(L_1, \dots, L_I) - Applicant_i's_Application_Costs(L_i)] \\
 &+ Total_Program_Welfare(L_1, \dots, L_I) - Total_Screening_Costs(L_1, \dots, L_I)
 \end{aligned}$$

where $Applicant_i's_Welfare(L_1, \dots, L_I)$ is the welfare i enjoys from her match when the equilibrium action profile is (L_1, \dots, L_I) ; $Applicant_i's_Application_Costs(L_i)$ is i 's application cost when submitting L_i and does not depend on other applicants' behavior; the total program welfare from match quality and total screening costs are $Total_Program_Welfare(L_1, \dots, L_I)$ and $Total_Screening_Costs(L_1, \dots, L_I)$, respectively.

Without loss of generality, let us focus on applicant 1 and consider L'_1 , such that L'_1 drops one program from L_1 and thus $|L_1| - |L'_1| = 1$. We consider the effects of i applying to one more program and switches the profile from (L'_1, L_2, \dots, L_I) to (L_1, L_2, \dots, L_I) :

$$\begin{aligned}
& \text{Total_Social_Welfare}(L_1, L_2, \dots, L_I) - \text{Total_Social_Welfare}(L'_1, L_2, \dots, L_I) \\
&= \sum_{i=1}^I [\text{Applicant}_i\text{'s_Welfare}(L_1, L_2, \dots, L_I) - \text{Applicant}_i\text{'s_Welfare}(L'_1, L_2, \dots, L_I)] \\
&\quad - [\text{Applicant}_1\text{'s_Application_Costs}(L_1) - \text{Applicant}_1\text{'s_Application_Costs}(L'_1)] \\
&\quad + \text{Total_Program_Welfare}(L_1, L_2, \dots, L_I) - \text{Total_Program_Welfare}(L'_1, L_2, \dots, L_I) \\
&\quad - [\text{Total_Screening_Costs}(L_1, L_2, \dots, L_I) - \text{Total_Screening_Costs}(L'_1, L_2, \dots, L_I)].
\end{aligned}$$

As L_1 is optimal for applicant 1, the private benefit to applicant 1 when switching from L'_1 to L_1 is non-negative:

$$\begin{aligned}
& \text{Applicant}_1\text{'s_Welfare}(L_1, L_2, \dots, L_I) - \text{Applicant}_1\text{'s_Welfare}(L'_1, L_2, \dots, L_I) \\
&\quad - [\text{Applicant}_1\text{'s_Application_Costs}(L_1) - \text{Applicant}_1\text{'s_Application_Costs}(L'_1)] \geq 0. \quad (1)
\end{aligned}$$

Importantly, programs' total screening costs will be higher, leading to a potential reduction in social welfare. That is,

$$-[\text{Total_Screening_Costs}(L_1, L_2, \dots, L_I) - \text{Total_Screening_Costs}(L'_1, L_2, \dots, L_I)] < 0. \quad (2)$$

The increase in the total screening costs (Equation 2) can dominate the increase in applicant 1's welfare (Equation 1). This negative effect can further outweigh the total changes in the match quality of programs and other applicants. Hence, the total social welfare may decrease, which captures the externalities that are associated with applicant 1's behavior. A Pigouvian tax can be implemented through application costs such that the private decision of applying to more programs will take into account the associated externalities.

3 Experiment at the Toulouse School of Economics

3.1 Background and Experimental Design

TSE organizes its master's programs into two years of study, M1 and M2. In the first year, it admits approximately 150 students, who are placed into three M1 programs: Law and

economics, statistics and econometrics, and economics. M1 students are allowed to apply to the seven M2 programs for their second year of study. The names of the programs, which indicate their differentiated foci, are described in Table 1. In the rest of the paper, the programs are randomly ordered and labeled as P1 to P7. Our study focuses on the matching between applicants and programs.

Table 1: Program Names and Prerequisites

Program Acronym	Prerequisites	Program Name in English (<i>French</i>)
ECL	Yes	Economics and Competition Law (<i>Economie et Droit de la Concurrence</i>)
EMO	No	Economics of Markets and Organizations (<i>Economie des Marchés et des Organisations</i>)
ERNA	No	Environmental and Natural Resources Economics (<i>Economie de l'environnement et des ressources naturelles</i>)
ETE	No	Economic Theory and Econometrics – PhD Track (<i>Economie Mathématique et Econométrie</i>)
MIF	No	Financial Markets and Intermediaries (<i>Marchés et Intermédiaires Financiers</i>)
PPD	No	Development Economics and Public Policies (<i>Politique Publique et Développement</i>)
STA	Yes	Statistics and Econometrics (<i>Statistique et Econométrie</i>)

Notes: The programs are listed in alphabetical order, according to their acronyms. This order does not correspond to P1-P7. The prerequisites for Program STA include some courses in statistics and econometrics; those for Program ECL include certain courses in law.

In partnership with the TSE administration, we conducted an experiment in May 2013 aimed at improving the market design for the applicant-program match. Previously, for example, in 2011 and 2012, the match was organized in a semi-centralized but rather ad-hoc fashion. Applicants submitted two ranked choices without being provided explicit guidelines on how these two choices would be used. Given applicants' choices and academic files, the program directors met to decide who could be matched with which program. During the meeting (and sometimes multiple meetings), each program director screened its candidates and decided whether to accept them. Because every applicant was guaranteed the opportunity to continue his/her study at TSE, the costliest part of this process was finding solutions for applicants who were ranked low and were rejected by their two choices. As its main disadvantages, this market design was time-consuming for the program directors and even created adversarial relationships among some of them.

In our experiment, the “subjects” come from two sides of the matching market in 2013: The seven M2 programs, which are represented by their directors, and the 129 applicants who were finishing their M1 study at TSE and were applying for admission to TSE's M2 programs. As depicted in Figure 2, the matching game in the experiment is played in five

stages:

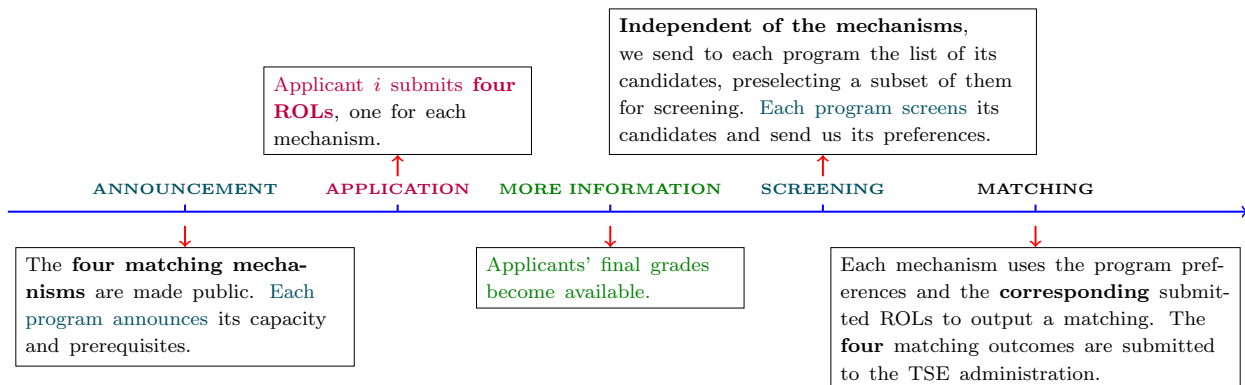


Figure 2: Timeline of Matching Applicants with M2 Programs at TSE in the Experiment

- (i) **Announcement:** The programs announce their capacity and prerequisites (in terms of courses that an applicant must have taken during M1 study). Only two programs have such prerequisites. The programs also disclose that an applicant’s final grade – a weighted average of all courses taken during the academic year – will be used in the screening process. Applicants are informed that one of the three aforementioned mechanisms (DA, DA-T, and DA-C) and the Immediate Acceptance (IA) will be chosen to match applicants with programs.⁴ The version of DA-T is DA-T-4; under the designed DA-C, applicants can freely rank their top-3 choices and must write a motivation letter for each additional choice.
- (ii) **Application:** *Under each mechanism*, applicants submit an ROL on an official university website that is also used for all coursework.
- (iii) **More Information.** After applicants submit their ROLs, they take their final exams and obtain their final grades, which are weighted averages of all courses taken during the academic year.
- (iv) **Screening.** To each program, we send a list of the applicants who have included the program in at least one of her four submitted ROLs. We also attach the information on their final grades and their grades for each individual course, but we do not inform the programs about how each applicant ranks them or the motivation letters that may have been written under DA-C. To save the programs screening costs and convince them that engaging in the experiment is not too costly for them, we, as the market designer and the clearinghouse, pre-select a subset of the program’s candidates for

⁴The Immediate Acceptance mechanism, also known as the Boston mechanism, is sometimes used for school choice in practice. The definition of the mechanism is given in Appendix A.1. Although also being included in the experiment that we conduct, this mechanism is not the focus of the current study. We therefore do not present the results of this mechanism.

screening. Meanwhile, the program directors are explicitly encouraged to screen all applicants. The programs then rank applicants and send us their rankings.

- (v) **Matching.** After receiving the programs’ rankings over applicants, we calculate the matching under each of the four mechanisms and send the outcomes to the administration. The TSE administration chooses one of the four matchings to be implemented (which happens to be the one from DA).

The above information and the instructions of the experiment are all explained on the application website, and we provide screenshots in Appendix A.

“Sub-game perfect equilibrium.” It is announced that one of the mechanisms will be chosen for the final matching, so the applicants should play a “sub-game perfect equilibrium.”⁵ That is, under each mechanism, it is in applicants’ best interest to behave optimally as if that mechanism were actually implemented. On the website, we also make this point explicit (see the screenshots in Figures E.1–E.5 in Appendix E).

One may be concerned that, under DA-C, applicants must pay the cost before knowing if the mechanism will be chosen. In other words, the return to writing a motivation letter is discounted by the probability that DA-C is chosen, which amounts to inflating the application cost. In any case, the application costs under DA-C remain subjective.

Information on the mechanisms. We provide applicants with definitions of the mechanisms and explicit tips on how to play the game under each mechanism. Under DA, we emphasize that truth-telling, i.e., ranking programs according to their true preferences, is a dominant strategy. Under DA-T-4 and DA-C, it is noted that while truth-telling is no longer a dominant strategy, it is still in the applicants’ best interest to order the ranked programs truthfully (Haeringer and Klijn, 2009). The screenshots in Appendix E (Figures E.1–E.5) show what applicants could see on the website.

Information to the program directors. When screening applicants, the directors are not informed about the submitted ROLs or any signals thereof or which mechanism will be

⁵When implementing the experiment, no one in any party involved in the experiment (including the experimenters, the TSE administration, the program directors, and the applicants) had a clear prediction on which mechanism would be chosen. Although the TSE administration had the final say, the choice could be influenced by various factors. In particular, the program directors, with opposing objectives, could lobby or protest. Besides, the administration would also be choosing a mechanism for the future. Our understanding of the objectives of the TSE administration is that both match quality and screening cost were important. These *ex-ante* considerations convinced us to state the following in the experiment instructions “*to the best of our knowledge, we believe that every mechanism has an almost equal probability to be chosen.*” See Figures E.1 in Appendix E for a screenshot of the instructions.

chosen. This makes it more plausible that they behave truthfully and that Assumption 1 is satisfied.

Screening cost depends on whether the programs can screen applicants round by round, and crucially, this possibility relies on how programs can use information about applicants' submitted ROLs. For various reasons, some programs may prefer to take into account the applicants' preferences or their ROLs when screening them. This would, however, create incentive problems for the applicants. To preserve the strategy-proofness of DA, it is usually required that applicants' ROLs are not seen by the programs. For example, the National Resident Matching Program (NRMP) – the clearinghouse that matches medical residents with hospitals – states that “[y]our rank order list is confidential and never will be shared with the programs.” One way to keep ROLs confidential is to ask the programs to screen the applicants all at once before running the mechanisms, which is the practice in, for example, university admissions in Germany and Victoria, Australia and the NRMP, where certain variants of DA are used.

Pre-selecting candidates to be screened. To reduce screening costs, we use applicants' submitted ROLs to identify a subset of applicants over whom the program's strict preferences are required to calculate the matching.⁶ This identification is achieved by running each of the four mechanisms with a series of approximations of programs' preferences – first by sorting applicants by their average M1 grades and then by perturbing the ranking numerous times by adding independent noises. In this process, we also take into account whether an applicant meets the prerequisites of a program, if there are any. The variance of noises is significant, leading to a large enough subset of applicants to be screened by each program.

The number of pre-selected applicants is program-specific and ranges from 27 to 52, with an average of 41; Appendix A.3 provides more details. When the mechanisms were actually run, no program had to accept or reject an applicant who was not pre-selected for that program. That is, there was no need for the programs to screen more applicants than the pre-selected ones. Additionally, no program screened more candidates beyond the pre-selected ones, even though they were encouraged to do so. In our estimation and counterfactual analyses, we sometimes need the complete preference ranking of each program over all applicants. For this, we expand the ranking over the pre-selected applicants to all applicants following the procedure explained in Appendix B.

⁶This pre-selection resembles the service provided by recruiting firms to employers. Although it may have potentials in reducing congestion, this practice is beyond the scope of our study.

It should be emphasized that the applicants are *not informed* of this pre-selection at any stage. It is therefore impossible that their behavior is affected by this practice. Moreover, every program is encouraged to screen everyone, and the properties of each market design presented in Section 2.2 are thus invariant to pre-selection.

Randomized orders of programs and mechanisms. To prevent potential framing effects, we randomly assign the 129 applicants into 7 groups, each of which has a unique order of programs presented on the website. Each of the 7 programs is presented as the first to a unique group of applicants and presented as the last to another. Moreover, applicants in a given group play the four mechanisms in one of the following four sequences (DA, DA-T-4, DA-C, IA), (DA, DA-C, DA-T-4, IA), (IA, DA, DA-T-4, DA-C), and (IA, DA, DA-C, DA-T-4). We decide to have DA-T-4 and DA-C immediately follow DA in order to avoid potential confusion about the mechanisms.

3.2 Length of ROL and Unacceptable Programs

In the experiment, applicants are required to rank all 7 programs under DA, 4 programs under DA-T, and at least 3 programs under DA-C, as the web design makes it cumbersome to allow applicants to rank a flexible number of programs. To alleviate this constraint, we also elicit information from every applicant on whether a program is acceptable to him/her. When evaluating the programs' screening costs, we assume that applicants never apply to unacceptable programs (Assumption 2).

Before beginning the application process, every applicant must answer the following question for each program:

*“If you are accepted (and only accepted) by **program name**, will you stay at TSE and register for the program in September 2013?”*

In the survey, “**program name**” is replaced by the full name of one of the seven programs. An applicant must tick one of the three possible answers, “Yes,” “No,” and “I Don’t Know.” We relabel these responses as “definitely acceptable,” “unacceptable,” and “possibly acceptable”, respectively, whereby a possibly acceptable program is less preferred than a definitely acceptable one but more preferred than an unacceptable one. The same question is repeated for all seven programs, and we clarify that this information is not used in the *actual* matching process, although we use it in this research. In the following, we sometimes refer to the “definitely acceptable” and “possibly acceptable” categories collectively as “acceptable”.

The responses to these questions are summarized in Table 2. Two findings stand out from the table, which shows the heterogeneity of applicant preferences: Program P4 is definitely acceptable to 81% of the applicants and unacceptable to only 7% of them; in contrast, program P6, which is more specialized and has prerequisites, is unacceptable to the majority of applicants (78%).

Table 2: How the Applicants Consider the Programs: Percentages

Program	Unacceptable	Possibly Acceptable	Definitely Acceptable
P1	22 (3.63)	10 (2.62)	67 (4.06)
P2	29 (3.98)	19 (3.43)	53 (4.39)
P3	38 (4.30)	18 (3.34)	44 (4.40)
P4	7 (2.23)	12 (2.89)	81 (3.46)
P5	33 (4.15)	24 (3.74)	43 (4.38)
P6	78 (3.69)	7 (2.26)	16 (3.17)
P7	40 (4.27)	18 (3.34)	43 (4.30)

Notes: This table shows the percentages of 129 applicants who consider each program as “unacceptable,” “possibly acceptable,” and “definitely acceptable”. The results are calculated from the applicants’ responses to survey questions that are not used in the actual matching process. Bootstrap standard errors from 10,000 samples are in parentheses.

When assuming that applicants do not apply to unacceptable programs, we notice that one applicant is matched with an unacceptable program under all mechanisms (and is enrolled in this program in September 2013). Since dropping the applicant’s unacceptable programs from his/her ROLs would change the matching outcome under every mechanism, we instead re-categorize the program to be “possibly acceptable” to this applicant. Furthermore, there are 9 applicants whose ROLs under DA are not consistent with the survey responses, provided that they report true preferences under DA. For example, one ranks an unacceptable program before a possibly acceptable one. For these applicants, we update some of the programs’ acceptability to restore consistency. More details on this data cleaning are available in Appendix B.

4 Experimental Data and Direct Evaluation

In addition to observing applicant behaviors in the experiment, we also collect administrative data on applicants’ grades from M1 courses (from the first semester and from both

semesters), demographic information, and scholarship status. The key variables are summarized in Table 3.

Table 3: Summary Statistics

Variable	Mean	S.D.	Min	Max
M1 Grades: First semester	11.66	2.21	6.01	17.40
M1 Grades: Final (2 semesters)	12.27	2.04	8.02	17.68
Age	24.47	1.73	20.96	33.57
M1 Program: Economics	0.57	0.50	0	1
M1 Program: Statistics	0.34	0.48	0	1
Female	0.45	0.50	0	1
Scholarship	0.28	0.45	0	1

Notes: In total, there are 129 observations. Scholarship equals one if the applicant holds a scholarship from TSE or from the French government. Out of a total of 20 points, the grades are the credit-weighted average of grades from the courses taken during their first-semester or two-semester M1 study.

4.1 Direct Evaluation

We now turn to applicant strategies and outcomes under each mechanism. As discussed in Section 2.3, we focus on two types of measures, match quality and screening cost. The first type includes stability, unmatched applicants, applicant welfare, and program welfare. Screening cost, or congestion, is measured by the number of candidates to be screened.

Under the assumption that applicants report truthfully under DA, we take applicants' reported ROLs under DA as their true ordinal preferences. There is a recent literature documenting that some applicants do not report truthfully under DA (e.g., Hassidim, Romm and Shorrer, 2016; Artemov, Che and He, 2017; Shorrer and Sóvágó, 2017). The truth-telling assumption is plausible in our setting for the following reasons: (i) On the experiment website, we explicitly announce that truth-telling under DA is in their best interest (see the screen shots in Figure E.3), and provide the details of the algorithm as well as an example on a linked help page; together with the fact that our applicants are graduate students in economics and have a few weeks to make a decision, it is plausible that the applicants understand/believe the strategy-proofness of DA. (ii) In contrast to other settings, our applicants do rank all programs. (iii) Comparing their submitted ROLs under DA with the survey responses on program acceptability, we identify only 9 instances of inconsistency, among which 8 involves the response "I don't know" (see Appendix B).

In this section, we do not make further assumptions on applicants' strategies under DA-T-4 or DA-C and instead evaluate what is observed. Recall that we assume that applicants do not apply to the programs that are deemed unacceptable or for which they do not meet the prerequisites. Therefore, we remove these programs from the ROLs when analyzing

congestion. This removal does not affect the matching outcome in the data and therefore mitigates market congestion only, especially for DA (without cost).

4.1.1 Statistical Inferences with “One Observation”

The focus of the study is the matching game consisting of agents on both sides, applicants and programs. Therefore, in an ideal experiment, we would have the same game independently played multiple times by different sets of agents under different mechanisms. In other words, the unit of observation in the experiment is a game play or a matching market. However, our data contain outcomes from only one market under several mechanisms. Thus, careful consideration is necessary when making statistical inferences.

This leads us to resampling (or bootstrapping): We randomly resample applicants with replacement; the size of the resample is equal to the size of the original dataset; and in each resample, the programs are always kept the same. Implicitly, we assume that the empirical distribution of applications (types and strategies) is a good approximation of the theoretical distribution; by drawing from the empirical distribution independently, we create an independent play of the same game in each resample. In a given resample, multiple observations may originate from the same applicant in the original data; when programs rank the applicants, the ties among them are broken randomly.⁷ Standard errors and testing results in this section are based on 10,000 resamples.

4.1.2 Applicant Strategies

Table 4 presents the distribution of the number of programs ranked under each mechanism. By imposing a cost, both DA-T-4 and DA-C discourage the applicants from ranking too many programs. As a result, every applicant ranks 0.7 (0.9) programs fewer on average under DA-T-4 (DA-C) than under DA. The differences between these numbers of applications are all statistically significant (p-value < 0.01), as is the difference between DA-T-4 and DA-C (p-value < 0.05).

DA-T-4 allows applicants to freely rank four programs, and therefore, the mode of the number of ranked programs is four, accounting for 57 percent of the applicants. Moreover, 22 percent of them do not adopt a “truncation” strategy; that is, these applicants’ submitted

⁷We use a lottery to break ties for each program. Because ties involve the same applicant being sampled multiple times in a resample, how ties are broken does not affect our outcome. Relatedly, because some variables, X , are assumed to be common knowledge among all applicants, re-sampling should be conditional on X . As this would deflate standard errors and potentially lead to small sample issues, we present the results from random resampling which would tend to be conservative.

