

Incorporating Reference-Dependence Considerations in Deferred Acceptance*

PRELIMINARY. PLEASE DON'T CITE WITHOUT PERMISSION

Bnaya Dreyfuss

Ofer Glicksohn

Ori Heffetz

Assaf Romm

April 30, 2022

Abstract

Can incorporating expectations-based-loss-aversion (EBLA) considerations help reduce seemingly dominated choices in strategy-proof mechanisms? We run an experiment ($N = 500$) where each participant is randomly assigned into one of four different variations of Deferred Acceptance—{static vs. dynamic} $\times$ {student proposing vs. student receiving}—and play ten different problems of a simulated school-assignment game. While a standard, no-EBLA model predicts the same straightforward behavior across the ten problems and the four variations, an EBLA model predicts stark differences in behavior across both variations. We find that (i) across matching problems, behavior is more consistent with the EBLA model, and (ii) across mechanisms, the *Dynamic student receiving* variant lead to significantly fewer dominated choices.

1 Introduction

A growing body of empirical evidence documents puzzling behavior in strategy-proof matching mechanisms: many participants play seemingly dominated strategies, both in real, large-stakes applications and controlled lab experiments. A prominent example of this behavior is documented in centralized clearinghouses that use rank-order lists (ROLs) submitted by schools and students to match them together based on the deferred-acceptance algorithm (DA; [Gale and Shapley 1962](#)). Even though the mechanism is strategy-proof for students, i.e., submitting a truthful ranking of schools is a weakly dominant strategy, a non-trivial share of supposedly informed students who participate in real-world matches appear to make dominated choices. For example, evidence from different countries shows that students do not rank a school program with funding above the *same* program without the funding. Similarly, in simple allocation experiments, a substantial fraction of participants rank smaller amounts of money above larger ones. For a recent review of field evidence see [Dreyfuss et al. \(forthcoming\)](#). For a recent review of lab evidence see [Hakimov and Kübler \(2019\)](#).

Recently, [Dreyfuss et al. \(forthcoming\)](#) and [Meisner and von Wangenheim \(2019\)](#) show that expectations-based loss aversion (EBLA) can be a potential explanation for such non-straightforward behavior. Intuitively, participants with EBLA preferences may intentionally downrank (or, in some cases, completely omit) a high-value, low-probability school to avert the likely disappointment

*Dreyfuss: Department of Economics, Harvard University (e-mail: bdreyfuss@g.harvard.edu); Glicksohn: Bogen Family Department of Economics and Federmann Center for the Study of Rationality, The Hebrew University of Jerusalem (e-mail: ofer.glicksohn@mail.huji.ac.il); Heffetz: S.C. Johnson Graduate School of Management, Cornell University, Bogen Family Department of Economics and Federmann Center for the Study of Rationality, The Hebrew University of Jerusalem, and NBER (email: oh33@cornell.edu); Romm: Bogen Family Department of Economics and Federmann Center for the Study of Rationality, The Hebrew University of Jerusalem (e-mail: assaf.romm@mail.huji.ac.il). We thank Shengwu Li, Matthew Rabin, JESC-lab group and seminar participants at Cornell University, Hebrew University, and Harvard University for invaluable comments; Yonatan Rahimi, Tal Asif, Ronny Gelman, and Ilana Lavie for excellent research assistance, and especially Guy Lichtinger for thoughtful and creative research assistance in the development of the project. This research was supported by the Israel Science Foundation (grant No. 2968/21), the Maurice Falk Institute, and the Johnson School at Cornell.

	Static: Submit a list in advance	Dynamic: Decide at each step
Proposing: Students apply to schools, schools retain the highest-ranked applications.	SP	DP
Receiving: Schools send admission offers to students by ranking, students can retain at most one offer in each step.	SR	DR

Table 1: The four different DA variants

of rejection. Both papers show how an EBLA model, as formulated by Kőszegi and Rabin (2006, 2007, 2009), is consistent with various empirical patterns from the lab and the field. Dreyfuss et al. (forthcoming) also structurally analyzes existing lab data from Li (2017) and shows that the EBLA model indeed fits the data significantly better than the classical-preferences benchmark.