Table 4: Number of Programs Ranked by Applicants (in percentages)

Mechanism	# of programs ranked							Average number of programs ranked ^a
	1	2	3	4	5	6	7	
DA	4 (1.69)	10 (2.62)	17 (3.31)	26 (3.87)	28 (3.95)	15 (3.14)	0 -	4.09 (0.12)
DA-T-4	4 (1.69)	11 (2.70)	29 (3.98)	57 (4.34)	- -	- -	- -	3.38 (0.07)
DA-C	5 (1.84)	14 (3.02)	56 (4.42)	12 (2.87)	9 (2.54)	5 (1.85)	0 -	3.21 (0.10)

Notes: This table shows the distribution of the numbers of programs in the submitted ROLs under each mechanism. Each number – except those in the last column – represents the percentages of the applicants who rank that many programs in their ROLs. We remove from the ROLs the unacceptable programs and those for which the applicant is not qualified. Bootstrap standard errors from 10,000 samples are in parentheses. ^a From t -tests, the differences in the numbers of applications under DA, DA-T-4, and DA-C are all statistically significant (p -value < 0.05). Specifically, the difference is 0.71 (s.e. 0.07) between DA and DA-T-4, 0.88 (s.e. 0.09) between DA and DA-C, and 0.17 (s.e. 0.07) between DA-T-4 and DA-C.

ROLs do not coincide with the top portion of their true ordinal preferences. In other words, many applicants “skip” a program in their application. Under DA-C, this statistic is 7 percent.

Recall that it is free to rank the first three programs under DA-C, while the applicant has to pay a cost by writing a motivation letter for each of the fourth to seventh choices. Unsurprisingly, the mode is three; however, 26 percent of the applicants choose to pay some costs and rank more than three programs.

4.1.3 Match Quality

Applicants’ submitted ROLs and programs’ rankings over applicants allow us to calculate the matching outcome under each mechanism. Table 5 reports the number of matched applicants by program and mechanism. Programs P5-P7 do not meet their capacities mainly because the total capacity (142) exceeds the total number of applicants (129). Only under DA-C is an applicant unmatched.

In terms of applicant welfare, 85 percent of the applicants are assigned to their most preferred program, as shown in Table 6. DA and DA-T-4 perform equally well on this dimension, while DA-C performs worse due to the applicant left unmatched with this mechanism. The ordinal welfare distribution barely varies across the 10,000 resamples under either mechanism. These statistics imply that applicant preferences are sufficiently heterogeneous. One would expect that in such a market, congestion should be less of a concern. Yet, as we shall see in Section 4.1.4, program screen costs can still be significantly lowered without harming match quality.

Table 7 further investigates individual matches. The left side of the table reports the

Table 5: Number of Matched Applicants by Program under Each Mechanism

Program	Capacity	#Matched Applicants under		
		DA	DA-T-4	DA-C
P1	14	14 (0.06)	14 (0.12)	14 (0.06)
P2	22	22 (1.30)	22 (1.28)	22 (1.48)
P3	22	22 (2.30)	22 (2.01)	22 (2.00)
P4	28	28 (0.02)	28 (0.02)	28 (0.02)
P5	22	12 (3.08)	12 (3.21)	11 (3.09)
P6	12	10 (2.31)	10 (2.31)	10 (2.31)
P7	22	21 (2.07)	21 (2.15)	21 (2.35)
Total Unmatched		0 (1.27)	0 (1.90)	1 (1.99)

Notes: This table shows the number of applicants matched with each program under each mechanism. The total capacity is 142. Three programs, P5-P7, do not meet their capacities. Bootstrap standard errors from 10,000 samples are in parentheses.

Table 6: Ordinal Welfare of the Applicants under Each Mechanism

	Fraction of Applicants Matched with				Unmatched
	Most Preferred	2nd Preferred	3rd Preferred	4th Preferred	
DA	85	10	4	1	0
DA-T-4	85	10	4	1	0
DA-C	85	10	4	0	1

Notes: This table shows the numbers of applicants matched with their most preferred program, second most preferred one, etc., under each mechanism. The applicants' true ordinal preferences are what they reveal under DA. No applicant is matched with his/her 5th or less preferred program. Bootstrap standard errors from 10,000 samples are not reported, but are all less than 0.05.

frequency of blocking pairs under the two mechanisms. DA-T-4 leads to one blocking pair only, and DA-C results in two blocking pairs formed by two applicants and two programs.

Furthermore, we investigate individual welfare across mechanisms. The middle part of Table 7 compares the matchings from DA-T-4 and DA-C to that from DA. DA gives us the applicant-optimal stable matching, which also happens to coincide with the program-optimal matching in our data. DA-T-4 and DA-C lead to 3 and 4 applicants having different outcomes, respectively. Moreover, the number of applicants who are worse off is greater than the number of those who are better off.

The right side of Table 7 shows the number of programs that have different matching outcomes under DA-T-4/DA-C relative to DA. 3 (4) programs' outcomes are affected by DA-T-4 (DA-C), and more programs are worse off than better off under DA-T-4/DA-C.⁸

⁸Given two matchings, μ and μ' , program j is better off if the newly matched applicants, $\mu'^{-1}(j) \setminus (\mu^{-1}(j) \cap \mu'^{-1}(j))$, dominate the displaced applicants, $\mu^{-1}(j) \setminus (\mu^{-1}(j) \cap \mu'^{-1}(j))$, i.e., element-wise $\mu'^{-1}(j) \setminus (\mu^{-1}(j) \cap \mu'^{-1}(j))$.

Table 7: Deviation from the Optimal Stable Matching under Each Mechanism

Mechanism	Blocking Pairs			App. w/ Diff. Match			Prog. w/ Diff. Match		
	#pairs	#app.	#prog.	Total	Worse off	Better off	Total	Worse off ^a	Better off ^a
DA-T-4	1 (1.87)	1 (1.73)	1 (0.85)	3 (3.36)	2 (2.24)	1 (1.45)	3 (1.74)	2	1
DA-C	2 (2.64)	2 (1.92)	2 (1.07)	4 (3.41)	3 (2.32)	1 (1.46)	4 (1.62)	3	1

Notes: This table shows how the matching outcome under each mechanism is different from the applicant-optimal stable matching (i.e., the DA outcome). There are 129 applicants in total. Blocking pairs are defined with respect to applicants' true preferences revealed under DA. Bootstrap standard errors from 10,000 samples are in parentheses. ^a Standard errors are not calculated for these two statistics because a program's welfare change cannot always be labelled as better off or worse off in every bootstrap sample (although it is feasible in the experimental data). Recall that a program is better off if all the matched applicants are (weakly) better than those matched in the old matching.

In summary, both applicant welfare and program welfare differ across the mechanisms, but the magnitude is small or even negligible. However, as we emphasize, programs' screening cost is another important factor, and we now investigate the screening costs that programs must pay under each mechanism.

4.1.4 Programs' Screening Costs

Recall that each program screens its candidates all at once and that we make no assumptions about screening cost except that it strictly increases with the number of candidates to be screened (see Section 2 for more details). Given these assumptions, Table 8 shows, on average, how many candidates a program screens under each mechanism or, equivalently, how many applicants include a given program in their submitted ROLs under each mechanism. When the application cost increases from DA to DA-T-4 (DA-C), the number decreases from 75 to 62 (59), while the differences are all significant at the 2% level based on *t*-tests.

Table 8: Summary Statistics: Number of Candidates to Be Screened

	mean	s.d.	min	max
DA	75.29	37.39	11	120
DA-T-4	62.29	33.97	11	117
DA-C	59.14	31.14	11	107

Notes: Each row of this table shows the summary statistics of the number of candidates to be screened across the seven programs in the experimental data. In other words, these statistics are not based on resampling. Treating each program under a given mechanism as an observation (and not using resampling), *t*-tests show that the difference between any two mechanisms is significant at the 2% level.

We further normalize each program's screening cost by its capacity. Figure 3 shows the number of candidates per opening (i.e., per available seat) to be screened by each program under each mechanism. Again, the screening cost is negatively correlated with

$\mu'^{-1}(j) \succ_j \mu^{-1}(j) \setminus (\mu^{-1}(j) \cap \mu'^{-1}(j))$ when the matched applicants in these two sets are ordered according to *j*'s ordinal preferences, \succ_j .

the application cost. Under DA, a program has to screen 3.69 candidates per opening on average. DA-T-4 reduces the cost by approximately 0.65 candidate per opening; DA-C is even more effective. Moreover, when moving from DA-T-4 to DA-C, this reduction occurs for every program except P6.




Figure 3: Number of Candidates Screened by Each Program per Opening

Notes: This figure shows the total number of candidates screened by each program per opening under each mechanism. Programs must screen all their candidates at once before the mechanism is implemented. The average number of candidates screened across the seven programs is 3.69 under DA, 3.04 under DA-T-4, and 2.90 under DA-C. The error bars indicate the 90% confidence intervals from 10,000 bootstrap samples.

In conclusion, among the three mechanisms, the increases in application cost does not significantly affect match quality but greatly reduces programs' screening costs.

5 Evaluation Based on Structural Estimation

The set of market designs that we can directly evaluate is necessarily restricted because we implement only one instance of DA-T and one instance of DA-C in the experiment. In order to draw a complete picture of the trade-offs between screening costs and match quality, we would like to evaluate many other market designs with different parameters governing truncation and application costs. Building on the rich information on both programs and applicants we observe, this section implements a structural estimation and conducts evaluations using counterfactual analysis.

To achieve this goal, our structural estimation contains several stages. When DA is replaced by DA-T or DA-C, an applicant's optimal strategy is no longer truthful reporting, and instead, it depends on other applicants' actions. Recovering the cardinal preferences

of all agents becomes necessary. First, we use experimental data to estimate applicants' cardinal preferences. In particular, we combine applicants' submitted ROLs under DA with the survey data on program acceptability; we also make use of the fact that truthfully ranking acceptable, qualified programs is the unique equilibrium (Assumption 3). Recall that applicant ordinal preferences are truthfully revealed under DA (see the discussion in Section 4.1). Second, we estimate program preferences over applicants from programs' submitted rankings of applicants, given the assumption that programs truthfully report their preferences (Assumption 1). We next use data on final grades to estimate applicants' beliefs about their probabilities of acceptance by each program.

Optimal strategies of applicants depend on the solution concept that we retain. We specify the structure of incomplete information of the game that applicants are assumed to play, leading to use a Bayesian Nash equilibrium as defined in Section 2. That is, when submitting ROLs, we assume that each applicant knows her own preferences but not others', even though she is aware of their distribution. Moreover, in spite of knowing what criteria programs use to rank applicants, she is uncertain about how she is ranked by programs. This game may seem combinatorially complex, because an applicant can choose from a large set of ROLs. However, under DA-T and DA-C, optimal strategies are to drop some programs from the true preference order (Haeringer and Klijn, 2009), which makes it simpler to play.

Finally, we simulate the counterfactual equilibria attained when different versions of DA-T and DA-C are used.

The rest of this section formalizes the game of incomplete information. We first specify the following key elements: (i) The estimation of applicants' cardinal preferences, (ii) the estimation of program preferences, (iii) the information structure and solution concept, and (iv) applicants' beliefs about program preferences. For counterfactual analyses, we also explain the procedure for computing equilibria under each of these mechanisms.

5.1 Applicant Preferences

For each applicant, we observe a submitted ROL under DA, $(l_i^{*,1}, \dots, l_i^{*,J})$, which corresponds to her true ordinal preferences and from the survey, we observe whether each program is definitely acceptable, possibly acceptable, or unacceptable to the applicant. To reiterate, we do *not* use the data from DA-T-4 and DA-C, because the applicants may not play an equilibrium strategy in the experiment (see the discussion in Section 4.1).

Econometric model. Suppressing subscript i , we first postulate that the utility function associated with each program, $j \in \{1, \dots, J\}$, is:

$$v_j = x\beta_j + \varepsilon_j, \quad (3)$$

in which ε_j is extreme-value distributed and independent across programs. The parameters, β_j , are program-specific but not individual-specific, an assumption that we maintain throughout the paper. Furthermore, we model acceptability using an outside option whose value is the sum of two terms. The first term is observed at the moment of the application decision and is written as

$$v_0 = x\beta_0 + \varepsilon_0. \quad (4)$$

To model programs as “definitely acceptable”, “possibly acceptable”, or “unacceptable”, we posit that final acceptability is determined by adding to v_0 the value of another random variable, $\eta \in [x\beta_f + \underline{\eta}, x\beta_f + \bar{\eta}]$, which is revealed after the matching process.⁹ Program j is said to be **definitely acceptable** if and only if it has a value higher than the best possible outside option:

$$v_j > v_0 + x\beta_f + \eta \text{ for any } \eta,$$

that is,

$$v_j > v_0 + x\beta_f + \bar{\eta}.$$

Similarly, the program is deemed (definitely) **unacceptable** if and only if it is always dominated by the worst possible outside option:

$$v_j < v_0 + x\beta_f + \underline{\eta}.$$

Otherwise, the applicant considers the program **possibly acceptable** if

$$v_j \in [v_0 + x\beta_f + \underline{\eta}, v_0 + x\beta_f + \bar{\eta}]. \quad (5)$$

Let the ranking information be described as the ROL complemented with the max/min outside options (\bar{O}, \underline{O}) so that the observation is now an extended ROL:

$$(l^{*,1}, \dots, l^{*,\bar{J}}, \bar{O}, l^{*,\bar{J}+1}, \dots, \underline{O}, l^{*,\underline{J}}, \dots, l^{*,J}),$$

⁹This formalization to take into account the three types of acceptability differs from our theoretical framework in Section 2. However, as long as we categorize both “definitely acceptable” and “possibly acceptable” as acceptable, our theoretical results still hold true.

in which $l^{*,\bar{J}} \in \mathcal{J}$ for $1 \leq \bar{J} \leq J$ is the lowest-ranked definitely acceptable program. When $\bar{J} = 0$, there is no definitely acceptable program, and when $\bar{J} = J$, all programs are definitely acceptable. Similarly, $l^{*,\underline{J}} \in \mathcal{J}$ for $1 \leq \underline{J} \leq J$ is the highest-ranked unacceptable program so that $\underline{J} > \bar{J}$ and that $\underline{J} > J$ implies an absence of unacceptable programs.

Identification and likelihood function. A location normalization is needed, and the simplest one sets the lower bound of the outside option to zero:

$$(\beta_0 + \beta_f) = 0, \underline{\eta} = 0.$$

The upper bound, $\bar{\eta}$, is to be estimated. The scale normalization is given by the usual logit assumption. The parameter vector is thus described by $\theta = ((\beta_j)_{j=1,\dots,J}, \bar{\eta})$. If the survey information on program acceptability were not available, the choice probability of ROL (l^1, \dots, l^J) would be described by an exploded logit (Beggs, Cardell and Hausman, 1981):

$$\Pr(v_{l^1} = \max_{j \in \{1, \dots, J\}} v_j, v_{l^2} = \max_{j \in \{2, \dots, J\}} v_{l^j}, \dots, v_{l^{J-1}} = \max_{j \in \{J-1, J\}} v_{l^j}) = \prod_{j=1, \dots, J-1} \frac{\exp(x\beta_{l^j})}{\sum_{k=j}^J \exp(x\beta_{l^k})}.$$

When we introduce information on program acceptability, the likelihood function can be written as a sum of exploded logit terms (see Appendix C).

Results. We consider various sets of explanatory variables taken from administrative information on grades, M1 program (Economics, Econometrics and Statistics, Economics and Law), gender, age, and having a scholarship or not. There are many course grades that can be used at a fine level of detail. Because of the timing of the application decision, we opt to use the M1 first-semester grade as the main explanatory variable. At the time of application, no applicant knows his/her final grades, the weighted sum of his/her first- and second-semester grades.

The results are presented in Table 9. A clear pattern is that one's M1 program have some power to predict her preferences over the M2 programs, as the coefficients of "M1 Economics" and "M1 Statistics" have different signs across M2 programs.

5.2 Program Preferences

Program j sets a latent score for each applicant, $s_{i,j}$, according to which applicants are ranked. Conditional on meeting prerequisites $p_{i,j}$ (a binary variable), $s_{i,j}$ is assumed to

Table 9: Estimation of Applicant Preferences

	P1	P2	P3	P4	P5	P6	P7
Grade: First semester	0.28 (0.08)	0.10 (0.07)	0.00 (0.07)	0.10 (0.06)	0.00 (0.06)	-0.04 (0.08)	-0.01 (0.08)
M1 Program: Economics	0.58 (0.50)	-0.37 (0.49)	0.83 (0.98)	-0.12 (0.78)	-0.027 (0.55)	-6.65 (1.16)	0.016 (0.51)
M1 Program: Statistics	1.13 (0.51)	-0.71 (0.52)	2.73 (1.02)	-0.31 (0.81)	-0.02 (0.59)	-6.43 (1.18)	0.58 (0.54)
Female	-0.39 (0.31)	0.95 (0.30)	0.55 (0.30)	-0.54 (0.31)	0.672 (0.27)	0.51 (0.36)	-0.04 (0.32)
Scholarship	-0.57 (0.36)	-0.99 (0.37)	-0.024 (0.27)	-0.58 (0.35)	0.00 (0.27)	-0.32 (0.40)	-0.70 (0.33)
Intercept	1.11 (0.16)	0.71 (0.16)	0.377 (0.18)	1.88 (0.16)	0.55 (0.14)	-0.87 (0.19)	0.40 (0.16)
Upper bound of shocks to the outside option:			0.707	(0.07)			

Notes: Estimation is based on applicants' submitted ROLs under DA supplemented with survey information on program acceptability. The likelihood function is specified in Appendix C and can be considered as an extended version of the multinomial logit with rank-ordered data combined with acceptability. The McFadden pseudo-R-squared is equal to 0.078.

depend only on the final grade (the credit-weighted average of grades from two semesters), $FinalGrade_i$, and some noise,

$$s_{i,j} = FinalGrade_i + \sigma\xi_{i,j}, \quad (6)$$

in which $\xi_{i,j}$ is extreme value distributed. The term $\xi_{i,j}$ stands for the additional information that programs can use beyond final grades. $\xi_{i,j}$ is assumed to be independent of $FinalGrade_i$ and the covariates, x_i . Based on our robustness checks, including additional covariates leads to similar results.

Program prerequisites are common knowledge, and applicants condition their decision on $p_{i,j}$. Program preferences are therefore lexicographical and rank the subsample $p_{i,j} = 1$ first, the subsample $p_{i,j} = 0$ second, and according to $s_{i,j}$ within the first subsample.¹⁰

Given the structure of the experiment, we can estimate the parameters governing program preferences using a rank-ordered logit on the subsample of applicants satisfying the prerequisites, $p_{i,j} = 1$. Nonetheless, the rank-ordered logit procedure is fragile when many alternatives – or applicants – are considered, as we do here, in a small sample of seven programs. Therefore, we consider a limited information procedure using only length- K rankings among applicants with $K \geq 2$.

Specifically, let the full ranking of applicants for program j be $(i_1^{(j)}, \dots, i_I^{(j)})$, and partition this length- I vector into $I(K)$ vectors starting with $(i_{1,\cdot}, i_K)$, then $(i_{K,\cdot}, i_{2K-1})$, etc. We

¹⁰Because no one with $p_{i,j} = 0$ is accepted by program j , how we rank the applicants with $p_{i,j} = 0$ does not matter, as long as they are ranked after those with $p_{i,j} = 1$.

write the pseudo-likelihood function corresponding to these $J \times I(K)$ observations treated as ROLs while neglecting possible correlations across observations. This process does not affect the consistency of the estimates (Avery, Hansen and Hotz, 1983), although standard errors should be computed using a sandwich formula.¹¹

Results. The estimation of the latent score, equation (6), depends on how we choose K to form length- K ROLs. We present results for three different K 's, while the results from $K = 2$ are used in the counterfactual analysis. When $K = 2$, the estimates are as follows:

$$\tilde{s}_{i,j} = FinalGrade_i + \underset{(.022)}{0.058} \times \xi_{i,j}, \text{McFadden pseudo } R^2 = 0.337;$$

when $K = 3$:

$$\tilde{s}_{i,j} = FinalGrade_i + \underset{(.020)}{0.056} \times \xi_{i,j}, \text{McFadden pseudo } R^2 = 0.391;$$

when $K = 5$:

$$\tilde{s}_{i,j} = FinalGrade_i + \underset{(.015)}{0.049} \times \xi_{i,j}, \text{McFadden pseudo } R^2 = 0.496.$$

5.3 Solution Concept and Information Structure

We first clarify the solution concept and evaluation approach that are used in counterfactual analysis. Solution concepts depend on the information that each applicant has, as Panel A of Table 10 shows. A possibility is to assume complete information in our matching game. That is, both applicant and program preferences are common knowledge to every applicant, which leads us to Nash equilibrium as the solution concept. However, the complete-information assumption may be too restrictive in our setting because it is unlikely that every applicant knows everyone else's cardinal preferences or can predict her final grade and program preferences perfectly (cf. Figure 1). Indeed, for this reason, our model in Section 2 specifies a game of incomplete information, which leads to Bayesian Nash equilibrium.

In addition, multiple methods are available to evaluate counterfactual outcomes, depending on the use of information available to researchers. First, we always use common-knowledge information X . As Panel B of Table 10 shows, the outcome under any given market design can be evaluated with or without information on applicants' and programs'

¹¹An alternative is to consider sequences $(i_{1,\cdot}, i_K)$, $(i_{K+1,\cdot}, i_{2K})$, etc., in which independence may be satisfied but information is more limited.

Table 10: Solution Concepts and Evaluation Methods

<i>Panel A. Solution concepts</i>	
<u>Complete information?</u>	
A1. Yes (Nash equilibrium)	Applicant and program preferences are common knowledge. Every applicant best responds to others' actions in equilibrium.
A2. No (Bayesian Nash)	Applicant and program preferences are private information, but their distributions and X are common knowledge. Every applicant best responds to the distribution of others' actions in equilibrium, conditional on common knowledge.
<i>Panel B. Methods for Outcome Evaluation (conditional on common knowledge X)</i>	
<u>Conditional on observed applicant & program ordinal preferences?</u>	
B1. Yes (sample-specific)	An outcome is evaluated conditional on X and on the observed ordinal preferences of applicants and programs. That is, we consider only random draws of preference shocks consistent with the observed ordinal preferences. Each simulated applicant (with simulated cardinal preferences) plays an equilibrium strategy according to either Nash equilibrium (A1) or Bayesian Nash (A2).
B2. No (population-specific)	An outcome is evaluated conditional on X but unconditional on the observed ordinal preferences of applicants and programs. That is, we consider random draws of shocks from the assumed distribution. Each simulated applicant (with simulated cardinal preferences) plays an equilibrium strategy according to either Nash equilibrium (A1) or Bayesian Nash (A2).

observed ordinal preferences. If the information in the data on ordinal preferences is utilized, we have a “sample-specific” evaluation; otherwise, when this information is not used, the method is “population-specific”.

In addition to choosing Bayesian Nash equilibrium (A2) as our solution concept, we adopt the “sample-specific” evaluation (B1). This evaluation facilitates the comparison between our experimental results and the counterfactual analyses, as the former is obviously conditional on applicants' and programs' realized ordinal preferences.

To summarize, when solving for a Bayesian Nash equilibrium, we need every applicant to best respond to others with incomplete information on other applicants' and programs' preferences. The distributions of applicant and program preferences are modeled and estimated in the previous two subsections. However, programs form preferences with information on final grades that is not available to applicants at the time of application (cf. Figure 2). It is therefore necessary to consider the uncertainty in final grades when we estimate applicants' beliefs about program preferences, which is what we turn to next.

5.4 Applicant Beliefs on Program Preferences

When forming beliefs about program preferences, applicant i observes her own characteristics, x_i , including her grades during the first semester and the prerequisites of each program ($p_{i,j} = 1$ if i satisfies the prerequisites of program j and 0 otherwise). These prerequisites consist of choices of courses that are made well in advance and are partially determined by one’s M1 program. We therefore consider them as fixed in the statistical model.

We assume that applicants have rational expectations and expect that

- (i) The information structure detailed above is common knowledge and, in particular, program preferences are formed according to the model above.
- (ii) Applicant i predicts her final grade $FinalGrade_i$ as a function of characteristics x_i , using the following equation:

$$FinalGrade_i = x_i\gamma + \nu_i, \tag{7}$$

in which x_i includes the characteristics used in the modified exploded logit model above and also the first-semester grade in particular; ν_i is assumed to have a normal distribution.

The specifications of applicant and program preferences and the grade equation guarantee that any possible applicant preference ranking and program preference ranking among qualified applicants have a positive probability of occurrence. It is straightforward to show that equilibrium acceptance probabilities are non-degenerate for any qualified applicant, and therefore, Assumption 3 is satisfied.

Estimation. The rational-expectation assumption implies that applicant beliefs are consistent with the distribution of final grades and program preferences.

Equation (7) is estimated by standard ordinary least squares (OLS), and the results are shown in Table 11. Not surprisingly, the main determinant of the final grade is the first-semester grade. One may notice that the coefficient on “Scholarship” is significantly negative, which may be explained by the fact that scholarship is often need-based. For the residual ν_i in the grade equation, the assumption of normal distribution is reasonable, as shown by a quantile-quantile plot of the residuals, which is available upon request.

Based on these results, in the counterfactual analysis, we simulate applicant beliefs about program preferences by drawing ν_i from a normal distribution and $\xi_{i,j}$ from extreme value distributions. Furthermore, the prerequisites of P3 and P6 are held constant, because they

Table 11: Estimation of the Grade Equation

	Estimate	Std. Error	t Statistics	Pr(> t)
Grade: First semester	0.86	0.03	27.65	0.00
M1 Program: Economics	-0.42	0.25	-1.70	0.09
M1 Program: Statistics	-0.73	0.26	-2.77	0.02
Female	0.08	0.14	0.57	0.57
Scholarship	-0.37	0.15	-2.46	0.02
Intercept	2.83	0.42	6.77	0.00

Notes: Number of observations: 129; Residual standard error: 0.76 on 123 degrees of freedom; R-squared: 0.87; Adjusted R-squared: 0.86; F-statistic: 158.1 on 5 and 123 degrees of freedom, p-value: $< 2.2e - 16$.

are decided well in advance. As programs’ strategic behavior is assumed away, the above procedure describes the set-up for simulating applicant beliefs about program preferences.

5.5 Counterfactual Analysis of Market Designs

5.5.1 Simulating Equilibrium Outcomes: Procedures

We solve for a Bayesian Nash equilibrium (case A2 in Table 10) under each mechanism and then evaluate the sample-specific counterfactual outcomes (case B1 in Table 10).

Counterfactuals under DA-T and DA-C. We perform counterfactual analysis under DA-T with degrees of truncation varying from 1 to 6. Under DA-C, the application cost has the following form in the counterfactual analysis:

$$C(|L|) = c \times (|L| - 1) \times \mathbf{1}\{|L| > 1\},$$

in which c is the constant marginal cost for second and further choices and c can take one of the 11 values in $\{10^{-5}, 3.1 \times 10^{-5}, 10^{-4}, 3.1 \times 10^{-4}, \dots, 0.01, 0.31, 1\}$, increasing on a logarithmic scale. We have $C(|L|) = 0$ for $|L| \leq 1$ to avoid distorting the applicants’ participation decisions.¹² Although application costs are homogenous across applicants, applicant responses can be heterogenous. They depend on the marginal benefit of an application which is determined jointly by applicant preferences and acceptance probabilities.

To make sense of the magnitude of application costs, one may consider the following thought experiment: A “possibly acceptable” program can be improved to “definitely acceptable” by adding at most 0.707, which is the estimate of the maximum shock to outside option ($\bar{\eta}$ in equation 5); the same program can be demoted to “unacceptable” by subtract-

¹²This restriction can be relaxed in a more general setting. If it is costly to submit a one-choice ROL, some applicants may choose not to participate in the matching when their expected benefit is less than the application cost.

ing 0.707 at most. Moreover, the variance of utility shocks is $\pi^2/6 (\approx 1.645)$. From this perspective, most of the above costs, especially those below 0.031, are quite small.

Computation of equilibrium. For both DA-T and DA-C, an equilibrium must be solved as a fixed point in terms of cutoff distribution. A program’s *cutoff* in a given matching is the lowest rank among the applicants accepted by the program if it has no vacancy; otherwise, the cutoff is zero. A Bayesian Nash equilibrium can be summarized by the joint distribution of the cutoffs of programs (Azevedo and Leshno, 2016; Fack et al., Forthcoming). Given a particular draw of random shocks, an applicant ranked above a program’s cutoff can be accepted by this program. Solving for a Bayesian Nash equilibrium can then be achieved by finding a fixed point of the cutoff distribution.

We briefly outline the algorithm here, and further details are provided in Appendix D.1. The algorithm works with M (4000 in our implementation) simulated samples of size I (the number of applicants) in which the random shocks that contribute to the construction of applicant preferences ($\varepsilon_i^{(m)}$), to final grades ($\nu_i^{(m)}$), and to program preferences ($\xi_{i,j}^{(m)}$) are drawn from their respective distributions. These random shocks remain fixed in the whole procedure.

Over the M simulations, our procedure aims at finding a distribution of cutoffs and outcomes whose randomness comes from random applicant preferences, random grades, and random program preferences. We adopt the following iterative procedure:

- *Initialization:* The distribution of cutoffs across the M simulated samples, say, $\Phi_0 = \{\Phi_0^{(1)}, \dots, \Phi_0^{(M)}\}$, is initialized at the value obtained under DA in the observed sample. Therefore, Φ_0 is degenerate because it is constant across the simulated samples. We also calculate applicants’ and programs’ preferences in each simulation sample m :
 - (i) We compute applicants’ cardinal and therefore ordinal preferences by inputting the simulated shocks $\varepsilon_i^{(m)}$, estimated coefficients, and observed characteristics into the utility functions specified in equations (3) and (4). Note that shocks to applicant preferences ($\varepsilon_i^{(m)}$) are drawn **unconditionally** on the ordinal preferences observed in the sample.
 - (ii) We attribute to each program the same capacity as in the experiment and compute its “cardinal” preferences over applicants by plugging the simulated shocks ($\nu_i^{(m)}$ and $\xi_{i,j}^{(m)}$), estimated coefficients, and observed first-semester grades into the grade function (equation 7) and the latent score function (equation 6). These

preferences are then translated into rankings over applicants in increasing order of preference from 1 to I , denoted by $r_{i,j}^{(m)}$.

- *Iteration:* For $t = 1, 2, \dots$, iteration t begins with a distribution of cutoffs, Φ_{t-1} , and has two steps to obtain Φ_t :
 - (i) In each simulation sample m , each applicant submits the ROL that maximizes her expected utility given the “current” distribution of cutoffs, Φ_{t-1} . Necessarily, the optimal ROL takes into account the application cost in either DA-T or DA-C. This computation is further detailed in Appendix D.1. Recall that no applicant applies to unacceptable programs or programs whose prerequisites are not met by the applicant (Assumption 2). Based on the theoretical results in Section 2.2, an applicant’s optimal ROL can always be found by dropping some programs from her true preference order.
 - (ii) In each simulation sample m , we run the DA algorithm using ROLs submitted by applicants and simulated program preferences ($r_{i,j}^{(m)}$). This process results in a matching outcome and a cutoff for each program, which leads us to the updated cutoff distribution, Φ_t .
- *Stopping rule:* Stop when $\|\Phi_t - \Phi_{t-1}\|$ is small enough.

The resulting distribution of cutoffs, $\Phi_e = \{\Phi_e^{(1)}, \dots, \Phi_e^{(M)}\}$, characterizes a Bayesian Nash equilibrium under the given mechanism.

Sample-specific counterfactual actions and outcomes. We evaluate the counterfactual welfare conditional on applicants’ and programs’ observed ordinal preferences, which is case B1 in Table 10 in Section 5.3. That is, the shocks to applicant preferences (ε) are simulated conditional on the observed ROLs submitted by applicants under DA, and program ordinal preferences are also held constant at their observed values. In this way, we answer the following question: *Based on all the information observed by the researcher, what is the best prediction of the counterfactual outcome if a counterfactual market design had been implemented at TSE in 2013?* There are certainly many alternative counterfactual questions that one can ask. Our choice makes it easier to compare simulated equilibrium outcomes with the observed outcomes, because the latter certainly are derived from the realized ordinal preferences of applicants and programs.

More specifically, we adopt the following procedure to simulate a new set of samples:

- (i) In each simulation sample m , we compute applicants' cardinal and therefore ordinal preferences by inputting the simulated shocks $\varepsilon_i^{(m)}$, estimated coefficients, and observed characteristics into the utility functions specified in equations (3) and (4). As spelled out in Section 5.3, we now compute shocks to applicant preferences ($\varepsilon_i^{(m)}$) **conditionally** on the ordinal preferences observed in the sample using a Gibbs sampler described in Appendix D.2.
- (ii) Each applicant submits the ROL that maximizes her expected utility, given the equilibrium distribution of the cutoffs, Φ_e . Again, an applicant's optimal ROL can be found by dropping some programs from her true preference order. This process is conducted separately for each version of DA-T and DA-C.
- (iii) In each simulation sample, we run the DA algorithm with the optimal ROLs submitted by applicants and the observed program preferences.

Similar to the direct evaluation, we measure the performance of each design on several dimensions: (i) Match quality, which is summarized by stability (or the incidence of blocking pairs) and unmatched applicants,¹³ and (ii) screening costs. In addition to examining applicants' and programs' ordinal welfare, we evaluate their cardinal welfare, given that our preference estimates have cardinal implications.

5.5.2 Counterfactual Analysis: Results

Our first set of results is summarized in Figure 4, which shows the effects of market design on screening costs, measured as the number of candidates screened per opening, and on match quality, measured as the number of blocking pairs and the number of unmatched applicants.¹⁴

Subfigure (a) depicts the effects of DA-T of different degrees. DA-T- K , for $K \in \{1, 2, \dots, 7\}$, means that applicants are allowed to freely rank only up to K choices; therefore, DA-T-7 is DA (without truncation). The outcomes under DA with no cost are always constant in our sample-specific counterfactual analysis. In the figure, when increasing the application limit from DA-T-7 to DA-T-3, match quality remains (almost) the same as the applicant-optimal stable matching, while screening costs decrease monotonically from screening 4.16

¹³Note that the stable matching in our data is unique in all simulation samples, given the ordinal preferences of applicants and programs (c.f., Section 4). In other words, the applicant-optimal stable matching and the program-optimal stable matching coincide.

¹⁴The confidence intervals shown in Figure 4 consider the uncertainty in applicant preferences but not the uncertainty in the estimated coefficients, due to significant computational costs. Our simulations of DA-T taking into account both sources of uncertainty, being computationally less costly, show similar results to Figure 4a.

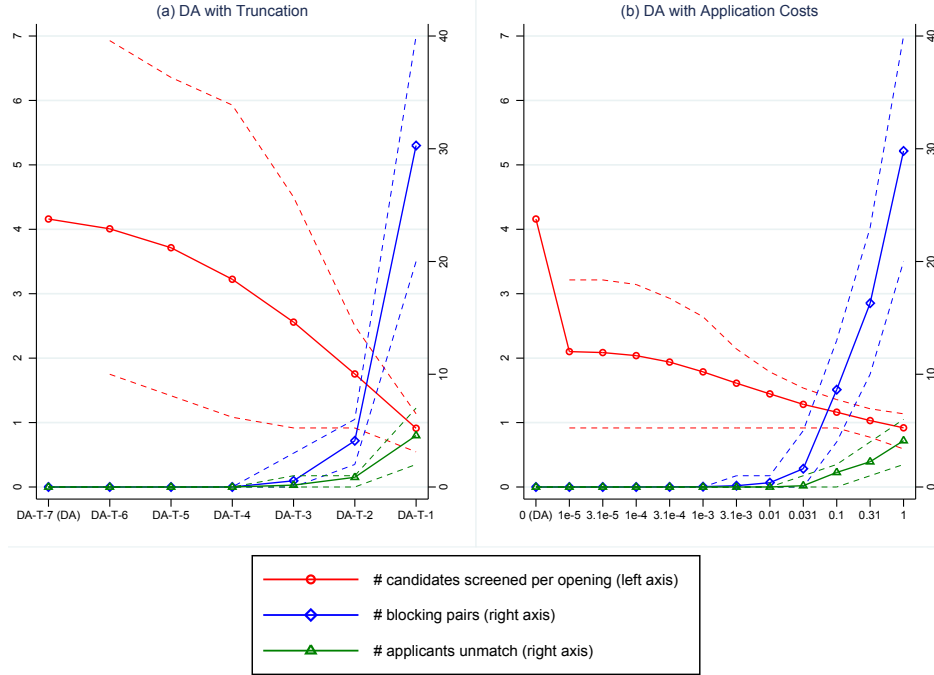


Figure 4: Match Quality under DA with Truncations or Finite Application Costs

Notes: This figure shows the counterfactual analysis of the effects of market design on match quality (measured by the number of blocking pairs and the number of unmatched applicants) and screening costs (measured by the number of candidates screened per opening). Subfigure (a) depicts the effects of DA-T with different degrees, where DA-T- K , $K \in \{1, 2, \dots, 7\}$, meaning that applicants are allowed to freely rank only up to K choices. Therefore, DA-T-7 is DA (without truncation). Subfigure (b) presents the effects of DA-C with a constant marginal cost for second and later choices and a zero cost of ranking the first choice. The constant marginal cost varies from 0 to 1 on a logarithmic scale, and it becomes DA when the cost is zero. The dashed lines are the 90% confidence intervals from the 4000 simulation samples. There is no variation in simulating the outcome under DA because all simulated applicant and program preferences are conditional on the observed ordinal preferences.

candidates per opening to screening 2.56 per opening. When the application limit tightens further, the screening costs decrease further to 1.75 and 0.91 candidates per opening under DA-T-2 and DA-T-1, respectively. However, this result comes at the expense of match quality: The number of blocking pairs increases from 0.5 under DA-T-3 to 4.10 under DA-T-2 and then skyrockets to 30.29 under DA-T-1; we also observe some unmatched applicants, 0.86 under DA-T-2 and 4.56 under DA-T-1.

Subfigure (b) examines the effects of DA-C with a constant marginal cost for second and further choices, the first choice being costless. As explained above, the marginal costs vary from 0 to 1.

Several interesting patterns emerge. First, the screening costs drop dramatically, even when there is a very low application cost: Imposing an application cost of 10^{-5} decreases the number of candidates screened per opening by 50%, while match quality remains the same. Second, the reduction in screening costs by this low application cost is larger than that of DA-T-3 (although smaller than that of DA-T-2). When the cost rises, screening costs further decrease, but the match quality begins to be impacted only when the cost is

above 0.0031. Third, when the application cost is 1, the outcome is almost identical to DA-T-1 because the cost is so high that it is not worth ranking a second program.

In summary, Figure 4 describes the trade-off between screening costs and match quality; most importantly, imposing a mild application cost in DA-T, and more so in DA-C, can significantly decrease screening costs without harming match quality. It should be emphasized that our equilibrium analysis considers all effects of application cost and measures the net effect. For example, in addition to the potentially negative effects on match quality, application costs can benefit some applications and increase their chance of being matched, because costs reduce the competition among some applicants when other applicants stop applying to some programs.

One may be concerned that our analysis of DA-C does not consider the costs that are actually paid by applicants. We note that application costs can be designed as a transfer, e.g., a fee paid to programs, rather than being purely wasteful, e.g., writing a motivation letter. In the following, we study how applicant welfare is affected by such costs.

5.5.3 Applicants Welfare

Cardinal welfare. Table 12 provides an evaluation based on the estimated cardinal preferences while explicitly taking into account the application costs. For each given market design, we first calculate cardinal welfare per applicant in each simulation sample and then measure its mean and standard deviation across all simulation samples. When doing so, we ignore the “strategizing” cost that an applicant has to incur to find an optimal ROL.

Panel A first presents per-applicant cardinal welfare across the DA-T of different degrees under which an applicant can submit a limited number of choices. The results confirm Figure 4 and show that DA-T-6 to DA-T-4 achieve the exact DA outcome, while DA-T-3 incurs a small welfare loss. Surprisingly, DA-T-2 obtains a welfare gain that is comparable to re-assigning 1.5 applicants from an “almost” unacceptable program to an “almost” definitely acceptable one (see Section 5.5.1 for a discussion on this cardinality measure). This outcome occurs even though DA-T-2 leaves 0.86 applicants unmatched on average, as partially reflected in the increase in the standard deviation of applicant welfare. DA-T-1 brings the lowest applicant welfare with the highest variance, consistent with the large number of unmatched applicants (4.56). Relative to DA, the welfare loss in DA-T-1 is comparable to moving 2.9 applicants from an almost unacceptable program to an almost definitely acceptable one.

The analysis of DA-C can also be found in Panel A of Table 12. In addition to describing

the per-applicant welfare produced by matching outcomes, we also report the application costs paid by applicants and their net welfare. In terms of matching outcomes, per-applicant welfare has a non-linear relationship with application cost: Welfare increases with the application cost when the cost is not too high (≤ 0.31), but it drops to the lowest when it is prohibitively high ($= 1$). In contrast, the standard deviation of per-applicant welfare monotonically increases with cost, consistent with the monotonically increasing number of unmatched applicants.

Among all DA-T and DA-C configurations, the maximum welfare gain relative to DA is obtained for DA-C with a cost 0.31, equivalent to re-assigning 5.8 applicants from an almost unacceptable program to an almost definitely acceptable one. The largest welfare loss (under DA-T-1) amounts to moving 1.1 applicants from an almost definitely acceptable program to an almost unacceptable one.

The costs actually paid by applicants follow the same pattern as the welfare. Applicants pay higher costs when the cost increases but remains moderate, and they stop applying to additional programs once the cost is prohibitively high. The maximum costs are paid when the marginal cost is 0.31: The magnitude is equivalent to re-assigning 7.8 applicants from an almost definitely acceptable program to an almost unacceptable one.

As discussed earlier, we can consider the application cost as a transfer from applicants to programs, e.g., a transfer of monetary fees. Nonetheless, one may be interested in applicants' net welfare, given the application cost they pay. Panel A of Table 12 shows that the net welfare is monotonically decreasing in application cost. Relative to DA, the maximum welfare loss (at marginal cost $= 1$) amounts to re-assigning 2.5 applicants from an almost definitely acceptable program to an almost unacceptable one, which is still slightly lower than the loss under DA-T-1.

In summary, Panel A confirms that the welfare loss due to application costs (in the form of DA-T or DA-C) is negligible when the cost is low. Moreover, combined with Figure 4, DA-C appears to be a better choice than DA-T: We can set the marginal cost as high as 0.01 so that there is almost no loss in match quality, and the number of candidates screened is 1.44 per opening. To reach a comparable screening cost, DA-T has to be of degree 2 (DA-T-2), under which programs screen 1.75 candidates per opening and 0.86 applicants are unmatched.