In this paper we ask: Is there a DA implementation that induces straightforward behavior even when students are loss averse? This question is both theoretically important and has practical, real-world implications.

To answer this question, we study, both theoretically and empirically, the potential role of EBLA in four different variants of DA. Summarized in table 1, we analyze the predicted effect of varying two features of DA: (i) *static* (normal-form) versus *dynamic* (extensive-form) implementation, and (ii) *student-proposing* versus *student-receiving* role designation. In the industry-standard DA variant, Static Student-Proposing (SP), students propose to schools in an order determined by rank-order lists (ROLs) they submit in advance. The dynamic version of this variant, Dynamic Student-Proposing (DP), requires students to actively propose to schools, in real-time, in any order they choose. In the Static Student-Receiving (SR) variant, students respond to school admission offers according to ROLs they submit in advance. Lastly, in the Dynamic Student-Receiving (DR) variant, students are required to actively respond to admission offers from schools, in real-time, by retaining at most one offer at any given moment.

We make three contributions: theoretical, empirical, and practical. First, we theoretically show that in large markets (where, following Azevedo and Leshno (2016), there is a continuum of students and a fixed number of schools), the Dynamic Student-Receiving (DR) variant is “EBLA-proof” in that it eliminates non-straightforward behavior for any level of loss aversion. (As we discuss below, in concurrent work, Meisner and von Wangenheim (2019) prove a related result.) In contrast, EBLA-driven non-straightforward behavior is predicted to be quite prevalent under the other three variants, especially in highly competitive settings. Second, we experimentally test the model’s predictions, both across the four algorithm variants (by randomly assigning different participants to different variants) and across matching settings (by exogenously varying the degree of competitiveness across ten matching problems that each participant faces). Finally, for practitioners who may care less about theoretical results and their empirical tests and more about pragmatics, we show that of the four mechanism variants, DR in fact stands out in minimizing non-straightforward behavior in our experiment. As we discuss below, we view this result as a valuable empirical fact regardless of its theoretical drivers (and regardless of whether DR is feasible to implement in practice).

We report our theoretical analysis in section 2. We focus on large-market economies, with a continuum of students and a finite set of capacity-constrained schools, and rely on the theoretical characterization of Azevedo and Leshno (2016); Abdulkadiroğlu et al. (2015). We show that in such economies, DR is EBLA-proof: for any degree of loss aversion and any belief distribution, the EBLA model predicts that participants will always play what we call *straightforward strategies* (“truthful strategies”). Straightforward strategies are strategies consistent with a ranking of alternatives according to their consumption values. In DR, straightforwardness implies that in each step, the student will always pick the highest-value offer from the set of offers they receive.

In contrast with DR, the other three DA variants we examine are not EBRD-proof: in highly competitive settings, the model predicts loss-averse individuals under SP, SR, and DP to behave in

a non-straightforward way.

The theory in section 2 thus makes two types of testable predictions. First, for each of the three non-EBRD-proof variants, straightforwardness should decrease with the competitiveness of the matching setting. Second, in sufficiently noncompetitive environments, straightforwardness should be similarly high across all four variants. These two predictions, together with some given distribution of the degree of loss aversion in a population, yield specific quantitative predictions both across matching settings and across DA variants.

We outline our experimental design in section 3. Subjects are randomly assigned to one of four treatments corresponding to the four DA variants, and play the same ten randomly ordered matching problems, each simulating a different large-market matching economy. Using the population distribution of loss aversion from a previous study, the ten matching problems are designed to create variation in the predicted prevalence of straightforward behavior in the three non-DR treatments. Three of the ten problems ("weak student" problems) are designed to predict a moderately high fraction of non-straightforward behavior (24–31%); another four ("medium student") problems predict a lower fraction of such behavior (12–24%), and the remaining three ("strong student") problems predict no non-straightforward behavior at all.

In section 4 we report experimental results from two independent samples: Cornell students ($N = 196$) and Israeli psychology graduate-school candidates ($N = 304$) who participated in the (DA-based) Israeli Psychology Master's Match (IPMM), which we view as a highly relevant population. Our main predictions and results are summarized in figure 8 in section 4.

The figure shows that, as the theory predicts, in the non-DR treatments the share of observed non-straightforward behavior monotonically decreases when moving from weak- (33–48%) to medium- (30–37%) to strong-student problems (19–22%). In the DR treatment (where the model predicts no NSF) we find lower levels and a flat trend (16–19%). Taken together, we find that consistent with our predictions, there is a large reductions in observed NSF behavior in DR compared to the other three treatments in competitive environments, but not in noncompetitive environments. Moving beyond aggregate NSF share predictions, we show that (i) the specific NSF strategies that the model predicts are indeed those who are common in the data and vice-versa, and (ii), in problems in which the model predicts specific NSF strategies (e.g., ranking the third-highest value school on the top of the list), these specific NSF ROLs are indeed the most common. In section 5 we conclude and discuss the potential implications of our findings.

2 Theoretical Analysis

2.1 General Framework

Consider a two-sided market with n capacity-constrained schools and a unit mass continuum of students, governed by a DA mechanism that matches (masses of) students to different schools. The set of schools is denoted by $S = \{s_1, \dots, s_{n+1}\}$, where s_{n+1} represents the outside option. Each student i has strict (reference-independent) preference over schools represented by a utility function $m_i(\cdot)$, and summarized by the vector $\mathbf{m}_i$. We call this component *consumption utility*. Later in this section we add a *news-utility* (Kőszegi and Rabin, 2009) component. Schools rank students by assigning each of them a *priority score* $\rho_{ij} \in [0, 1]$ so that ρ_{ij} represents student i 's relative rank at school j . Each school j has some capacity $c_j \in (0, 1]$ that represents the share of students it can admit.

As shown in Azevedo and Leshno (2016), given some students' and schools' submitted rankings, the match result from running DA in this economy is characterized by a vector of thresholds $T = \{T_1, \dots, T_{n+1}\}$, one for each school, where each student i is assigned to her highest-ranked school such that $\rho_{ij} \geq T_j$.¹ Moreover, student-proposing and student-receiving DA result in the same match

¹A problem that might arise in our setting is that while in the continuum economy the DA algorithm converges in the limit and the allocation in the limit is well defined, the process might not be complete in finite time. We ignore this issue and assume that matching is always achieved in finite time.

(and associated thresholds vector T). We assume throughout $c_{n+1} = 1$, which implies $T_{n+1} = 0$, so that the outside option is always available with certainty.

Students know the thresholds $T = \{T_1, \dots, T_n\}$ and (correctly) treat them as exogenous. However, we assume that students do not know their priority score at each school, ρ_{ij} , but rather have a probability distribution over their score which we denote by the CDF G_{ij} . This captures the idea that students know how selective a school is but not necessarily their exact ranking relative to other students. We denote student i 's joint distribution of priority scores in all schools by G_i . In what follows we focus on the decision from a single-student's perspective fixing schools' and other students' behavior and therefore we suppress the index i . Moreover, WLOG we index schools by student i 's preferences so that $m_i(s_j) > m_i(s_k)$ for all $j < k$ and $m_i(s_{n+1})$ is normalized to 0.

As explained in the introduction we focus on four different *variants* of the DA algorithm resulting from the product $\{\text{proposing, receiving}\} \times \{\text{static, dynamic}\}$. We denote the DA variant governing the market by M . A *matching problem* from a student perspective can therefore be summarized by $\langle M, G, \mathbf{m} \rangle$, where M defines the matching process, G is the joint distribution over priority scores at said schools, and $\mathbf{m}$ is a vector of school consumption utilities. Fixing all other players' actions, we can treat each DA variant M as a game student i plays against Nature, where Nature's moves are (probabilistically) determined by G .

We formally define the game induced by each of the four variants in Appendix A, and provide here the necessary notation. A history in the set of history is denoted by $h \in \mathcal{H}$ where $z \in \mathcal{Z}$ is a terminal history in the set of terminal histories. Subhistories are denoted with subscripts (h_t include the first t periods in h). A strategy is denoted by $l \in \mathcal{L}$. A strategy l and a history h induce (i) belief distribution over terminal histories, denoted by $\tilde{F}_{l|h}$, and (ii) a CDF over final matches, denoted by $F_{l|h}$.

We now introduce our definition of non-straightforward behavior ("misrepresentation") applicable to all four DA variants.