Panel B of Table 12 evaluates the cardinal welfare of the game play observed in the experiment. To simulate these results, we hold constant applicants' actions across the 4000 simulation samples. DA-T-4 reaches almost identical welfare levels both in equilibrium and

Table 12: Cardinal Welfare per Applicant under Different Market Designs

DA-T of diff. degrees			DA-C w/ diff. marginal application cost ^c						
Welfare per app.		Marginal cost	Welfare per app.		Costs paid		Net welfare		
mean	s.d.		mean	s.d.	mean	s.d.	mean	s.d.	
<i>A. Counterfactual Analysis^a</i>									
DA	3.2489	0.1064	0 (DA)	3.2489	0.1064	0	0	3.2489	0.1064
DA-T-6	3.2489	0.1064	0.00001	3.2489	0.1064	0.0000	0.0000	3.2489	0.1064
DA-T-5	3.2489	0.1064	0.000031	3.2489	0.1064	0.0000	0.0000	3.2488	0.1064
DA-T-4	3.2489	0.1064	0.0001	3.2489	0.1064	0.0001	0.0000	3.2488	0.1064
DA-T-3	3.2479	0.1064	0.00031	3.2490	0.1064	0.0004	0.0000	3.2486	0.1064
DA-T-2	3.2558	0.1070	0.001	3.2492	0.1064	0.0010	0.0001	3.2482	0.1064
DA-T-1	3.2328	0.1136	0.0031	3.2499	0.1065	0.0025	0.0003	3.2475	0.1064
			0.01	3.2522	0.1065	0.0061	0.0009	3.2461	0.1065
			0.031	3.2594	0.1068	0.0134	0.0028	3.2460	0.1068
			0.1	3.2721	0.1083	0.0292	0.0092	3.2430	0.1086
			0.31	3.2806	0.1103	0.0425	0.0294	3.2381	0.1135
			1	3.2426	0.1136	0.0075	0.0950	3.2351	0.1477
<i>B. In the Experiment^b</i>									
DA-T-4	3.2494	0.1061	DA-C	3.2409	0.1061	-	-	-	-

Notes: The calculations are based on 4000 simulations. In each simulation, the shocks to applicants' cardinal preferences are drawn conditional on the observed true ordinal preferences (i.e., the submitted ROLs under DA, supplemented with the program acceptability information). More details are provided in Appendix D.2. In each simulation sample, we calculate cardinal welfare per application (i.e., the average of applicant welfare); the table reports the means and standard deviations of per-applicant welfare across simulation samples. Therefore, the standard deviation measures the variation across simulation samples but not across applicants. When simulating matching outcomes, program preferences are always fixed at the observed values. ^a In counterfactual analysis, applicants play an equilibrium strategy in each simulation sample, while the equilibrium is solved numerically, as in Appendix D.1. Given the realization of cardinal preferences, applicants may play different actions across samples, as dictated by the equilibrium strategy. ^b For these calculations, applicants always take the actions played in the experiment across all simulation samples. In the experiment, DA-C allows the applicant to rank up to three choices, and a motivation letter is required for each additional choice. However, we do not have a measure of the magnitude of the cost.^c In the counterfactual analysis, DA-C allows applicants to rank one choice but requires a constant marginal cost for additional choices.

in the experiment. However, DA-C in the experiment has a different cost configuration from that in the counterfactual. In fact, DA-C in the experiment is the same as DA-T-3 plus a cost for additional choices, and therefore, the former should weakly dominate DA-T-3. While this prediction is true in simulations of equilibrium outcomes, it is not true in the actual game played.¹⁵ The welfare loss is due largely to the unmatched applicant under DA-C in the experiment, which may be explained by the fact that applicant actions, which are held constant across simulation samples, can be suboptimal in some samples. Moreover, this loss may indicate that applicants do not use equilibrium strategies in the experiment.

Welfare effect on individual applicants. In addition to examining cardinal welfare, we further investigate how many applicants are better off or worse off. The advantage of this welfare measure is that it does not compare welfare across applicants. The results are summarized in Table 13, which has a similar structure as Table 12.

¹⁵When simulating the experimental configuration of DA-C, almost no applicants rank more than three choices, even when the cost is negligible. Therefore, this simulation leads to the same outcome as DA-T-3, which obtains (almost) the same outcome as DA.

Table 13: Applicant Ordinal Welfare under Different Market Designs:
Counterfactual Analysis; Relative to the Applicant-Optimal Stable Matching

	DA-T of diff. degrees				DA-C w/ diff. marginal application cost					
	Better off		Worse off		Marginal cost	Better off		Worse off		
	Mean	s.d.	Mean	s.d.		Mean	s.d.	Mean	s.d.	
DA-T-6	0.00	0.00	0.00	0.00	0.00001	0.00	0.00	0.00	0.00	
DA-T-5	0.00	0.00	0.00	0.00	0.000031	0.00	0.04	0.00	0.04	
DA-T-4	0.00	0.00	0.00	0.00	0.0001	0.00	0.05	0.00	0.05	
DA-T-3	0.04	0.20	0.21	0.41	0.00031	0.01	0.09	0.01	0.11	
DA-T-2	1.40	0.88	2.16	0.86	0.001	0.03	0.18	0.03	0.20	
DA-T-1	6.38	1.34	10.16	2.12	0.0031	0.11	0.34	0.12	0.38	
					0.01	0.36	0.60	0.40	0.67	
					0.031	1.13	1.00	1.24	1.11	
					0.1	3.56	1.57	3.71	1.47	
					0.31	5.55	1.46	6.87	1.89	
					1	6.33	1.34	10.25	2.22	

Notes: The counterfactual simulation is the same as that in Table 12; the notes therein provide more details. In each of the 4000 simulation samples, we first calculate the number of applicants who are better off (or worse off) relative to the DA outcome, which is the applicant-optimal stable matching. The table then reports the means and standard deviations of these two ordinal welfare measures under a given market design across simulation samples.

In each simulation sample, we calculate how many applicants are better off (or worse off) relative to their individual outcome under DA – the applicant-optimal stable matching. The table then reports the means and standard deviations across simulation samples. We do not consider the costs paid by applicants.

Consistent with the previous results, applicants (almost) always obtain the applicant-optimal stable outcome when the application cost is low (i.e., under DA-T-3 to DA-T-6 or DA-C with a marginal cost of no more than 0.01). However, whenever the outcomes deviate from the applicant-optimal stable outcome, there are more “losers” than “winners.” This finding indicates that the gain in per-applicant welfare under DA-T-2 or DA-C with costs 0.1 or 0.31 comes at the expense of some applicants, as shown in Table 12. That is, “winners” on average benefit more in cardinal welfare than what “losers” lose.

5.5.4 Program Welfare in Matching Outcomes

Given the two-sided nature of the matching game, we are also interested in program welfare from matching outcomes without taking into account screening costs.

We define two measures of program welfare while maintaining the assumption that programs have responsive preferences.¹⁶ The first measure is the sum of the final grades of

¹⁶We choose not to use the estimate of the latent score (equation 6) to measure programs’ cardinal preferences because (i) only the ordinal information in program preferences matter in the game (not the cardinality) and (ii) the only observable determining the latent score is applicants’ final grade. The latter justification motivates our first measure of program welfare.

all of the program’s matched applicants. That is, for program j , its welfare given matching μ is $\sum_{i \in \mu^{-1}(j)} FinalGrade_i$. The second measure takes into account programs’ ordinal preferences, $\sum_{i \in \mu^{-1}(j)} r_{i,j}$ for each j . Namely, j receives 129 points if matched with its top applicant and 1 point if matched with its lowest-ranked applicant. By construction, both measures take into account the quantity and quality of matched applicants.

Similar to the analysis of applicants’ cardinal welfare, we first study per-program welfare (Table 14) and then investigate the number of “winners” and “losers” according to the two welfare measures (Table 15). It should be emphasized that this welfare analysis does not consider screening costs.

Table 14: Programs’ Cardinal Welfare under Different Market Designs

DA-T of diff. degrees					DA-C w/ diff. marginal application cost				
Welfare per program					Welfare per program				
Def. 1: Grades		Def. 2: Rank			Marginal cost	Def. 1: Grades		Def. 2: Rank	
Mean	s.d.	Mean	s.d.			Mean	s.d.	Mean	s.d.
<i>A. Counterfactual Analysis</i>									
DA	226.10	0	1456.14	0	0 (DA)	226.10	0	1456.14	0
DA-T-6	226.10	0.00	1456.14	0.00	0.00001	226.10	0.00	1456.14	0.00
DA-T-5	226.10	0.00	1456.14	0.00	0.000031	226.10	0.00	1456.14	0.00
DA-T-4	226.10	0.00	1456.14	0.00	0.0001	226.10	0.00	1456.14	0.10
DA-T-3	225.85	0.55	1455.63	1.03	0.00031	226.10	0.00	1456.14	0.33
DA-T-2	224.83	0.52	1453.44	1.18	0.001	226.10	0.02	1456.12	0.52
DA-T-1	219.49	2.37	1435.68	11.38	0.0031	226.10	0.05	1456.06	0.96
					0.01	226.09	0.10	1455.92	1.50
					0.031	225.90	0.55	1454.52	4.09
					0.1	224.02	1.21	1444.81	7.35
					0.31	222.77	1.61	1442.72	8.55
					1	220.17	2.30	1436.83	10.57
<i>B. In the Experiment</i>									
DA-T	226.10	0	1448.14	0	DA-C	224.63	0	1445.43	0

Notes: The counterfactual simulation is the same as that in Table 12; the notes therein provide more details. In each of the 4000 simulation samples, we first calculate the per-program welfare according to two measures. The first is the sum of final grades of all applicants matched with the given program, and the second is the sum of matched applicants’ rankings by the program (129 points if matched with the top-ranked applicant, 1 point if matched with the lowest-ranked applicant). The table then reports the mean and standard deviation of these two welfare measures under a given market design across simulation samples. The standard deviation is zero for program welfare under DA, DA-T and DT-C in the experiment because there is no variation in matching outcomes across simulation samples.

A clear pattern evident in Table 14 is that program welfare decreases when application costs increase under either DA-T or DA-C. This pattern is explained mainly by the number of unmatched applicants. If an average applicant – with a final grade of 12.27 and ranked by the program as 65th – becomes unmatched, the per-program welfare decreases by 1.75 or 9.29, according to the two welfare measures, respectively. For example, under DA-T-1, 4.11 applicants on average are unmatched (Figure 4), which leads to a welfare loss of approximately 7 or 38. The actual loss is smaller because unmatched applicants tend to

have lower final grades and/or to be ranked lower by programs.

Nonetheless, Table 14 shows that the loss in program welfare is negligible when the application cost is low. Even in the extreme, under DA-T-1 or DA-C with a marginal cost equal to one, on average, the loss to a program’s welfare amounts to losing an applicant with a final grade ranging from 5.93 to 6.61 or one ranked 109th. On the other hand, the standard deviation indicates that the effect on program welfare varies more significantly across simulation samples when the application cost increases.

Comparing the welfare of each program under a given market design with that under DA (Table 15), we find more “losers” than “winners” when the application cost increases. This pattern holds true with both of the welfare measures. Once again, when the application cost is low, the effects on program welfare are negligible.

Table 15: Program Welfare under Different Market Designs
Counterfactual Analysis; Relative to the Program-Optimal Stable Matching

	DA-T of diff. degrees				DA-C w/ diff. marginal application cost ^c	DA-C w/ diff. marginal application cost ^c				
	Better off		Worse off			Marginal cost	Better off		Worse off	
	Mean	s.d.	Mean	s.d.			Mean	s.d.	Mean	s.d.
<i>A. Program Welfare Definition 1: Grades^b</i>										
DA-T-6	0.00	0.00	0.00	0.00	0.00001	0.00	0.00	0.00	0.00	
DA-T-5	0.00	0.00	0.00	0.00	0.000031	0.00	0.01	0.00	0.01	
DA-T-4	0.00	0.00	0.00	0.00	0.0001	0.00	0.01	0.00	0.01	
DA-T-3	0.01	0.03	0.03	0.07	0.00031	0.00	0.01	0.00	0.02	
DA-T-2	0.12	0.06	0.29	0.10	0.001	0.00	0.03	0.01	0.04	
DA-T-1	0.26	0.13	0.55	0.12	0.0031	0.02	0.05	0.03	0.08	
					0.01	0.05	0.08	0.08	0.13	
					0.031	0.11	0.10	0.21	0.16	
					0.1	0.18	0.11	0.41	0.11	
					0.31	0.27	0.12	0.48	0.11	
					1	0.28	0.13	0.54	0.11	
<i>B. Program Welfare Definition 2: Rank^b</i>										
DA-T-6	0.00	0.00	0.00	0.00	0.00001	0.00	0.00	0.00	0.00	
DA-T-5	0.00	0.00	0.00	0.00	0.000031	0.00	0.01	0.00	0.01	
DA-T-4	0.00	0.00	0.00	0.00	0.0001	0.00	0.01	0.00	0.01	
DA-T-3	0.01	0.03	0.03	0.07	0.00031	0.00	0.01	0.00	0.02	
DA-T-2	0.16	0.09	0.26	0.11	0.001	0.00	0.03	0.01	0.04	
DA-T-1	0.36	0.15	0.45	0.12	0.0031	0.02	0.05	0.03	0.08	
					0.01	0.05	0.08	0.08	0.13	
					0.031	0.12	0.10	0.21	0.16	
					0.1	0.20	0.11	0.39	0.11	
					0.31	0.33	0.13	0.42	0.12	
					1	0.36	0.13	0.45	0.12	

Notes: The counterfactual simulation is the same as that in Table 12; the notes therein provide more details. In each of the 4000 simulation samples, we first calculate how many programs are better off (or worse off) relative to the program-optimal matching according to two measures. The first is the sum of final grades of all applicants matched with the given program, and the second is the sum of matched applicants’ rankings by the program (129 points if matched with the top-ranked applicant, 1 point if matched with the lowest-ranked applicant). Under a given market design, the table then reports the means and standard deviations of the measures on welfare changes across simulation samples.

6 Conclusion and Discussion

This paper investigates the market design for a two-sided many-to-one matching market with the application cost as a Pigouvian tax to reduce congestion and lower the screening costs of programs that are recruiting applicants. Both forms of application cost, DA-T (an application limit) and DA-C (a positive marginal cost of application), are effective in reducing congestion. However, some key differences should be noted.

First, under DA-T of degree K , the marginal cost of applying to an additional program is zero when the total number of applications is below K . However, the marginal benefit of applying to an additional acceptable program, which can be small, is always positive because of the full-support assumption on applicant and program preferences. Therefore, applicants always optimally apply to K programs under DA-T- K whenever they have at least K acceptable programs. In comparison, under DA-C, an applicant considers the positive marginal cost for additional applications and therefore may choose not to apply to an acceptable program.

Second, the programs to which an applicant chooses not to apply under DA-C must have marginal benefits lower than the marginal costs, implying that the expected loss in match quality is bounded by application cost. DA-T does not offer such an opportunity for bounding losses. This observation implies the possibility that the market designer, with little information on applicants' and programs' preferences, chooses a very low application cost and reduces congestion significantly without sacrificing match quality.

Third, although neither DA-T nor DA-C is strategy-proof, DA-C offers an opportunity for less strategic applicants to play safe by paying application costs. This is a desirable feature as non-strategic applicants may lose to sophisticated applicants in a non-strategy-proof mechanism (Pathak and Sönmez, 2008). DA-T can only provide such a safety measure when it does not substantially restrict the number of applications, in which case it is not effective in reducing congestion.

Fourth, DA-C is also more flexible in the sense that application costs can be non-linear or program-specific. For instance, a popular program can impose a high application cost.

Lastly, although application costs and pollution taxes are essentially the same, application limits, or DA-T, do not resemble a cap-and-trade scheme, because there is no trading of rights to submit applications in a matching market. Therefore, the equivalence between regulating prices and fixing quantities in the setting of Weitzman (1974) no longer holds in our matching model.

6.1 Experiment and Structural Analyses

Our study marries real-life, experimental data with a structural analysis. This choice of research design deserves some discussion, especially in terms of experimental design, behavioral assumptions in the structural analysis, and external validity.

One might argue that, given the realized preferences of both applicants and programs, our choice of mechanisms, DA-T-4 and a version of DA-C under which it is free to submit up to three choices, is too conservative. In the end, 99 percent of the applicants are matched with their top-3-preferred programs. At the time of designing the experiment, we were not willing to gamble on the match quality in exchange for uncertain screening costs, given that it was the first time this formal centralized match was implemented in our specific setting. It is true that, had we known the structure of preferences, we might have tried more restrictive DA-T and DA-C. In a sense, because it is a field experiment, we have mimicked what would happen in real life in order to learn about the structural parameters.

In the structural analysis, we assume that every applicant behaves rationally. However, our data from the experiment indicate that applicants are not fully optimizing, probably because of the complexity of the game. Importantly, the conclusions drawn from the experiment and those from the structural analysis are similar, both qualitatively and quantitatively. This implies that deviations from full rationality, which is a topic outside the scope of this paper, have no significant consequences on the conclusions of our study.

The question of external validity is difficult to address without considering specific benchmarks. In our experiment, applicant preferences are highly heterogeneous (i.e., “horizontal” applicant preferences), while programs have similar ordinal preferences over applicants (i.e., “vertical” program preferences). Such preference patterns are also documented or considered in the literature, for example, Agarwal (2015) on the medical match in the US, Carvalho et al. (forthcoming) on university admissions in Brazil, and Che and Koh (2016) on decentralized college admissions. One may expect that congestion should be less likely to happen in these settings, and yet it does occur in our experiment. One may further argue that application costs would be less effective in other settings, e.g., those with more homogeneous applicant preferences. Although the exact effects of application costs would depend on the context, the way how application costs work remains the same: an application cost will discourage an applicant from applying to programs that are unlikely to accept her and/or are not so attractive to her.

6.2 Market Design and Congestion in Real-Life Markets

With due caution in interpreting its external validity, our study may shed light on market design in other centralized or decentralized matching. In the literature, the Gale-Shapley DA mechanism has been used to approximate equilibrium outcomes of decentralized settings for marriage and other markets (Chiappori and Salanié, 2016). It is therefore a natural starting point to focus on variants of DA.

Various forms of application costs are used to combat congestion across matching markets. Table 16 presents examples of application costs used in practice in both centralized and decentralized markets, including DA-T and DA-C.

Table 16: Examples of Various Market Designs in Practice

Mechanism	Centralized	Decentralized
DA (cost ≈ 0)	School choice in Boston, MA	Research articles and academic conferences Job market for Econ PhDs Online platforms: Craigslist and Zillow
DA-T	NYC high schools (ranking 12 choices) Lycées in Paris (ranking 8 choices) Univ. admissions in Australia/Brazil/Chile	Airbnb (1 request to book) “Early Decision” in college admissions (apply to only one college) Research articles and academic journals (except law journals)
DA-C	University admission in Hungary National Resident Matching Program	Unit cost (marginal cost > 0): Upwork; College admissions in the U.S.; Graduate admissions in the U.S. Fixed cost: eHarmony (18-hour questionnaire)

While DA-T in centralized markets usually allows applicants to submit multiple applications (e.g., school choice in New York City and Paris), in decentralized markets, DA-T-1 is the version most often applied in practice when DA-T is adopted: Airbnb allows one “request to book” at any point in time; one can apply to only one college for early decision in the U.S.; academic journals, except those in law, do not permit an article to be considered by multiple journals simultaneously. A possible explanation for these practices is that screening costs are significantly high in these markets.¹⁷

The usual configuration of DA-C in centralized markets is similar to that developed in the paper: Applicants pay a constant marginal cost if one applies to more programs beyond a limited number of choices, e.g., NRMP in the U.S. and university admissions in Hungary.

In decentralized markets, however, the application cost in DA-C can be either a marginal cost or a fixed cost. An applicant pays an application fee or makes program-specific investments when applying to an additional college or graduate program, which amounts to paying a marginal cost. Similarly, one may have to pay some fees to apply to an additional job on

¹⁷Interestingly, academic journals in law, which allow an article to be considered simultaneously by several journals, ask law students to review submitted articles, which may reduce screening costs.

Upwork, a platform matching freelancers and tasks. An example of a fixed cost is the costly registration for the eHarmony dating platform – it may take up to 18 hours to fill out the questionnaire appropriately. Application costs have great potentials in combatting congestion in decentralized market. For example, Pallais (2015) documents that application costs significantly influence the number of per-student applications in the U.S. college admissions.

Meanwhile, there are many markets with (almost) zero application cost (Table 16). This practice can be justified when screening costs are low. For example, academic conferences in economics allow multiple submissions of the same article, while the review process is less demanding than that for academic journals. Other examples are more commonly a result of the recent development of decreasing application costs. For instance, the rise of platforms such as EconJobMarket for economics PhDs has led to an increase in the number of applications that economics departments have received. This almost-zero “application cost” is also observed on online platforms such as Craigslist and Zillow.

In summary, our results can help us understand the current design in these markets and, more importantly, how to improve them when congestion is a concern due to non-negligible screening costs.