Definition 1. In a DA variant, a strategy l that conforms to a rank-order list (ROL) ordering schools by their consumption value is a **Straightforward (SF)** strategy. A **Non-straightforward (NSF)** strategy is any strategy inconsistent with this ROL.

The SF strategy is simply the ROL that ranks schools by their values for both static variants. In DP, the SF strategy sequentially applies to schools by order of their values. In DR, it retains the highest-valued offer at any given point for any given history. Under classical preferences, a straightforward strategy can be seen as *truthful*, as it reveals the actual ranking of schools by their consumption value.

2.2 EBLA Preferences

As mentioned earlier, in addition to classical consumption utility, students' utility functions have a news-utility component as formalized in Kőszegi and Rabin (2009), and reviewed here. The formulation may look cumbersome at times, but it is aimed to intuitively capture a basic feature in people's preferences: the utility and disutility of belief updating. In short, learning about a higher likelihood of a good outcome (relative to previously held beliefs) entails positive utility, and learning about a lower possibility of a good outcome entails disutility. In this model, "loss aversion" means putting a higher weight on the latter.

News utility is defined over belief updates. Formally, during the matching process, when moving from the history h_t to the history h_{t+1} , the student rationally updates her beliefs over final outcomes from $F_l(\cdot|h_t)$ to $F_l(\cdot|h_{t+1})$ where $F_l(\cdot|h)$ is the CDF over final payoffs (i.e., realizations of m_i) conditional on reaching history h while following strategy l . Let F and F' be two belief distributions over the outcomes in the game. The news utility function is given by:

$$N(F'|F) = \int_0^1 \mu(m(c_{F'}(p)) - m(c_F(p))) dp, \quad (1)$$

where $c_F(p)$ denotes the consumption level (school) at percentile p of the cumulative distribution function resulting from F , and $\mu(\cdot)$ is a piece-wise gain-loss utility function defined in the following way:

$$\mu(x) = \begin{cases} x & \text{if } x \geq 0 \\ \lambda x & \text{if } x < 0 \end{cases} \quad (2)$$

Where $\lambda \geq 1$ represents a loss-aversion parameter that differs across students. The function N describes how old beliefs are compared to new beliefs, percentile by percentile: in each period t , the student compares, for each percentile, the outcome under F to the outcome under F' , with negative surprises given a higher weight, parameterized by $\lambda \geq 1$.

The total utility from a strategy l is therefore given by:

$$U(l) = \mathbb{E}_{F_l}[m(s)] + N(F_{l|z_1} | F_{l|\emptyset}) + \mathbb{E}_{\tilde{F}_l} \left[\sum_{t=2}^{\bar{t}_z} N(F_{l|z_t} | F_{l|z_{t-1}}) \right] \quad (3)$$

The first term is the expected consumption utility described above. The last term is the expected sum of news utility streams from the terminal history, given a probability distribution over possible terminal histories $\tilde{F}_l$. Note that beliefs over payoff prior to period 1 (i.e., prior to choosing a strategy) are not defined. Therefore the second term includes an undefined object, $F_{l|\emptyset}$. To close the model, we assume that prior to entering the mechanism, the student believes that she will consume the outside option with certainty, i.e., $F_{l|\emptyset} = F_0$ for all l , where $F_0(x) = 1$ for all $x \geq 0$ and 0 otherwise.²

To derive predictions, we use Preferred Personal Equilibrium (PPE; [Kőszegi and Rabin, 2009](#)) as our solution concept for all variants. We provide a formal definition in appendix A.

Under this solution concept, given the expectation induced by the strategy l , the student must find it optimal to follow through with it in all possible contingencies. Among all strategies that fulfill this criteria, the one that provides the most utility is selected.

As we mentioned above, under our large markets assumption, all four variants are strategy-proof for classical-preferences students and equivalently to loss neutral students ($\lambda = 1$). The following definition extends strategy-proofness to economies with EBLA agents.

Definition 2. A DA variant is **EBLA strategy-proof** if for any degree of loss aversion λ , and any belief G , the optimal (PPE) strategy is SF.

Equipped with this new notion of strategy proofness, the following proposition differentiates between the model's prediction in the DR variant and the other three variants.