References

- Abdulkadiroglu, Atila and Tayfun Sönmez**, “Matching markets: Theory and practice,” *Advances in Economics and Econometrics*, 2013, 1, 3–47.
- Abdulkadiroğlu, Atila, Nikhil Agarwal, and Parag A Pathak**, “The welfare effects of coordinated assignment: evidence from the New York City high school match,” *American Economic Review*, 2017, 107 (12), 3635–89.
- Agarwal, Nikhil**, “An empirical model of the medical match,” *American Economic Review*, 2015, 105 (7), 1939–78.
- and **Paulo Somaini**, “Demand Analysis using Strategic Reports: An Application to a School Choice Mechanism,” *Econometrica*, Forthcoming.
- Ajayi, Kehinde**, “School Choice and Educational Mobility: Lessons from Secondary School Applications in Ghana,” 2013. Manuscript.
- Arnosti, Nick, Ramesh Johari, and Yash Kanoria**, “Managing congestion in dynamic matching markets,” 2016.
- Artemov, Georgy, Yeon-Koo Che, and Yinghua He**, “Strategic ‘Mistakes’: Implications for Market Design Research,” 2017.
- Ashlagi, Itai, Mark Braverman, Yash Kanoria, and Peng Shi**, “Clearing Matching Markets Efficiently: Informative Signals and Match Recommendations,” 2018.

- Avery, Robert B, Lars Hansen, and V. Joseph Hotz**, “Multi-period Probit Models and Orthogonality Condition Estimation,” *International Economic Review*, 1983, *24* (1), 21–35.
- Azevedo, Eduardo and Jacob Leshno**, “A Supply and Demand Framework for Two-Sided Matching Markets,” *Journal of Political Economy*, 2016.
- Baumol, William J.**, “On Taxation and the Control of Externalities,” *American Economic Review*, 1972, *62* (3), 307–322.
- Beggs, Steven, Scott Cardell, and Jerry Hausman**, “Assessing the potential demand for electric cars,” *Journal of econometrics*, 1981, *17* (1), 1–19.
- Blundell, Richard**, “What Have We Learned from Structural Models?,” *American Economic Review*, 2017, *107* (5), 287–292.
- Calsamiglia, Caterina, Chao Fu, and Maia Güell**, “Structural Estimation of a Model of School Choices: the Boston Mechanism vs. Its Alternatives,” 2014. Barcelona GSE Working Paper No. 811.
- Carvalho, José-Raimundo, Thierry Magnac, and Qizhou Xiong**, “College Entry and Allocation Mechanisms,” *Quantitative Economics*, forthcoming.
- Chade, Hector, Gregory Lewis, and Lones Smith**, “Student portfolios and the college admissions problem,” *Review of Economic Studies*, 2014, *81* (3), 971–1002.
- Che, Yeon-Koo and Youngwoo Koh**, “Decentralized College Admissions,” *Journal of Political Economy*, 2016, *124* (5), 1295–1338.
- Chiappori, Pierre-Andre and Bernard Salanié**, “The Econometrics of Matching Models,” *Journal of Economic Literature*, 2016.
- Coles, Peter, Alexey Kushnir, and Muriel Niederle**, “Preference signaling in matching markets,” *American Economic Journal: Microeconomics*, 2013, *5* (2), 99–134.
- Dubins, Lester E. and David A. Freedman**, “Machiavelli and the Gale-Shapley Algorithm,” *American Mathematical Monthly*, 1981, *88* (7), 485–494.
- Fack, Gabrielle, Julien Grenet, and Yinghua He**, “Beyond Truth-Telling: Preference Estimation with Centralized School Choice and College Admissions,” *American Economic Review*, Forthcoming.
- Fradkin, Andrey**, “Search frictions and the design of online marketplaces,” *NBER Working Paper*, 2014.
- Gale, David E. and Lloyd S. Shapley**, “College Admissions and the Stability of Marriage,” *American Mathematical Monthly*, 1962, *69* (1), 9–15.
- Haeringer, Guillaume and Flip Klijn**, “Constrained School Choice,” *Journal of Economic Theory*, 2009, *144* (5), 1921–1947.
- Hafalir, Isa E, Rustamdjan Hakimov, Dorothea Kübler, and Morimitsu Kurino**, “College admissions with entrance exams: Centralized versus decentralized,” Technical Report, SFB 649 Discussion Paper 2016.

- Halaburda, Hanna, Mikolaj Jan Piskorski, and Pinar Yildirim**, “Competing by restricting choice: the case of search platforms,” *Harvard Business School Strategy Unit Working Paper*, 2015, (10-098).
- Hassidim, Avinash, Assaf Romm, and Ran Shorrer**, ““Strategic” Behavior in a Strategy-proof Environment,” 2016. mimeo.
- Hastings, Justine, Thomas Kane, and Douglas Staiger**, “Heterogeneous Preferences and the Efficacy of Public School Choice,” 2008. Manuscript.
- He, Yinghua**, “Gaming the Boston School Choice Mechanism in Beijing,” 2015. TSE Working Paper No. 15-551.
- Horton, John J.**, “Supply constraints as a market friction: Evidence from an online labor market,” 2015.
- Hwang, Sam Il Myoung**, “How does heterogeneity in beliefs affect students in the Boston Mechanism,” *Working Paper*, 2017.
- Kanoria, Yash and Daniela Saban**, “Facilitating the search for partners on matching platforms: Restricting agent actions,” 2017.
- Knittel, Christopher R and Ryan Sandler**, “The welfare impact of indirect pigouvian taxation: Evidence from transportation,” Technical Report, National Bureau of Economic Research 2013.
- Lee, Robin S. and Michael Schwarz**, “Interviewing in two-sided matching markets,” *The RAND Journal of Economics*, 2017, 48 (3), 835–855.
- Milgrom, Paul and Robert Weber**, “Distributional Strategies for Games with Incomplete Information,” *Mathematics of Operations Research*, 1985, 10 (4), 619–632.
- Muller, Nicholas Z. and Robert Mendelsohn**, “Efficient Pollution Regulation: Getting the Prices Right,” *American Economic Review*, December 2009, 99 (5), 1714–39.
- Nguyen, Kim, Michael Peters, and Michel Poitevin**, “Can EconJobMarket help Canadian universities?,” *Canadian Journal of Economics*, 2018, 50 (5), 1573–1594.
- Pallais, Amanda**, “Small Differences That Matter: Mistakes in Applying to College,” *Journal of Labor Economics*, 2015, 33 (2), 493–520.
- Pathak, Parag A.**, “The mechanism design approach to student assignment,” *Annu. Rev. Econ.*, 2011, 3 (1), 513–536.
- and **Peng Shi**, “How Well Do Structural Demand Models Work? Counterfactual Predictions in School Choice,” Technical Report, National Bureau of Economic Research 2017.
- and **Tayfun Sönmez**, “Leveling the playing field: Sincere and sophisticated players in the Boston mechanism,” *American Economic Review*, 2008, 98 (4), 1636–52.
- Pigou, Arthur C.**, *The Economics of Welfare*, London: Macmillan, 1920.
- Rogerson, Richard, Robert Shimer, and Randall Wright**, “Search-Theoretic Models of the Labor Market: A Survey,” *Journal of Economic Literature*, 2005, 43 (4), 959–988.

- Roth, Alvin E.**, “The Economics of Matching: Stability and Incentives,” *Mathematics of Operations Research*, 1982, 7 (4), 617–628.
- , “A Natural Experiment in the Organization of Entry-Level Labor Markets: Regional Markets for New Physicians and Surgeons in the United Kingdom,” *American Economic Review*, 1991, 81 (3), 415–40.
- and **Marilda A. Oliveira Sotomayor**, *Two-Sided Matching: A Study in Game-Theoretic Modeling and Analysis* Econometric Society Monograph Series, Cambridge University Press, 1990.
- and **Xiaolin Xing**, “Turnaround Time and Bottlenecks in Market Clearing: Decentralized Matching in the Market for Clinical Psychologists,” *Journal of Political Economy*, 1997, 105 (2), 284–329.
- Sandmo, Agnar**, “Optimal Taxation in the Presence of Externalities,” *The Swedish Journal of Economics*, 1975, 77 (1), 86–98.
- Seabright, Paul and Ananya Sen**, “Statistical Externalities and the Labour Market in the Digital Age,” 2015.
- Shorrer, Ran I and Sándor Sóvágó**, “Obvious Mistakes in a Strategically Simple College-Admissions Environment,” Technical Report, Working paper 2017.
- Weitzman, Martin L.**, “Prices vs. Quantities,” *Review of Economic Studies*, 1974, 41 (4), 477–491.

(For Online Publication)

Appendix to

A Pigouvian Approach to Congestion in Matching Markets

Yinghua He Thierry Magnac

November 9, 2018

List of Appendices

Appendix A: Experiment: Design and Implementation	A.2
Appendix B: Data	A.5
Appendix C: Likelihood Function based on the Extended Rank-Order List	A.7
Appendix D: Counterfactual Analysis	A.18
Appendix E: Screen Shots from the Web Site for the Experiment	A.24

Appendix A Experiment: Design and Implementation

A.1 The Boston Immediate Acceptance Mechanism

A popular mechanism is the Immediate Acceptance mechanism (IA), also known as the Boston mechanism, which solicits ROLs from applicants, uses programs' strict rankings over applicants as inputs, and includes multiple rounds:

Round 1. Each program considers all the applicants who rank it first and assigns its positions in the order of the program's preferences until either no positions remain or no applicant who has listed it as his/her first choice remains.

Generally, in

Round k ($k > 1$). The k^{th} choice of applicants who have not yet been assigned is considered. Each program that still has available positions assigns the remaining positions to applicants who rank the program as their k^{th} choice in the order of that program's preferences until either no positions remain or no applicant who has listed it as his/her k^{th} choice remains.

The process terminates after any round k in which every applicant is assigned a position at some program or if the only applicants who remain unassigned listed no more than k choices.

A.2 Design: Interface for Applicants

The experiment provides applicants with an interface hosted on the official website that Université Toulouse 1 Capitole uses for coursework. With screenshots from the website (Figures E.1–E.5), this appendix details each step of the experiment involving applicants.

To minimize the potential “framing” effects, we randomize on three dimensions: (1) The order of the programs in the survey questionnaire on acceptability, (2) the order of the programs when applicants are asked to rank them under the four mechanisms, and (3) the order of mechanisms in which applicants play. First, applicants are randomly and (almost) evenly divided into seven groups. Second, for each group, we assign an order of the seven programs to be presented when asking applicants to report acceptability. The seven program orders are such that each program is presented once first, once last, and the rest at random positions. Third, each group faces the “reverse” program order when asked to rank the programs under the four mechanisms. That is, if P1 is the first to be shown in the acceptability survey, it is presented last when the applicant decides her ROL, and vice versa. Moreover, these seven program orders are the same ones used in the survey. Finally, each group of applicants plays the four mechanisms following one of the four orders conditional on DA-C and DA-T always being after DA.

The experiment begins with the welcome page (Figure E.1) to which applicants have access after being notified by email. It describes in both English and French the purpose of the experiment (i.e., trying four mechanisms instead of adopting only one mechanism) and explains important dates and other information. The above content is the same for every applicant; then, every applicant clicks the link that takes her to the next screen, which differs across applicants.

On the second webpage, we implement a survey on program acceptability (Figure E.2). For each given program, applicants reveal their enrollment decision under the assumption that they were only accepted by that program. It is emphasized that the answers to survey questions are not used in the actual matching mechanisms. The order of the programs is randomized to avoid potential framing effects. Similarly, the link to the next step is also randomized such that one may play either IA or DA first.

Once they have reached the webpage to submit an ROL under a given mechanism, applicants are given suggestions on how to play the specific game (Figures E.3-E.5). It is also emphasized that each ROL submitted under a given mechanism is used if and only if the mechanism is selected for the final assignment. Under every mechanism, all the programs are listed for applicants to rank and the order of the programs is randomized.

A.3 Programs' Involvement in the Experiment

When designing the experiment, all the program directors were informed about the experiment and its intended purposes. The seven directors were involved at all stages of the experiment.

Pre-experiment stage. Each director was asked to provide a short description of the program, prerequisites (if any), and a capacity. This information was then uploaded to the website for the experiment.

Screening candidates. We did not simply forward applicants' submitted ROLs under the four mechanisms to the program directors. There were several concerns. First, we intended to prevent the directors from learning applicants' ROLs (i.e., how each applicant ranked the programs). Second, providing all ROLs would require the directors to screen too many candidates (c.f., Table 8).

Instead, in the files containing all the applicants, their grades and their M1 program enrollment information, we **pre-selected** a subset of applicants for each program. Each program director received a program-specific file and was asked to rank all the applicants within it. We explained that the pre-selected applicants were those who would be considered

by the program with a high probability, while the others were unlikely to be considered. At the same time, we strongly encouraged them to screen as many candidates as possible. When ranking candidates, the program directors had flexibility: They did not have to respect the ranking by grades; the ranking could be weak; and if necessary, ties in the mechanisms would be broken by grades.

We pre-selected applicants by running the four mechanisms with “fictitious” program preferences and capacities. We created those program preferences by adding noises to applicants’ final grades. For each program, the pre-selected applicants were those who had ever been “considered” (either accepted or rejected) by the program in these simulations.

Among the seven programs, the number of pre-selected candidates ranged from 27 to 52, with a mean of 41. No director screened more candidates beyond the pre-selected ones. This process was proven “successful” in the sense that beyond the pre-selected applicants for a program, no one was ever “considered” by the program under any mechanism.

Motivation letters. The directors did not have access to the motivation letters because receiving a letter from an applicant would imply that the applicant did not rank the program within their top three choices. We planned to transmit the letters to the directors if (a) there were a large number of applicants assigned to the fourth or lower choices under DA-C and (b) if the administration would choose the DA-C outcome upon seeing the tentative DA-C outcome together with the outcomes from the other three mechanisms. Ultimately, all applicants except one – who was unmatched – were matched with one of their top three choices (see Table 6). Therefore, the tentative DA-C outcome was very similar to the “real one” that could have been obtained had the directors had access to the letters.^{A.1} Moreover, upon seeing the results, the TSE administration did not choose DA-C.

^{A.1}Had the directors seen the letters, the unmatched applicant would not have had a different outcome because she/he submitted only three choices.

Appendix B Data

B.1 Programs' Ranking over All Applicants

As discussed in Appendix A.3, the program directors did not rank all applicants, while a complete order that ranks all applicants by each program is needed for counterfactual analysis. We construct the complete order by the following steps:

- (i) The partial orders of applicants that are submitted by the program directors are always respected in the constructed complete orders.
- (ii) For the two programs with prerequisites, every applicant who does not meet the prerequisites is ranked at the bottom in a descending order of their final grades.
- (iii) For each program, the complete order by applicants' final grades is respected whenever it is not in conflict with the program's submitted partial order. When there is a conflict, we do the following adjustment: Suppose applicant i is moved up in the submitted partial order and is just above applicant i^* . This implies that i 's final grade is lower than i^* 's. In the constructed complete order, i is also just above i^* . We iterate this process for each conflict between the ranking by final grades and the submitted partial order.

B.2 Data Cleaning

B.2.1 Acceptability and Applicant Strategies

The survey data on program acceptability and applicants' submitted true preferences under DA are not always consistent with each other. Because we emphasize to applicants that the survey data is not used in the matching mechanisms, the DA data should have a greater weight. Below, we describe how we make some modifications, but as few as possible, in the survey data to restore consistency.

According to the original survey data, one applicant is matched with an unacceptable program (the third choice in the submitted DA ROL) under the three versions of the DA mechanism and with another under IA (the fourth choice in the submitted DA ROL). Moreover, we observe that the applicant enrolls in the matched program under DA in September 2013. We therefore change the applicants' third and fourth choices to be "possibly acceptable."

Another type of inconsistency occurs when an applicant ranks an unacceptable program above an acceptable or puts a "possibly acceptable" program above a "definitely acceptable" in the submitted DA ROL. There are in total nine applicants who have this inconsistency.

We detail our modifications in Table B.1. The principle is as following: Given an applicant, her submitted ROL is taken as the true preference order; therefore, the ROL should always rank “definitely acceptable” programs above “possibly acceptable” ones which are above “unacceptables.” In other words, program acceptability should never increase when we move down in the ROL. When there is a violation between a pair, say l^k and l^{k+1} such that l^{k+1} 's acceptability is higher than l^k 's, we change l^{k+1} 's acceptability to l^k 's. It turns out that we only need to make one modification like this in each of the nice ROLs (Table B.1).

Table B.1: Modifications to the Survey Data on Program Acceptability

App.	Submitted ROL under DA	Original Program Acceptability							Modified Program Acceptability						
		P1	P2	P3	P4	P5	P6	P7	P1	P2	P3	P4	P5	P6	P7
1	(P3, P1, P4, P7, P2, P6 , P5)	1	-1	1	1	-1	0	0	1	-1	1	1	-1	-1	0
2	(P3, P4, P7, P2 , P5, P6, P1)	-1	1	1	1	-1	-1	0	-1	0	1	1	-1	-1	0
3	(P4, P3, P2, P5, P6, P1 , P7)	0	0	0	1	0	-1	-1	-1	0	0	1	0	-1	-1
4	(P3, P1, P7, P4 , P2, P5, P6)	1	-1	1	1	-1	-1	0	1	-1	1	0	-1	-1	0
5	(P2, P1, P4, P5, P3, P7, P6)	1	1	0	0	0	0	-1	1	1	0	0	0	-1	-1
6	(P2, P5, P3, P1, P4 , P6, P5)	0	1	1	1	0	0	1	0	1	1	0	0	0	1
7	(P4, P5, P1, P2, P7, P3 , P6)	1	1	1	1	1	-1	0	1	1	0	1	1	-1	0
8	(P1, P4, P3, P2, P7 , P5, P6)	1	0	1	1	0	-1	1	1	0	1	1	0	-1	0
9	(P2, P4, P1, P3, P5, P6 , P7)	1	1	1	1	-1	1	-1	1	1	1	1	-1	-1	-1

Notes: Modified program acceptability is in red boldface.

B.2.2 Repeated Choices in DA-C

Under DA-C, the applicant interface on the web site allows one to repeatedly rank the same program multiple times. Indeed, there are two applicants (i and i') following this pattern. We change their strategies in the following way:

$$\begin{aligned}
 (l_i^1, l_i^2, l_i^3, l_i^1, l_i^1, l_i^1) &\rightarrow (l_i^1, l_i^2, l_i^3), \\
 (l_{i'}^1, l_{i'}^2, l_{i'}^3, l_{i'}^4, l_{i'}^4, l_{i'}^5) &\rightarrow (l_{i'}^1, l_{i'}^2, l_{i'}^3, l_{i'}^4, l_{i'}^5).
 \end{aligned}$$

The two applicants do not provide additional motivation letters for repeated choices, and both of them are assigned to their respective first choice under DA-C.

Appendix C Likelihood Function based on the Extended Rank-Order List

This appendix details the derivation of the likelihood function based on the extended rank-order list which is the ROL submitted under DA supplemented by information on program acceptability.

C.1 Notations and Properties

Let:

$$G(x, \lambda, \mu) = \mu \exp(-\lambda \exp(-x)).$$

Note that a Type I extreme value random variable is such that:

$$\Pr(\varepsilon < y) = G(y, 1, 1) = \int_{-\infty}^y \exp(-x) \exp(-\exp(-x)) dx,$$

and:

$$G(x + a, \lambda, \mu) = G(x, \lambda \exp(-a), \mu).$$

We also have:

$$G(x, \lambda_1, \mu_1) G(x, \lambda_2, \mu_2) = G(x, \lambda_1 + \lambda_2, \mu_1 \mu_2),$$

and

$$G(+\infty, \lambda, \mu) = \mu$$

Furthermore:

$$\begin{aligned} \int_{-\infty}^y \exp(-x) \exp(-\exp(-x)) G(x, \lambda, \mu) dx &= \mu \int_{-\infty}^y \exp(-x) \exp(-(1 + \lambda) \exp(-x)) dx, \\ &= \frac{\mu}{1 + \lambda} \exp(-(1 + \lambda) \exp(-y)) \\ &= G(y, 1 + \lambda, \frac{\mu}{1 + \lambda}), \end{aligned}$$

so that:

$$\mathbb{E}_x(G(x, \lambda, \mu) \mathbf{1}\{x \leq y\}) = G(y, 1 + \lambda, \frac{\mu}{1 + \lambda}).$$

These properties are used below.

C.2 The Rank-Ordered Multinomial Logit

To simplify notations, we define $v_j \equiv x\beta_j + \varepsilon_j \equiv b_j + \varepsilon_j$, omitting the index for applicants. Furthermore, consider:

$$\begin{aligned}
\Pr(v_{J-1} > v_J) &= \mathbb{E}_{J-1} \left(\int_{-\infty}^{\varepsilon_{J-1} + b_{J-1} - b_J} \exp(-t) \exp(-\exp(-t)) dt \right) \\
&= \mathbb{E}_{J-1} (G(\varepsilon_{J-1} + b_{J-1} - b_J, 1, 1)), \\
&= \mathbb{E}_{J-1} (G(\varepsilon_{J-1}, \exp(b_J - b_{J-1}), 1)), \\
&= G(+\infty, 1 + \exp(b_J - b_{J-1}), \frac{1}{1 + \exp(b_J - b_{J-1})}) \\
&= \frac{1}{1 + \exp(b_J - b_{J-1})} \\
&= \frac{\exp(b_{J-1})}{\exp(b_{J-1}) + \exp(b_J)}.
\end{aligned}$$

Iterating:

$$\begin{aligned}
&\Pr(v_{J-2} > v_{J-1}, v_{J-1} > v_J) \\
&= \mathbb{E}_{J-2} \left(G \left(\varepsilon_{J-2} + b_{J-2} - b_{J-1}, 1 + \exp(b_J - b_{J-1}), \frac{1}{1 + \exp(b_J - b_{J-1})} \right) \right) \\
&= \mathbb{E}_{J-2} \left(G \left(\varepsilon_{J-2}, \exp(b_{J-1} - b_{J-2}) + \exp(b_J - b_{J-2}), \frac{1}{1 + \exp(b_J - b_{J-1})} \right) \right) \\
&= G \left(+\infty, 1 + \exp(b_{J-1} - b_{J-2}) + \exp(b_J - b_{J-2}), \frac{1}{1 + \exp(b_J - b_{J-1})} \cdot \frac{1}{1 + \exp(b_{J-1} - b_{J-2}) + \exp(b_J - b_{J-2})} \right) \\
&= \frac{\exp(b_{J-2})}{\exp(b_{J-2}) + \exp(b_{J-1}) + \exp(b_J)} \cdot \frac{\exp(b_{J-1})}{\exp(b_{J-1}) + \exp(b_J)},
\end{aligned}$$

which is the usual rank-ordered logit (or exploded logit) formula. By extension:

$$\begin{aligned}
&\Pr(v_{J-K} > v_{J-K+1}, \dots, v_{J-1} > v_J) \\
&= \mathbb{E}_{J-K} (G(\varepsilon_{J-K}, \lambda_{J-K}, \mu_{J-K})) \\
&= G(+\infty, \lambda_{J-K-1}, \mu_{J-K-1}) = \mu_{J-K-1},
\end{aligned}$$

in which $\lambda_j = \exp(-(b_j - b_{j+1}))(1 + \lambda_{j+1})$ and $\lambda_J = 0$. This yields:

$$\lambda_j = \sum_{k>j} \exp(b_k - b_j).$$

Furthermore, $\mu_j = \frac{\mu_{j+1}}{1+\lambda_{j+1}}$ and $\mu_{J-1} = 1$ so that:

$$\mu_j = \prod_{m \geq j+1} \frac{1}{1 + \sum_{k>m} \exp(b_k - b_m)}.$$

Further iteration yields:

$$\Pr(v_1 > v_2, \dots, v_{J-1} > v_J) = \mu_0 = \prod_{m \geq 1} \frac{1}{1 + \sum_{k>m} \exp(b_k - b_m)} = \prod_{m \geq 1} \frac{\exp(b_m)}{\sum_{k \geq m} \exp(b_k)}.$$

C.3 Derivation of Likelihood Function

To further simplify notations, without loss of generality, we let the extended rank-order list L^f be $(1, \dots, \bar{J}, \bar{O}, \bar{J} + 1, \dots, \underline{J} - 1, \underline{O}, \underline{J}, \dots, J)$. We derive the following likelihood:

$$L = \Pr(v_1 \geq \dots \geq v_{\bar{J}} \geq \bar{v}_0 \geq v_{\bar{J}+1} \geq \dots \geq v_{\underline{J}-1} \geq \underline{v}_0 \geq v_{\underline{J}} \geq \dots \geq v_J),$$

where \bar{v}_0 and \underline{v}_0 are the max and min value of outside option. Moreover, by normalization, $\underline{v}_0 = \varepsilon_0$. From the above results, the term at the right of \underline{v}_0 equals:

$$\Pr(\varepsilon_0 > v_{\underline{J}}, \dots, v_{J-1} > v_J) = \mathbb{E}_{\varepsilon_0} G(\varepsilon_0, \lambda^{(0)}, \mu^{(0)}),$$

in which:

$$\lambda^{(0)} = \sum_{k \geq \underline{J}} \exp(b_k), \mu^{(0)} = \prod_{m \geq \underline{J}} \frac{\exp(b_m)}{\sum_{k \geq m} \exp(b_k)}.$$

Nonetheless, the term $\Pr(v_{\bar{J}} \geq \bar{v} \geq v_{\bar{J}+1} \geq \dots \geq v_{\underline{J}-1} \geq \underline{v} \geq v_{\underline{J}} \geq \dots \geq v_J)$ requires a different evaluation from above because ε_0 affects both \bar{v}_0 and \underline{v}_0 .

We start from the case in which there are more than two “possibly acceptable” alternatives and show that it applies as well to the case in which there are fewer “possibly acceptable” alternatives.

C.3.1 Two or More “Possibly Acceptable” Alternatives

Consider that there are h “possibly acceptable” alternatives i.e. $\bar{J} + 1 = \underline{J} - h, \dots, \underline{J} - 1$ are associated to don't know.

A backward induction mechanism. Start the induction from :

$$\begin{aligned}
M(\varepsilon_{\underline{J}-2}) &= G(\varepsilon_0, \lambda^{(0)}, \mu^{(0)}) \int_{\varepsilon_0 - b_{\underline{J}-1}}^{\varepsilon_{\underline{J}-2} + b_{\underline{J}-2} - b_{\underline{J}-1}} f(\varepsilon_{\underline{J}-1}) d\varepsilon_{\underline{J}-1} \\
&= G(\varepsilon_0, \lambda^{(0)}, \mu^{(0)}) (G(\varepsilon_{\underline{J}-2}, \exp(b_{\underline{J}-1} - b_{\underline{J}-2}), 1) - G(\varepsilon_0, \exp(b_{\underline{J}-1}), 1)), \\
&= G(\varepsilon_{\underline{J}-2}, \exp(b_{\underline{J}-1} - b_{\underline{J}-2}), 1) G_0(\varepsilon_0, \lambda^{(0)}, \mu^{(0)}) - G(\varepsilon_0, \exp(b_{\underline{J}-1}) + \lambda^{(0)}, \mu^{(0)}),
\end{aligned}$$

to posit that, by analogy:

$$\begin{aligned}
&M(\varepsilon_{\underline{J}-1-j}) \\
&= \int_{\varepsilon_0 - b_{\underline{J}-j}}^{\varepsilon_{\underline{J}-1-j} + b_{\underline{J}-1-j} - b_{\underline{J}-j}} f(\varepsilon_{\underline{J}-j}) d\varepsilon_{\underline{J}-j} \dots \left(\int_{\varepsilon_0 - b_{\underline{J}-1}}^{\varepsilon_{\underline{J}-2} + b_{\underline{J}-2} - b_{\underline{J}-1}} f(\varepsilon_{\underline{J}-1}) d\varepsilon_{\underline{J}-1} \cdot G(\varepsilon_0, \lambda^{(0)}, \mu^{(0)}) \right), \\
&= \sum_{p=1}^j \sigma_{\{p=1\}} G(\varepsilon_{\underline{J}-1-j}, \lambda_{p,j}^{(\varepsilon)}, \mu_{p,j}^{(\varepsilon)}) K_{p-1}(\varepsilon_0) - K_j(\varepsilon_0),
\end{aligned}$$

in which:

$$\sigma_{\{p=1\}} = \begin{cases} 1 & \text{if } p = 1, \\ -1 & \text{otherwise.} \end{cases}$$

The initialization is such that:

$$\begin{aligned}
\lambda_{1,1}^{(\varepsilon)} &= \exp(b_{\underline{J}-1} - b_{\underline{J}-2}), \mu_{1,1}^{(\varepsilon)} = 1, \\
K_0(\varepsilon_0) &= G_0(\varepsilon_0, \lambda^{(0)}, \mu^{(0)}), \\
K_1(\varepsilon_0) &= G(\varepsilon_0, \exp(b_{\underline{J}-1}) + \lambda^{(0)}, \mu^{(0)}).
\end{aligned} \tag{A.1}$$

By induction, for $j \geq 1$:

$$\begin{aligned}
M(\varepsilon_{\underline{J}-2-j}) &= \int_{\varepsilon_0 - b_{\underline{J}-1-j}}^{\varepsilon_{\underline{J}-2-j} + b_{\underline{J}-2-j} - b_{\underline{J}-1-j}} f(\varepsilon_{\underline{J}-1-j}) M(\varepsilon_{\underline{J}-1-j}) d\varepsilon_{\underline{J}-1-j} \\
&= \sum_{p=1}^j \sigma_{\{p=1\}} \left(\begin{array}{c} G(\varepsilon_{\underline{J}-2-j} + b_{\underline{J}-2-j} - b_{\underline{J}-1-j}, 1 + \lambda_{p,j}^{(\varepsilon)}, \frac{\mu_{p,j}^{(\varepsilon)}}{1 + \lambda_{p,j}^{(\varepsilon)}}) \\ -G(\varepsilon_0 - b_{\underline{J}-1-j}, 1 + \lambda_{p,j}^{(\varepsilon)}, \frac{\mu_{p,j}^{(\varepsilon)}}{1 + \lambda_{p,j}^{(\varepsilon)}}) \end{array} \right) K_{p-1}(\varepsilon_0) \\
&\quad - \left(\begin{array}{c} G(\varepsilon_{\underline{J}-2-j} + b_{\underline{J}-2-j} - b_{\underline{J}-1-j}, 1, 1) \\ -G(\varepsilon_0 - b_{\underline{J}-1-j}, 1, 1) \end{array} \right) K_j(\varepsilon_0) \\
&= \sum_{p=1}^j \sigma_{\{p=1\}} \left(\begin{array}{c} G(\varepsilon_{\underline{J}-2-j}, \exp(b_{\underline{J}-1-j} - b_{\underline{J}-2-j})(1 + \lambda_{p,j}^{(\varepsilon)}), \frac{\mu_{p,j}^{(\varepsilon)}}{1 + \lambda_{p,j}^{(\varepsilon)}}) \\ -G(\varepsilon_0, \exp(b_{\underline{J}-1-j})(1 + \lambda_{p,j}^{(\varepsilon)}), \frac{\mu_{p,j}^{(\varepsilon)}}{1 + \lambda_{p,j}^{(\varepsilon)}}) \end{array} \right) K_{p-1}(\varepsilon_0) \\
&\quad - \left(\begin{array}{c} G(\varepsilon_{\underline{J}-2-j}, \exp(b_{\underline{J}-1-j} - b_{\underline{J}-2-j}), 1) \\ -G(\varepsilon_0, \exp(b_{\underline{J}-1-j}), 1) \end{array} \right) K_j(\varepsilon_0).
\end{aligned}$$

which we equalize to:

$$M(\varepsilon_{\underline{J}-2-j}) = \sum_{p=1}^{j+1} \sigma_{\{p=1\}} G(\varepsilon_{\underline{J}-2-j}, \lambda_{p,j+1}^{(\varepsilon)}, \mu_{p,j+1}^{(\varepsilon)}) K_{p-1}(\varepsilon_0) - K_{j+1}(\varepsilon_0).$$

Identifying first the terms a function of $\varepsilon_{\underline{J}-2-j}$, we have, for $0 < p \leq j$:

$$\begin{aligned}
\lambda_{p,j+1}^{(\varepsilon)} &= \exp(b_{\underline{J}-1-j} - b_{\underline{J}-2-j})(1 + \lambda_{p,j}^{(\varepsilon)}), \\
\mu_{p,j+1}^{(\varepsilon)} &= \frac{\mu_{p,j}^{(\varepsilon)}}{1 + \lambda_{p,j}^{(\varepsilon)}},
\end{aligned}$$

and

$$\begin{aligned}
\lambda_{j+1,j+1}^{(\varepsilon)} &= \exp(b_{\underline{J}-1-j} - b_{\underline{J}-2-j}), \\
\mu_{j+1,j+1}^{(\varepsilon)} &= 1.
\end{aligned}$$

This yields for $0 < p \leq j$:

$$1 + \lambda_{p,j}^{(\varepsilon)} = \frac{1}{C(b_{\underline{J}-1-j}; (b_{\underline{J}-k})_{p \leq k \leq j+1})},$$

and, using the convention that the product $\prod_{p < m \leq j+1}$ running over an empty set is equal to 1,

$$\mu_{p,j+1}^{(\varepsilon)} = \prod_{p < m \leq j+1} C(b_{\underline{J}-m}; (b_{\underline{J}-k})_{p \leq k \leq m}).$$

Finally, identifying the terms in ε_0 :

$$\begin{aligned}
& K_{j+1}(\varepsilon_0) \\
&= \sum_{p=1}^j \sigma_{\{p=1\}} G(\varepsilon_0, \exp(b_{\underline{j}-1-j})(1 + \lambda_{p,j}^{(\varepsilon)}), \frac{\mu_{p,j}^{(\varepsilon)}}{1 + \lambda_{p,j}^{(\varepsilon)}}) K_{p-1}(\varepsilon_0) - G(\varepsilon_0, \exp(b_{\underline{j}-1-j}), 1) K_j(\varepsilon_0) \\
&= \sum_{p=1}^j \sigma_{\{p=1\}} G(\varepsilon_0, \exp(b_{\underline{j}-1-j})(1 + \lambda_{p,j}^{(\varepsilon)}), \mu_{p,j+1}^{(\varepsilon)}) K_{p-1}(\varepsilon_0) - G(\varepsilon_0, \exp(b_{\underline{j}-1-j}), 1) K_j(\varepsilon_0).
\end{aligned}$$

Let us postulate that:

$$K_j(\varepsilon_0) = G(\varepsilon_0, \sum_{1 \leq k \leq j} \exp(b_{\underline{j}-k}) + \lambda^{(0)}, \mu_j^{(0)}).$$

Because of equation (A.1), this property is true for $j = 0$ and $j = 1$ with:

$$\mu_0^{(0)} = \mu_1^{(0)} = \mu^{(0)}.$$

Suppose that it is true for j and write:

$$\begin{aligned}
& K_{j+1}(\varepsilon_0) \\
&= \sum_{p=1}^j \sigma_{\{p=1\}} G(\varepsilon_0, \exp(b_{\underline{j}-1-j})(1 + \lambda_{p,j}^{(\varepsilon)}), \mu_{p,j+1}^{(\varepsilon)}) G(\varepsilon_0, \sum_{1 \leq k \leq p-1} \exp(b_{\underline{j}-k}) + \lambda^{(0)}, \mu_{p-1}^{(0)}) \\
&\quad - G(\varepsilon_0, \exp(b_{\underline{j}-1-j}), 1) G(\varepsilon_0, \sum_{1 \leq k \leq j} \exp(b_{\underline{j}-k}) + \lambda^{(0)}, \mu_j^{(0)}).
\end{aligned}$$

By the above:

$$\exp(b_{\underline{j}-1-j})(1 + \lambda_{p,j}^{(\varepsilon)}) = \frac{\exp(b_{\underline{j}-1-j})}{C(b_{\underline{j}-1-j}; (b_{\underline{j}-k})_{p \leq k \leq j+1})} = \sum_{p \leq k \leq j+1} \exp(b_{\underline{j}-k}),$$

so that:

$$\begin{aligned}
K_{j+1}(\varepsilon_0) &= \sum_{p=1}^j \sigma_{\{p=1\}} G(\varepsilon_0, \sum_{p \leq k \leq j+1} \exp(b_{\underline{j}-k}) + \sum_{1 \leq k \leq p-1} \exp(b_{\underline{j}-k}) + \lambda^{(0)}, \mu_{p,j+1}^{(\varepsilon)} \mu_{p-1}^{(0)}) \\
&\quad - G(\varepsilon_0, \sum_{1 \leq k \leq j+1} \exp(b_{\underline{j}-k}) + \lambda^{(0)}, \mu_j^{(0)}) \\
&= G(\varepsilon_0, \sum_{1 \leq k \leq j+1} \exp(b_{\underline{j}-k}) + \lambda^{(0)}, \mu_{j+1}^{(0)}),
\end{aligned}$$

since second arguments of all G elements are equal. Additionally:

$$\begin{aligned}\mu_{j+1}^{(0)} &= \sum_{p=1}^j \sigma_{\{p=1\}} \mu_{p,j+1}^{(\varepsilon)} \mu_{p-1}^{(0)} - \mu_j^{(0)}, \\ &= \sum_{p=1}^j \sigma_{\{p=1\}} \prod_{p < m \leq j+1} C(b_{\underline{J}-m}; (b_{\underline{J}-k})_{p \leq k \leq m}) \mu_{p-1}^{(0)} - \mu_j^{(0)}.\end{aligned}\tag{A.2}$$

As a summary, we have:

$$\begin{aligned}M(\varepsilon_{\underline{J}-1-j}) &= \sum_{p=1}^j \sigma_{\{p=1\}} G(\varepsilon_{\underline{J}-1-j}, \lambda_{p,j}^{(\varepsilon)}, \mu_{p,j}^{(\varepsilon)}) G(\varepsilon_0, \sum_{1 \leq k \leq p-1} \exp(b_{\underline{J}-k}) + \lambda^{(0)}, \mu_{p-1}^{(0)}) \\ &\quad - G(\varepsilon_0, \sum_{1 \leq k \leq j} \exp(b_{\underline{J}-k}) + \lambda^{(0)}, \mu_j^{(0)}).\end{aligned}$$

Integration As in the previous section, we have:

$$\begin{aligned}G^*(\varepsilon_0) &= \int_{\varepsilon_0 - b_{\bar{J}+1}}^{\varepsilon_0 + \bar{\eta} - b_{\bar{J}+1}} M(\varepsilon_{\bar{J}+1}) f(\varepsilon_{\bar{J}+1}) d\varepsilon_{\bar{J}+1} \\ &= \int_{\varepsilon_0 - b_{\bar{J}+1}}^{\varepsilon_0 + \bar{\eta} - b_{\bar{J}+1}} \left[\sum_{p=1}^{h-1} \sigma_{\{p=1\}} G(\varepsilon_{\bar{J}+1}, \lambda_{p,h-1}^{(\varepsilon)}, \mu_{p,h-1}^{(\varepsilon)}) G(\varepsilon_0, \sum_{1 \leq k \leq p-1} \exp(b_{\underline{J}-k}) + \lambda^{(0)}, \mu_{p-1}^{(0)}) \right. \\ &\quad \left. - G(\varepsilon_0, \sum_{1 \leq k \leq h-1} \exp(b_{\underline{J}-k}) + \lambda^{(0)}, \mu_{h-1}^{(0)}) \right] f(\varepsilon_{\bar{J}+1}) d\varepsilon_{\bar{J}+1},\end{aligned}$$

in which we have set $\underline{J} - 1 - j = \bar{J} + 1$ i.e. $j = \underline{J} - 2 - \bar{J} = h - 1$ and $h \geq 2$. Note that $\underline{J} - h = \bar{J} + 1$. Then:

$$\begin{aligned}G^*(\varepsilon_0) &= \sum_{p=1}^{h-1} \sigma_{\{p=1\}} \left(\begin{array}{c} G(\varepsilon_0 + \bar{\eta} - b_{\bar{J}+1}, 1 + \lambda_{p,h-1}^{(\varepsilon)}, \frac{\mu_{p,h-1}^{(\varepsilon)}}{1 + \lambda_{p,h-1}^{(\varepsilon)}}) \\ -G(\varepsilon_0 - b_{\bar{J}+1}, 1 + \lambda_{p,h-1}^{(\varepsilon)}, \frac{\mu_{p,h-1}^{(\varepsilon)}}{1 + \lambda_{p,h-1}^{(\varepsilon)}}) \end{array} \right) G(\varepsilon_0, \sum_{1 \leq k \leq p-1} \exp(b_{\underline{J}-k}) + \lambda^{(0)}, \mu_{p-1}^{(0)}) \\ &\quad - \left(\begin{array}{c} G(\varepsilon_0 + \bar{\eta} - b_{\bar{J}+1}, 1, 1) \\ -G(\varepsilon_0 - b_{\bar{J}+1}, 1, 1) \end{array} \right) G(\varepsilon_0, \sum_{1 \leq k \leq h-1} \exp(b_{\underline{J}-k}) + \lambda^{(0)}, \mu_{h-1}^{(0)}) \\ &= \sum_{p=1}^{h-1} \sigma_{\{p=1\}} \left(\begin{array}{c} G(\varepsilon_0, \exp(b_{\bar{J}+1} - \bar{\eta})(1 + \lambda_{p,h-1}^{(\varepsilon)}) + \sum_{1 \leq k \leq p-1} \exp(b_{\underline{J}-k}) + \lambda^{(0)}, \mu_{p,h}^{(\varepsilon)} \mu_{p-1}^{(0)}) \\ -G(\varepsilon_0, \exp(b_{\bar{J}+1})(1 + \lambda_{p,h-1}^{(\varepsilon)}) + \sum_{1 \leq k \leq p-1} \exp(b_{\underline{J}-k}) + \lambda^{(0)}, \mu_{p,h}^{(\varepsilon)} \mu_{p-1}^{(0)}) \end{array} \right) \\ &\quad - \left(\begin{array}{c} G(\varepsilon_0, \exp(b_{\bar{J}+1} - \bar{\eta}) + \sum_{1 \leq k \leq h-1} \exp(b_{\underline{J}-k}) + \lambda^{(0)}, \mu_{h-1}^{(0)}) \\ -G(\varepsilon_0, \exp(b_{\bar{J}+1}) + \sum_{1 \leq k \leq h-1} \exp(b_{\underline{J}-k}) + \lambda^{(0)}, \mu_{h-1}^{(0)}) \end{array} \right)\end{aligned}$$

By the same principle as above:

$$\exp(b_{\bar{J}+1})(1 + \lambda_{p,h-1}^{(\varepsilon)}) = \frac{\exp(b_{\bar{J}+1})}{C(b_{\bar{J}+1}; (b_{\underline{J}-k})_{p \leq k \leq h})} = \sum_{p \leq k \leq h} \exp(b_{\underline{J}-k}),$$

so that the second and fourth terms can be combined to obtain:

$$\begin{aligned} & G(\varepsilon_0, \sum_{1 \leq k \leq h} \exp(b_{\underline{J}-k}) + \lambda^{(0)}, \sum_{p=1}^{h-1} \sigma_{\{p=1\}} \mu_{p,h}^{(\varepsilon)} \mu_{p-1}^{(0)}) - G(\varepsilon_0, \sum_{1 \leq k \leq h} \exp(b_{\underline{J}-k}) + \lambda^{(0)}, \mu_{h-1}^{(0)}) \\ &= G(\varepsilon_0, \sum_{1 \leq k \leq h} \exp(b_{\underline{J}-k}) + \lambda^{(0)}, \sum_{p=1}^{h-1} \sigma_{\{p=1\}} \mu_{p,h}^{(\varepsilon)} \mu_{p-1}^{(0)} - \mu_{h-1}^{(0)}) \\ &= G(\varepsilon_0, \sum_{1 \leq k \leq h} \exp(b_{\underline{J}-k}) + \lambda^{(0)}, \mu_h^{(0)}), \end{aligned}$$

where we use the definition of $\mu_h^{(0)}$ from equation (A.2).