Proposition 1. Dynamic Student-Receiving (DR) DA variant is EBLA strategy-proof while the other three DA variants (SP, SR, DP) are not EBLA strategy-proof.

All proofs are relegated to the online appendix. The second part of the proposition has been extensively discussed in DHR and [Meisner and von Wangenheim \(2019\)](#), and is easily proved by a counterexample. The first part of the proposition is proved by backward induction and has a very simple intuition. The model predicts that, no matter what she had planned to do, when a student is presented with a choice set of alternatives that she can get with certainty, she will always choose the one that gives her the highest consumption utility. In DR, the student chooses between schools that already sent her an offer, i.e., "sure things," and cannot do better than keeping the best one. Applying this argument iteratively shows that in the DR variant, the straightforward strategy is indeed the only one that is optimal in any decision node, in any possible history. This highlights

²This assumption is calibrationally, but not qualitatively, substantive. In particular, it does not affect our between-variant results, nor does it change our qualitative predictions within each mechanism. However, it is calibrationally and quantitatively substantive in that assuming more optimistic initial beliefs will require a higher value of λ for the optimal strategy to be NSF.

the importance of dynamic implementation: In the (static) SR variant, the student can lower her expectations and reduce potential disappointment by committing (via ROL submission) to reject some offers. In contrast, in DR, the student anticipates that she will keep desirable offers and is, therefore, is forced to choose an SF strategy.³

3 Experimental Design

The experiment consists of simulating a large matching market. Each subject is randomly assigned to one of 2×2 (static/dynamic $\times$ proposing/receiving) treatments and participates, following a tutorial, in ten order-randomized incentivized *matching problems*, each simulating a different large matching market. We denote the treatments by {SP, SR, DP, DR}. Each subject who finishes the experiment receives a participation payment of \$4 as well as the sum of values of the schools they matched with across the ten matching problems.

Each matching problem consists of five schools to which the subject can be matched with at the end of the matching process. Each school is given a dollar value—the payment the subject will gain if matched with said school.

The experimental design closely follows the theory section: Each subject is assigned a *priority score* at every school, which is a random, uniformly, and independently distributed integer number between 0 and 99. This score represents the school’s preference: the subject’s rank compared to other (hypothetical/simulated) candidates. For example, if the priority score of the subject is 78, it means that she is ranked higher than exactly 78 percent of candidates. Each school has a *threshold*, the minimal accepted priority score. That is, only candidates whose priority score is above a school’s threshold can be accepted to that school.

At the beginning of every problem, the subject learns about each school’s threshold but not about her own priority scores at the different schools. This creates an *unconditional* probability of acceptance at each school. Figure 1 shows the example matching problem given in the tutorial.

School	Threshold	Value	Your Priority Score	Chance that Your Priority Score $\geq$ Threshold
Pine Peak	60	\$0.75	?	40%
Birch Hill	0	\$0.25	?	100%
Hickory Bridge	50	\$0.50	?	50%
Maplecrest	80	\$1.25	?	20%
Elm South	70	\$1.00	?	30%

Figure 1: Screenshot of an the matching problem shown in the tutorial.

To make the process transparent and easier to monitor, at the end of every matching process subjects receive full information: they learn about their match as well as their priority score at each school (the “?”s in figure 1 are replaced with the realized priority scores).

The matching process varies by treatment. In both *proposing* treatments, subjects propose to schools, while in both *receiving* treatments subjects receive offers from schools. In both *static* treatments, subjects submit in advance a rank-order list (ROL) of schools, and the DA algorithm is run based on their submitted ROL. In both *dynamic* treatments, subjects do not submit ROLs but instead are required to actively make decisions at each step. Specifically, in the DP (dynamic proposing)

³This also illustrates a more subtle point about welfare: while we predict that DR maximizes the students’ ex-ante *consumption* utility, it actually *decreases* their overall (consumption and news) utility. We still find DR appealing since we view EBLA-driven behavior in these settings as a mistake. For further discussion see DHR.

treatment, subjects decide at each step which school to propose to. In the DR treatment, subjects decide at each step which offer to keep. Appendix F provides screenshots for all treatments.