Because:

$$\exp(b_{\bar{J}+1} - \bar{\eta})(1 + \lambda_{p,h-1}^{(\varepsilon)}) = \exp(-\bar{\eta}) \frac{\exp(b_{\bar{J}+1})}{C(b_{\bar{J}+1}; (b_{\underline{J}-k})_{p \leq k \leq h})} = \exp(-\bar{\eta}) \sum_{p \leq k \leq h} \exp(b_{\underline{J}-k}),$$

we can then write:

$$\begin{aligned} G^*(\varepsilon_0) &= \sum_{p=1}^{h-1} \sigma_{\{p=1\}} G(\varepsilon_0, \exp(-\bar{\eta}) \sum_{p \leq k \leq h} \exp(b_{\underline{J}-k}) + \sum_{1 \leq k \leq p-1} \exp(b_{\underline{J}-k}) + \lambda^{(0)}, \mu_{p,h}^{(\varepsilon)} \mu_{p-1}^{(0)}) \\ &\quad - G(\varepsilon_0, \sum_{1 \leq k \leq h} \exp(b_{\underline{J}-k}) + \lambda^{(0)}, \mu_h^{(0)}) \\ &\quad - G(\varepsilon_0, \exp(b_{\bar{J}+1} - \bar{\eta}) + \sum_{1 \leq k \leq h-1} \exp(b_{\underline{J}-k}) + \lambda^{(0)}, \mu_{h-1}^{(0)}) \end{aligned}$$

Final Derivation We end up with considering that:

$$\Pr(v_{\bar{J}} \geq \bar{v}_0 \geq v_{\bar{J}+1} \geq \dots \geq v_{\underline{J}-1} \geq \underline{v}_0 \geq v_{\underline{J}} \geq \dots \geq v_J) = \mathbb{E}_{\bar{J}} (H^{(h)}(\varepsilon_{\bar{J}})),$$

in which

$$H^{(h)}(\varepsilon_{\bar{J}}) = \int_{-\infty}^{\varepsilon_{\bar{J}} + b_{\bar{J}} - \bar{\eta}} f(\varepsilon_0) G^*(\varepsilon_0) d\varepsilon_0$$

and therefore:

$$H^{(h)}(\varepsilon_{\bar{J}}) = \sum_{p=1}^{h+1} \sigma_{\{p=1\}} G(\varepsilon_{\bar{J}}, \bar{\lambda}_p^{(0)}, \bar{\mu}_p^{(0)}),$$

in which:

$$\bar{\lambda}_p^{(0)} = \exp(\bar{\eta} - b_{\bar{J}})(1 + \exp(-\bar{\eta})) \sum_{p \leq k \leq h} \exp(b_{\underline{J}-k}) + \sum_{1 \leq k \leq p-1} \exp(b_{\underline{J}-k}) + \lambda^{(0)}$$

because $\underline{J} - h = \bar{J} + 1$. We can then write:

$$\bar{\lambda}_p^{(0)} = \exp(-b_{\bar{J}})(\exp(\bar{\eta}) + \sum_{p \leq k \leq h} \exp(b_{\underline{J}-k}) + \sum_{1 \leq k \leq p-1} \exp(b_{\underline{J}-k} + \bar{\eta}) + \exp(\bar{\eta})\lambda^{(0)})$$

and:

$$1 + \bar{\lambda}_p^{(0)} = \frac{1}{C((b_{\bar{J}}; \bar{\eta}, (b_{\underline{J}-k})_{p \leq k \leq h}), (b_{\underline{J}-k} + \bar{\eta})_{1 \leq k \leq p-1}, (b_k + \bar{\eta})_{k \geq \underline{J}})}.$$

We also have for $1 \leq p < h$:

$$\begin{aligned} \bar{\mu}_p^{(0)} &= \frac{1}{(1 + \exp(-\bar{\eta}) \sum_{p \leq k \leq h} \exp(b_{\underline{J}-k}) + \sum_{1 \leq k \leq p-1} \exp(b_{\underline{J}-k}) + \lambda^{(0)}) \mu_{p,h}^{(\varepsilon)} \mu_{p-1}^{(0)}} \\ &= C(\bar{\eta}; \bar{\eta}, (b_{\underline{J}-k})_{p \leq k \leq h}, (b_{\underline{J}-k} + \bar{\eta})_{1 \leq k \leq p-1}, (b_k + \bar{\eta})_{k \geq \underline{J}}) \\ &\quad \times \prod_{p < m \leq h} C(b_{\underline{J}-m}; (b_{\underline{J}-k})_{p \leq k \leq m}) \mu_{p-1}^{(0)}. \end{aligned}$$

which is also valid for $p = h$ because:

$$\begin{aligned} \bar{\mu}_h^{(0)} &= \frac{1}{1 + \exp(b_{\bar{J}+1} - \bar{\eta}) + \sum_{1 \leq k \leq h-1} \exp(b_{\underline{J}-k}) + \lambda^{(0)} \mu_{h-1}^{(0)}}, \\ &= C(\bar{\eta}; \bar{\eta}, (b_{\underline{J}-k})_{p \leq k \leq h}, (b_{\underline{J}-k} + \bar{\eta})_{1 \leq k \leq p-1}, (b_k + \bar{\eta})_{k \geq \underline{J}}) \mu_{h-1}^{(0)}. \end{aligned}$$

When $p = h + 1$, we have:

$$\bar{\mu}_{h+1}^{(0)} = \frac{1}{1 + \sum_{1 \leq k \leq h} \exp(b_{\underline{J}-k}) + \lambda^{(0)} \mu_h^{(0)}} = C(0; 0, (b_k)_{k \geq \bar{J}+1}) \mu_h^{(0)}.$$

To derive the likelihood $\Pr(v_1 \geq \dots \geq v_{\bar{J}} \geq \bar{v}_0 \geq v_{\bar{J}+1} \geq \dots \geq v_{\underline{J}-1} \geq \underline{v}_0 \geq v_{\underline{J}} \geq \dots \geq v_J)$, we reconsider the backward induction as in the previous sections:

$$\bar{\lambda}_{j,p} = \exp(-(b_j - b_{j+1}))(1 + \bar{\lambda}_{j+1,p}), \bar{\mu}_{j,p} = \frac{\bar{\mu}_{j+1,p}}{1 + \bar{\lambda}_{j+1,p}},$$

so that:

$$1 + \bar{\lambda}_{j,p} = \frac{1}{C((b_j; (b_k)_{j \leq k \leq \bar{J}}, \bar{\eta}, (b_{\underline{J}-k})_{p \leq k \leq h}, (b_{\underline{J}-k} + \bar{\eta})_{1 \leq k \leq p-1}, (b_k + \bar{\eta})_{j \geq \underline{J}})}),$$

and therefore if $p \leq h$:

$$\begin{aligned} \bar{\mu}_{0,p} &= \left[\prod_{1 \leq m \leq \bar{J}} C(b_m; ((b_k)_{m \leq k \leq \bar{J}}, \bar{\eta}, (b_{\underline{J}-k})_{p \leq k \leq h}, (b_{\underline{J}-k} + \bar{\eta})_{1 \leq k \leq p-1}, (b_k + \bar{\eta})_{k \geq \underline{J}}) \right] \\ &\quad \times C(\bar{\eta}; \bar{\eta}, (b_{\underline{J}-k})_{p \leq k \leq h}, (b_{\underline{J}-k} + \bar{\eta})_{1 \leq k \leq p-1}, (b_k + \bar{\eta})_{k \geq \underline{J}}) \\ &\quad \times \prod_{p < m \leq h} C(b_{\underline{J}-m}; (b_{\underline{J}-k})_{p \leq k \leq m}) \mu_{p-1}^{(0)}, \end{aligned}$$

and if $p = h + 1$:

$$\bar{\mu}_{0,h+1} = \left[\prod_{1 \leq m \leq \bar{J}} C(b_m; ((b_k)_{m \leq k \leq \bar{J}}), \bar{\eta}, (b_{\underline{J}-k} + \bar{\eta})_{1 \leq k \leq h}, (b_k + \bar{\eta})_{k \geq \underline{J}}) \right] C(0; 0, (b_k)_{k \geq \bar{J}+1}) \mu_h^{(0)},$$

Finally:

$$\Pr(v_1 \geq \dots \geq v_{\bar{J}} \geq \bar{v}_0 \geq v_{\bar{J}+1} \geq \dots \geq v_{\underline{J}-1} \geq \underline{v}_0 \geq v_{\underline{J}} \geq \dots \geq v_J) = \sum_{p=1}^{h+1} \sigma_{\{p=1\}} \bar{\mu}_{0,p}.$$

Possible mistakes are checked by summing over all possible cases under different scenarios of variables and coefficients and assessing that the sum of probabilities is equal to one. In an additional note which is available upon request, cases for $h = 0, 1$ and 2 are written directly, and we have checked that the formula above applies as well. Cases $h = 0$ and $h = 1$ are derived below.

h=0 We have by application of the above a single term $p = h + 1 = 1$:

$$\Pr(v_1 \geq \dots \geq v_{\bar{J}} \geq \bar{v}_0 \geq \underline{v}_0 \geq v_{\underline{J}} \geq \dots \geq v_J) = \bar{\mu}_{0,1},$$

in which:

$$\bar{\mu}_{0,1} = \left[\prod_{1 \leq m \leq \bar{J}} C(b_m; ((b_k)_{m \leq k \leq \bar{J}}), \bar{\eta}, (b_k + \bar{\eta})_{k \geq \underline{J}}) \right] C(0; 0, (b_k)_{k \geq \bar{J}+1}) \prod_{m \geq \underline{J}} C(b_m; (b_k)_{k \geq m}),$$

because

$$\mu_0^{(0)} = \mu^{(0)} = \prod_{m \geq \underline{J}} C(b_m; (b_k)_{k \geq m}).$$

h=1 We have by application of the above two terms $p = 1$ and $p = h + 1 = 2$:

$$\Pr(v_1 \geq \dots \geq v_{\bar{J}} \geq \bar{v}_0 \geq v_{\bar{J}+1} \geq \underline{v}_0 \geq v_{\underline{J}} \geq \dots \geq v_J) = \bar{\mu}_{0,1} - \bar{\mu}_{0,2},$$

in which:

$$\begin{aligned}\bar{\mu}_{0,1} &= \left[\prod_{1 \leq m \leq \bar{J}} C(b_m; ((b_k)_{m \leq k \leq \bar{J}}), \bar{\eta}, b_{\bar{J}+1}, (b_k + \bar{\eta})_{k \geq \bar{J}}) \right] \\ &\quad \times C(\bar{\eta}; \bar{\eta}, b_{\bar{J}+1}, (b_k + \bar{\eta})_{k \geq \bar{J}}) \prod_{m \geq \bar{J}} C(b_m; (b_k)_{k \geq m}), \\ \bar{\mu}_{0,2} &= \left[\prod_{1 \leq m \leq \bar{J}} C(b_m; ((b_k)_{m \leq k \leq \bar{J}}), \bar{\eta}, (b_k + \bar{\eta})_{k \geq \bar{J}+1}) \right] \\ &\quad \times C(0; 0, (b_k)_{k \geq \bar{J}+1}) \prod_{m \geq \bar{J}} C(b_m; (b_k)_{k \geq m}),\end{aligned}$$

because $\mu_1^{(0)} = \mu^{(0)} = \prod_{m \geq \bar{J}} C(b_m; (b_k)_{k \geq m})$.

Appendix D Counterfactual Analysis

D.1 Solving Equilibrium

This appendix specifies the algorithm to numerically solve an equilibrium under DA-T of different degrees or DA-C with various marginal costs.

As explained in Section 5.5, we begin with drawing $M (= 4000)$ samples of size I (number of applicants) in which random shocks to applicant preferences ($\varepsilon_i^{(m)}$), to applicants' final grades ($\nu_i^{(m)}$), and to program preferences ($\xi_{i,j}^{(m)}$) are generated. Using these simulated shocks as well as estimated coefficients and observed characteristics, we compute applicant and program preferences, both cardinal and ordinal; moreover, program preferences are translated into rankings over applicants from 1 to I denoted by $r_{i,j}^{(m)}$, with larger $r_{i,j}^{(m)}$ indicating being more preferred.

D.1.1 Solving Bayesian Nash Equilibrium under DA-T

For DA-T of degree $K \in \{1, 2, \dots, 6\}$, we adopt an iterative process defined by index t ($= 0, 1, \dots$):

- *Initialization*: We first initialize the cutoff distribution, Φ_0 . For each simulation sample m :
 - (i) We let the matching outcome $\mu_{(m,0)}$ (artificially) be $\hat{\mu}_{DA}$, where $\hat{\mu}_{(DA)}$ is the applicant-optimal stable matching observed in our real data (i.e., the outcome from DA as calculated in Section 4);
 - (ii) We then calculate the distribution of cutoffs Φ_0 by setting the cutoff of each program in each simulation sample,

$$\kappa_j^{(m,0)} = \begin{cases} \min_{i \in \mu_{(m,0)}^{-1}(j)} r_{i,j}^{(m)} & \text{if } |\mu_{(m,0)}^{-1}(j)| \geq q_j, \\ 0 & \text{if } |\mu_{(m,0)}^{-1}(j)| < q_j. \end{cases}$$

Because $\mu_{(m,0)} = \hat{\mu}_{(DA)}$ for all m , $\kappa_j^{(m,0)}$ is constant across all simulation samples. This leads to a degenerate distribution, Φ_0 .

We then calculate the simulated preferences of applicants and programs:

- (i) We compute applicants' cardinal and therefore ordinal preferences by inputting the simulated shocks $\varepsilon_i^{(m)}$, estimated coefficients, and observed characteristics into the utility functions specified in equations (3) and (4). Note that shocks to

applicant preferences $(\varepsilon_i^{(m)})$ are drawn **unconditionally** on ordinal preferences observed in the sample.

- (ii) We also identify all ROLs $\{L_{i,n}^{(m)}\}_{n=1,\dots,N_i}$ for all i such that (a) $L_{i,n}^{(m)}$ ranks no more than K programs; (b) $L_{i,n}^{(m)}$ is a partial order of the true preference order, i.e., the programs included in $L_{i,n}^{(m)}$ are ranked according to the true ordinal preferences; and (c) $L_{i,n}^{(m)}$ does not include any unacceptable or unqualified program. We only need to focus on these ROLs to find an optimal ROL because all other ROLs are strictly dominated (see Section 2).
- (iii) We attribute to each program the same capacity as in the experiment and compute its “cardinal” preferences over applicants by plugging the simulated shocks $(\nu_i^{(m)}$ and $\xi_{i,j}^{(m)})$, estimated coefficients, and observed first-semester grades into the grade function (equation 7) and the latent score function (equation 6). These preferences are then translated into rankings over applicants in increasing order of preference from 1 to I , denoted by $r_{i,j}^{(m)}$.
- *Iteration:* For $t \in \{1, 2, \dots\}$, given $\mu_{m,t-1}$ and $(\kappa_j^{(m,t-1)})_{j \in \mathcal{J}}$ for all m (i.e., Φ_{t-1}), iteration t goes through the following steps to find the optimal ROL for each applicant:

- (i) Define $A_i^{(m,t)} = \{j \in \mathcal{J} : r_{i,j}^{(m)} \geq \kappa_j^{(m,t-1)}\}$ for each i and each m . That is, $A_i^{(m,t)}$ is the set of programs that have lower cutoffs in iteration $t - 1$ than i 's ranks in sample m .
- (ii) Compute for each $L_{i,n}^{(m)}$ the expected utility:^{A.2}

$$W_{i,n}^{(m,t)} = \frac{1}{M} \sum_{l=1}^M \left[\sum_{j=1}^J \mathbb{1}_{\{j \text{ is } L_{i,n}^{(m)}\text{-preferred in } A_i^{(l,t)}\}} E_{\eta} \max\{x_i \beta_j + \varepsilon_{i,j}^{(m)}, \eta + \varepsilon_{i,j}^{(0)}\} \right], \quad (\text{A.3})$$

in which the expectation is taken with respect to the ex-post shock, η , on the outside option and “ j is $L_{i,n}^{(m)}$ -preferred in $A_i^{(l,t)}$ ” is defined such that j is ranked higher than j' for all $j' \in L_{i,n}^{(m)} \cap A_i^{(l,t)}$.

- (iii) Find the optimal ROL $L_{i,*}^{(m,t)} = \arg \max_{n=1,\dots,N_i} W_{i,n}^{(m,t)}$. If there are ties, we choose the unique longest list among them. Given Assumption 3, there is always a positive marginal benefit to apply to more programs, as long as the total number applications is not more than K . Tie occurs due to the discrete number of simulations.

^{A.2}This calculation amounts to assuming that i plays against a distribution of others' actions (which is approximated by their actions in iteration $t - 1$ in all samples). Notice that i is present in every sample, although ideally one could take i out. Besides, we use sample m also notwithstanding the introduction of spurious correlation. These two simplifications speed up the computation, while the error introduced is either $O(1/n)$ or $O(1/M)$ and is therefore negligible.

We then calculate the matching outcome in each sample and update the cutoff distribution:

- (i) With $L_{i,*}^{(m,t)}$ and $r_{i,j}^{(m)}$ for all applicants and programs, we run the DA algorithm to find the new matching $\mu^{(m,t)}$ in each simulation sample.
- (ii) We calculate the new cutoffs $\kappa_j^{(m,t)}$ for all j based on $\mu^{(m,t)}$ and the difference between two distributions of cutoffs:

$$d^{(t)} = \frac{1}{M} \sum_{j,m} \left(\kappa_j^{(m,t)} - \kappa_j^{(m,t-1)} \right)^2$$

- *Stopping rule:* Stop when $d^{(t)}$ is small enough.

Suppose that we stop at $t = t_f$. This leads to a joint distribution of cutoffs across simulation samples, $\{\kappa_j^{(m,t_f)}\}$ for all j and m , which is defined as Φ_e in the main text and “describes” an equilibrium under DA-T-K.

Counterfactual analysis of welfare. In counterfactual analysis, we use a **different set of simulation samples** and also need to find best response for each applicant with given cardinal preferences. Recall that ordinal preferences of applicants and programs are fixed at the value observed in the experimental data (cf. Table 10). However, across the simulation samples for counterfactual analysis, applicant cardinal preferences, as simulated in Appendix D.2, can differ.

Given the simulated cardinal preferences in a simulation sample ($x_i\beta_j + \varepsilon_{i,j}$ for all j), i ’s best response under DA-T-K is found through the following steps:

- (i) Given the equilibrium cutoff distribution solved above, $\{\kappa_j^{(m,t_f)}\}$ for all j and m , we define $A_i^{(m,t_f)} = \{j \in \mathcal{J} : r_{i,j}^{(m)} \geq \kappa_j^{(m,t_f)}\}$ for each i .
- (ii) We identify all ROLs $\{L_{i,n}^{(m)}\}_{n=1,\dots,N_i}$ for all i such that (a) $L_{i,n}^{(m)}$ ranks no more than K programs; (b) $L_{i,n}^{(m)}$ is a partial order of the true preference order; and (c) $L_{i,n}^{(m)}$ does not include any unacceptable or unqualified program.
- (iii) Compute for each $L_{i,n}^m$ the expected utility:

$$W_{i,n}^{(t_f)} = \frac{1}{M} \sum_{m=1}^M \left[\sum_{j=1}^J \mathbb{1}_{\{j \text{ is } L_{i,n}^m\text{-preferred in } A_i^{(m,t_f)}\}} E_{\eta} \max\{x_i\beta_j + \varepsilon_{i,j}^{(m)}, \eta + \varepsilon_{i,j}^{(0)}\} \right],$$

where “ j is $L_{i,n}^m$ -preferred in $A_i^{(m,t_f)}$ ” is defined such that j is ranked higher than j' for all $j' \in L_{i,n}^m \cap A_i^{(m,t_f)}$.

- (iv) i 's best response is the optimal ROL $L_{i,*}^m = \arg \max_{n=1,\dots,N_i} W_{i,n}^{(t_f)}$. If there are ties, we choose the unique longest list.

D.1.2 Solving Bayes-Nash equilibrium under DA-C

Similar to DA-T, equilibrium under DA-C also has to be solved as a fixed point. The application cost has the following form in the counterfactual analysis:

$$C(L) = \begin{cases} 0 & \text{if } |L| \leq 1, \\ c|L| & \text{if } |L| > 1; \end{cases}$$

in which c is the constant marginal cost for second and later choices; and c can take one of the 11 values in $\{10^{-5}, 3.1 \times 10^{-5}, 10^{-4}, 3.1 \times 10^{-4}, \dots, 0.01, 0.31, 1\}$.

For a given c , we adopt an iterative process almost the same as the one for DA-T and we do not repeat here except that the expected utility, equation (A.3), is replaced by:

$$\tilde{W}_{i,n}^{(m,t)} = W_{i,n}^{(m,t)} - \frac{1}{M} \sum_{m=1}^M c \times (|L_{i,n}^{(m)}| - 1),$$

which takes into account the application cost associated with $L_{i,n}^{(m)}$.

D.2 Simulating Preference Shocks: Conditional Drawings

Having solved equilibrium strategy numerically, we conduct counterfactual welfare analysis. Our simulation experiment is to answer the following question: *Had a counterfactual market design been implemented at TSE in 2013, what is our best prediction of the counterfactual outcome, based on all the information that is observed by the researcher?*

Therefore, we take the observed applicant and program ordinal preferences as given, although their cardinal preferences can vary while being compatible with the ordinal ones.

To simplify notations, we define $v_j \equiv x\beta_j + \varepsilon_j \equiv b_j + \varepsilon_j$ for $j = 1, 2, \dots, J$, omitting the index for applicants. Recall that $v_0 = \varepsilon_0$. For a given applicant, write $\varepsilon = (\varepsilon_0, \varepsilon_1, \dots, \varepsilon_J)$; L^f is an **extended** rank-order list (i.e., an applicant's submitted ROL under DA combined with program acceptability). Define $\varepsilon_{-j} = (\varepsilon_0, \dots, \varepsilon_{j-1}, \varepsilon_{j+1}, \dots, \varepsilon_J)$ and also the following set:

$$S(L^f) = \{\varepsilon \in \mathbb{R}^{J+1} : \varepsilon \text{ is compatible with } L^f \text{ given } (b_j)_{j \in \mathcal{J}}\}.$$

This means that we draw the vector ε into the density function given by:

$$\frac{\mathbf{1}\{\varepsilon \in S(L^f)\} \prod_{j=0}^J f(\varepsilon_j)}{\Pr(L^f)}.$$

Note that the conditional distribution of ε_j given ε_{-j} is given by:

$$\frac{f(\varepsilon_j)\mathbb{1}\{(\varepsilon_j, \varepsilon_{-j}) \in S(L^f)\}}{\Pr((\varepsilon_j, \varepsilon_{-j}) \in S(L^f) \mid \varepsilon_{-j})}.$$

For $j \neq 0$, the set $(\varepsilon_j, \varepsilon_{-j}) \in S(L^f)$ given ε_{-j} is an interval which can be open or not on the left or on the right according to the following cases:

- (i) Program j is ranked first and therefore if the second ranked is l , the only constraint that $S(L^f)$ imposes on ε_j is :

$$b_j + \varepsilon_j > b_l + \varepsilon_l.$$

- (ii) Program j is ranked last and therefore if the before-the-last ranked is u , the only constraint that $S(L^f)$ imposes on ε_j is :

$$b_j + \varepsilon_j < b_u + \varepsilon_u.$$

- (iii) If program j has two neighbors, from below, l and from above u , then the only constraint is:

$$b_l + \varepsilon_l < b_j + \varepsilon_j < b_u + \varepsilon_u.$$

For $j = 0$, the set $(\varepsilon_0, \varepsilon_{-0}) \in S(L^f)$ is slightly more complicated but follows the same principles. Suppose the ranking of programs is given by $(1, \dots, \bar{J}, \bar{O}, \bar{J} + 1, \dots, \underline{Q}, \underline{J}, \dots, J)$. It translates into the set of inequalities on ε_0 :

$$\begin{aligned} b_{\bar{J}} + \varepsilon_{\bar{J}} > \bar{\eta} + \varepsilon_0 > b_{\bar{J}+1} + \varepsilon_{\bar{J}+1}, \\ b_{\underline{J}-1} + \varepsilon_{\underline{J}-1} > \varepsilon_0 > b_{\underline{J}} + \varepsilon_{\underline{J}}, \end{aligned}$$

using the conventions that $b_0 = +\infty$ and $b_{J+1} = -\infty$. This means that:

$$\varepsilon_0 \in \left[\max(b_{\bar{J}+1} + \varepsilon_{\bar{J}+1} - \bar{\eta}, b_{\underline{J}} + \varepsilon_{\underline{J}}), \min(b_{\bar{J}} + \varepsilon_{\bar{J}} - \bar{\eta}, b_{\underline{J}-1} + \varepsilon_{\underline{J}-1}) \right],$$

and the interval is not empty since $\varepsilon \in S(L^f)$.

Defining all these domain intervals as $I_j(\varepsilon_{-j}, L^f)$, the conditional distribution of ε_j given ε_{-j} and L^f is given by:

$$\frac{f(\varepsilon_j)\mathbb{1}\{\varepsilon_j \in I_j(\varepsilon_{-j}, L^f)\}}{\Pr(\varepsilon_j \in I_j(\varepsilon_{-j}, L^f) \mid \varepsilon_{-j})}.$$

It is easy to simulate since $F(\varepsilon_j) = \exp(-\exp(-\varepsilon_j))$ and therefore for $\delta \in (0, 1)$:

$$F^{-1}(\delta) = -\log(-\log(\delta)).$$

Assume that $I_j(\varepsilon_{-j}, L^f) = [a_j(\varepsilon_{-j}), b_j(\varepsilon_{-j})]$ and draw $u \rightsquigarrow \mathcal{U}_{[0,1]}$ then:

$$\varepsilon_j^* = -\log(-\log(F(a_j(\varepsilon_{-j})) + u(F(b_j(\varepsilon_{-j})) - F(a_j(\varepsilon_{-j})))))) \quad (\text{A.4})$$

is distributed as ε_j given ε_{-j} and L^f .

For Gibbs sampling, the Markov chain that we use is the following:

- (i) Draw $\varepsilon_0^{(0)}$ in the type-I extreme value distribution.
- (a) Above \bar{O} , draw $\varepsilon_j^{(0)}$ sequentially in the reverse order of $1, \dots, \bar{J}$ imposing the constraints sequentially:

$$b_{\bar{J}} + \varepsilon_{\bar{J}} > \bar{\eta} + \varepsilon_0, b_{\bar{J}-1} + \varepsilon_{\bar{J}-1} > b_{\bar{J}} + \varepsilon_{\bar{J}}, \text{ etc.}$$

- (b) Between \bar{O} and \underline{O} , draw $\varepsilon_j^{(0)}$ sequentially in the order of $\bar{J} + 1, \dots, \underline{J} - 1$ imposing the constraints sequentially:

$$\bar{\eta} + \varepsilon_0 > b_{\bar{J}+1} + \varepsilon_{\bar{J}+1} > \varepsilon_0, b_{\bar{J}+1} + \varepsilon_{\bar{J}+1} > b_{\bar{J}+2} + \varepsilon_{\bar{J}+2} > \varepsilon_0, \text{ etc.}$$

- (c) Below \underline{O} , draw $\varepsilon_j^{(0)}$ sequentially in the order of \underline{J}, \dots, J imposing the constraints sequentially:

$$\varepsilon_0 > b_{\underline{J}} + \varepsilon_{\underline{J}}, b_{\underline{J}} + \varepsilon_{\underline{J}} > b_{\underline{J}+1} + \varepsilon_{\underline{J}+1}, \text{ etc.}$$

In this way, $\varepsilon^{(0)} \in S(L^f)$.

- (ii) For $t = \{1, 2, \dots\}$, repeat the steps in which $\varepsilon_0^{(t)}$ is drawn according to equation (A.4) given $\varepsilon_{-0}^{(t-1)}$, and then follow the same order for drawing $\varepsilon_j^{(t)}$ as in steps (ia) to (ic) that is the order of drawings given by:

$$\bar{J}, \dots, 1, \bar{J} + 1, \dots, J;$$

in the meantime, impose the exact intervals $I_j(\varepsilon_{-j}, L^f)$. Note that at every step, the ε of interest belongs to $S(L^f)$ since we draw in the correct intervals.

- (iii) If the chain is long enough, then the distribution of ε is the distribution of interest and expectations of all ε_j can be approximated.

Appendix E Screen Shots from the Web Site for the Experiment

(a) Top Part of the Screen

(b) Middle Part of the Screen

(c) Bottom Part of the Screen

Figure E.1: First Screen: Welcome Page of the Experiment

Notes: This is the first screen an applicant sees when participating in the experiment. It shows, in both English and French, (1) the purpose of the experiment (i.e., trying four mechanisms instead of adopting only one mechanism), (2) some information about the process (deadline, change application during the process, etc.), and (3) the link to the next step. This screen is the same for every applicant, except that the link to the next screen is randomized.

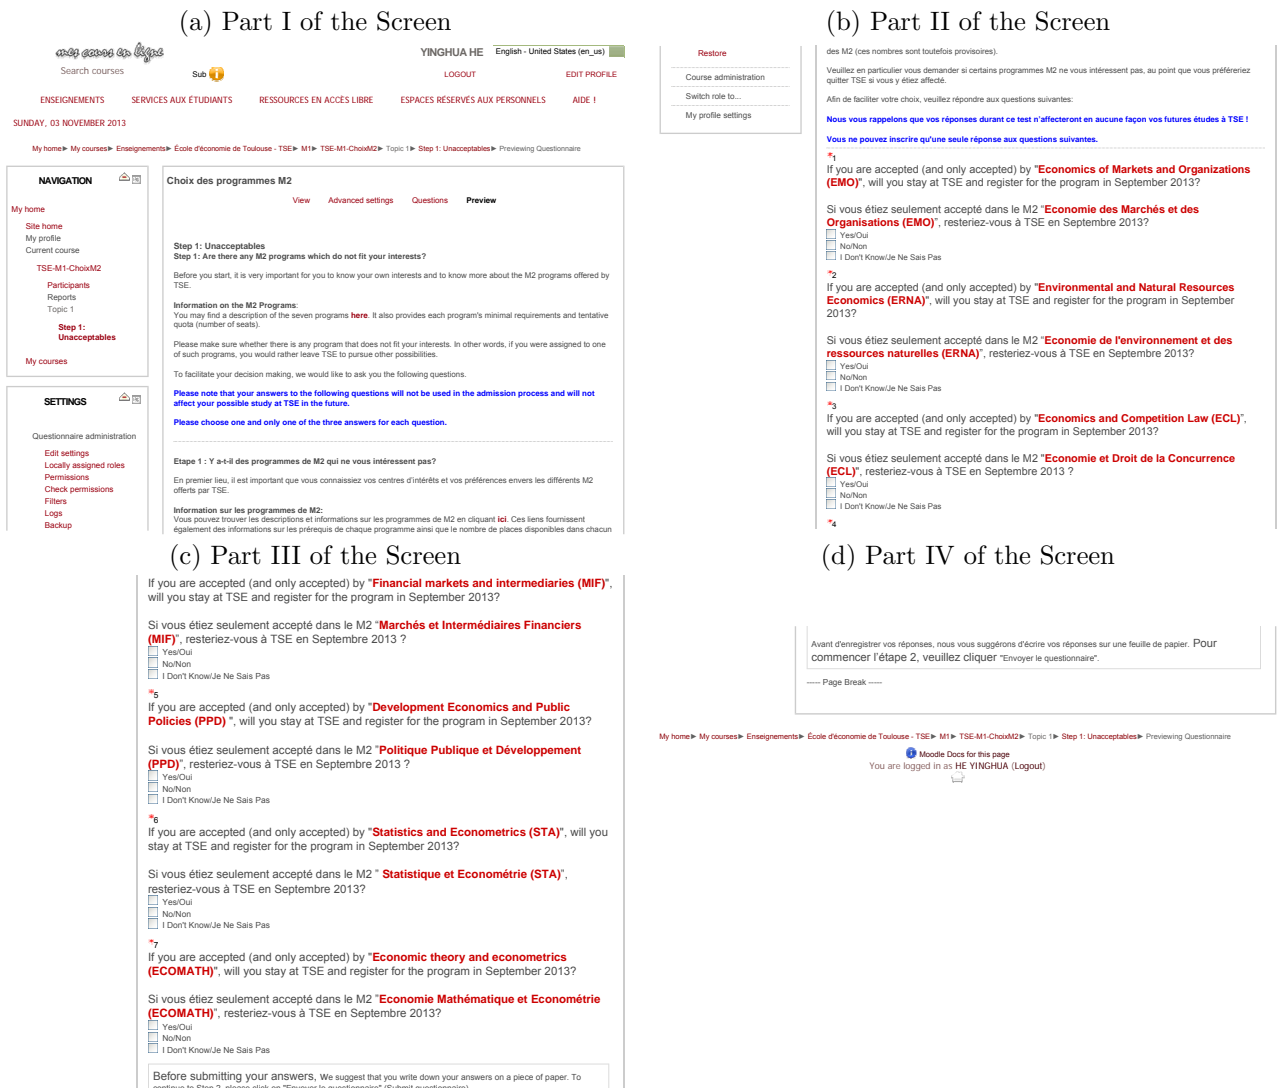


Figure E.2: Second Screen: Survey on Program Acceptability

Notes: This is the second screen an applicant sees when participating in the experiment. It provides, in both English and French, (1) a survey on program acceptability, (2) explanations that the survey answers are not used in the assignment mechanism, and (3) the link to the program descriptions and the link to the next step. The content of the screen is the same for every applicant, but the order of survey questions (i.e., the order of the programs) and the link to the next step are randomized. Therefore, applicants play the four mechanisms in a randomized order.

(a) Top Part of the Screen

(b) Middle Part of the Screen

(c) Bottom Part of the Screen

Figure E.3: The Deferred-Acceptance Mechanism

Notes: This is the screen where an applicant submits her application under the Deferred-Acceptance mechanism. It presents, in both English and French, (1) explanations that the application is used if and only if the Deferred-Acceptance mechanism is chosen, (2) some tips that can be derived from the literature, (3) the links to the pages describing the mechanism and the programs, and (4) the programs that each applicant can rank. This screen is the same for every applicant, except that (1) the link to the next screen is randomized, (2) the order of the programs presented is randomized, and (3) each applicant may reach this screen at different stage in the experiment.

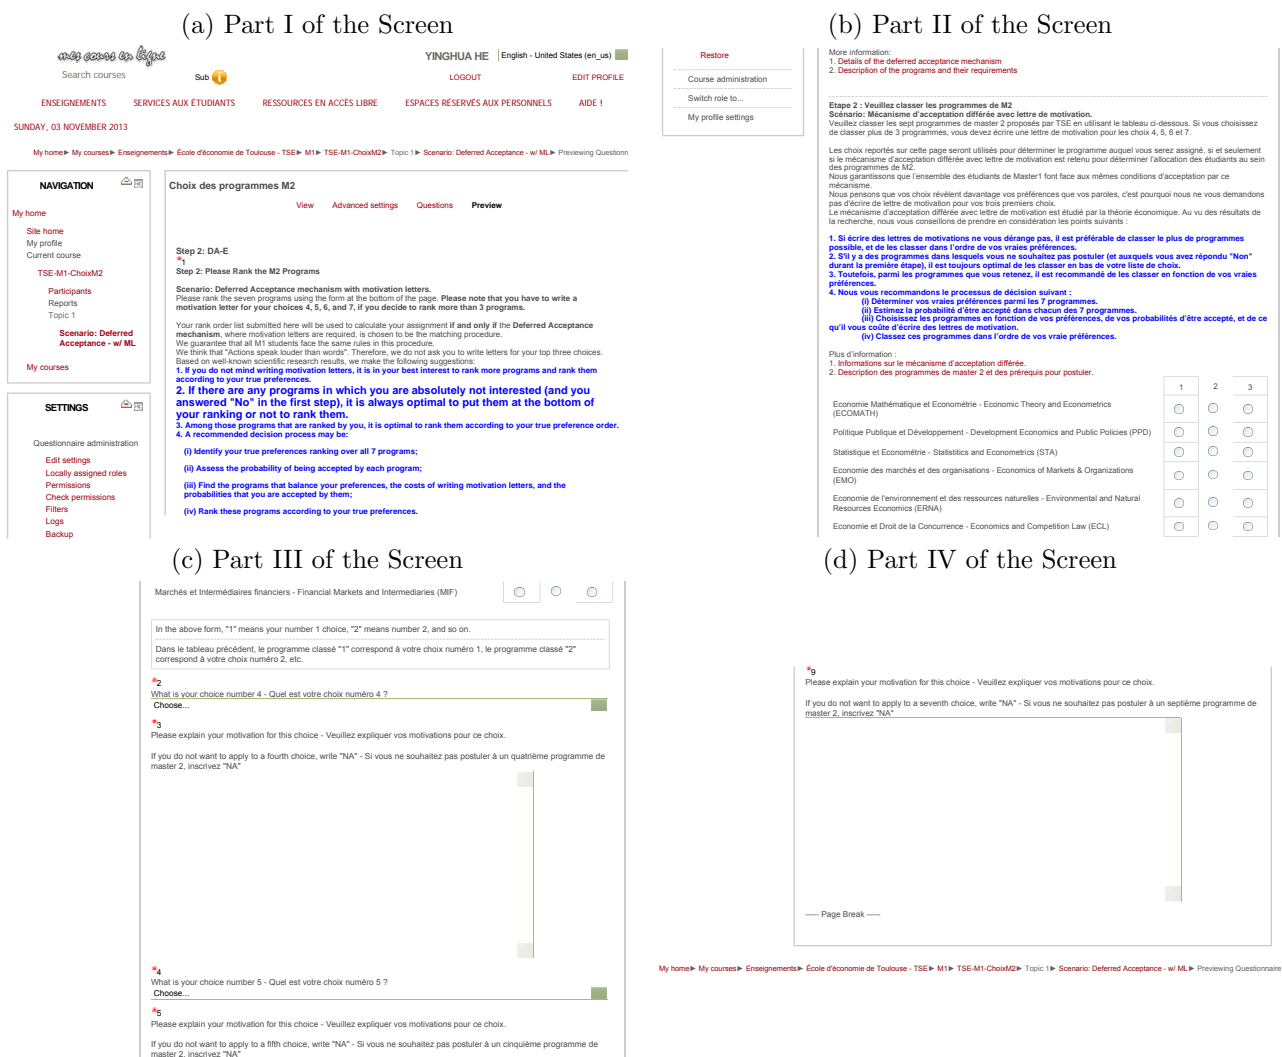


Figure E.4: The Deferred-Acceptance Mechanism with Costs

Notes: This is the screen where an applicant submits her application under the Deferred-Acceptance mechanism with costs (in terms of writing a motivation letter). Part of the screen, between parts III and IV is omitted in the figure, but the information on that part is a continuation of that on part III. It presents, in both English and French, (1) explanations that the application is used if and only if the Deferred-Acceptance mechanism with costs is chosen, (2) some tips that can be derived from the literature, (3) the links to the pages describing the mechanism and the programs, and (4) the programs that each applicant can rank. This screen is the same for every applicant, except that (1) the link to the next screen is randomized, (2) the order of the programs presented is randomized, and (3) each applicant may reach this screen at different stage in the experiment.

(a) Top Part of the Screen

(b) Middle Part of the Screen

(c) Bottom Part of the Screen

	1	2	3	4
Economie Mathématique et Econométrie - Economic Theory and Econometrics (ECONMATH)	<input type="radio"/>	<input type="radio"/>	<input type="radio"/>	<input type="radio"/>
Politique Publique et Développement - Development Economics and Public Policies (PPD)	<input type="radio"/>	<input type="radio"/>	<input type="radio"/>	<input type="radio"/>
Statistique et Econométrie - Statistics and Econometrics (STA)	<input type="radio"/>	<input type="radio"/>	<input type="radio"/>	<input type="radio"/>
Economie des marchés et des organisations - Economics of Markets & Organizations (EMO)	<input type="radio"/>	<input type="radio"/>	<input type="radio"/>	<input type="radio"/>
Economie de l'environnement et des ressources naturelles - Environmental and Natural Resources Economics (ERNA)	<input type="radio"/>	<input type="radio"/>	<input type="radio"/>	<input type="radio"/>
Economie et Droit de la Concurrence - Economics and Competition Law (ECL)	<input type="radio"/>	<input type="radio"/>	<input type="radio"/>	<input type="radio"/>
Marchés et Intermédiaires financiers - Financial Markets and Intermediaries (MF)	<input type="radio"/>	<input type="radio"/>	<input type="radio"/>	<input type="radio"/>

Figure E.5: The Deferred-Acceptance Mechanism with Truncation

Notes: This is the screen where an applicant submits her application under the Deferred-Acceptance mechanism with truncation (i.e., submitting no more than 4 choices). It shows, in both English and French, (1) explanations that the application is used if and only if the Deferred-Acceptance mechanism with truncation is chosen, (2) some tips that can be derived from the literature, (3) the links to the pages describing the mechanism and the programs, and (4) the programs that each applicant can rank. This screen is the same for every applicant, except that (1) the link to the next screen is randomized, (2) the order of the programs presented is randomized, and (3) each applicant may reach this screen at different stage in the experiment.