As explained in the theory section, subject i with loss-aversion parameter λ_i is faced with a matching problem $\langle M, G, \mathbf{m} \rangle$, where M defines the matching process, G is the joint distribution over priority scores (recall that in the experiment priority scores are uniform and independent), and $\mathbf{m}$ is a vector of school consumption utilities. We assume linear consumption utility, $m_i(s_j) = s_j$ for every school j and subject i .

We created ten matching problems given to each subject, held fixed across treatments. Each problem is defined as five school values and (unconditional) probabilities of acceptance. For the purpose of creating predictions, the most important feature in a matching problem is the probability of acceptance at the highest-value school. Specifically, this probability determines the predicted behavior of a subject with a moderately high loss-aversion parameter (say, $3 < \lambda < 6$). We group the ten matching problems into three that differ in the predicted share of behavior inconsistent with the straightforward strategy, which we call *NSF*. To create these predictions we use a previously-estimated population distribution of loss aversion. Specifically, in “weak-student” problems (three of the ten problems), the probability of acceptance at the highest-valued school is small (5%), which implies a large predicted share of NSF strategies (24%–31%). In “medium-student” problems (four out of ten), the probability of acceptance at the highest-value school is larger (20%–30%), and therefore the predicted share of NSF strategies is lower (12%–24%). Finally, in “strong-student” problems where the probability of acceptance at the highest-value school is high (60%–65%), which implies no predicted NSF strategies (under the past empirically estimated distribution of λ). These predictions for each of the three types of matching problems are summarized in the top panel of figure 8, which shows the predicted trend in the share of observed NSF strategies under all non-DR treatments, and a flat zero predicted share of NSF play under DR.

We provide further details both on the parameters used and the model’s predictions in each of the ten problems in appendix C.

It is important to note that these are predictions on chosen strategies and not on the observed behavior. Under the static treatments (SP, SR) there is no difference between strategies and behavior as subjects’ submitted ROLs are fully observed by us. However, under the dynamic treatments (DP, DR), we only observe *actions* dynamically taken by subjects, which only reveals partial information on their full strategy.

Specifically, under DP, we only observe the set and order of schools the subject applied to by the time the process ends (when either some school accepts, the subject decides to stop the process, or every school rejects the subject). Similarly, in DR we only observe the set of decisions regarding which of the offers a subject received they decided to keep. Such observed decisions can be either *SF-consistent* or *SF-inconsistent*.

Therefore, under the dynamic treatments, a SF-consistent behavior does not necessarily imply a chosen SF strategy, but a SF-inconsistent behavior necessarily implies a chosen NSF strategy.⁴ This means that under the dynamic treatments, *observed* NSF behavior is only a lower bound for *chosen* NSF strategies.

This, in turn, requires caution in comparing our primary measure of suboptimal behavior—the number of problems with SF-inconsistent behavior—across treatments. On the one hand, observing less NSF behavior under a dynamic than under a static treatment could merely result from our inability to observe full strategies under the former. On the other hand, if one of policymakers’ goals is to reduce observed NSF behavior, including potentially by reducing NSF opportunities, then such a cross-treatment comparison is informative. In appendices C and D we explore the use of and results from alternative measures of NSF behavior.

⁴The only exception is the SF-consistent case where, under DP, a subject is observed to send applications to all schools in order, consistent with only the ROL 12345.

4 Results

We ran the experiment on two independent samples.⁵ The first sample is drawn from the Cornell BSL (Business Simulation Lab) SONA system, an online pool of Cornell students. The second sample is recruited from the 2021 and 2020 participants Israeli Psychology Matching Mechanism (hereafter IPMM), a DA-based clearinghouse that matches students and graduate programs in psychology.

In the main text, our analysis pools together both samples.⁶ Figure 8 compares our main findings to our main predictions. Panel (a) presents predictions based on EBLA preferences compared to Classical preferences, by treatment and problem type. Our main within-treatments prediction is the downward trend described above. The theory predicts this trend in the three non-DR treatments. Across treatments, we predict that chosen NSF strategies will be less prevalent in the DR treatment, compared to the rest of the treatments, in all but strong matching-problem types. Panel (b) presents the empirical fraction of NSF behavior and the statistical comparisons.




Figure 2: Summary of predictions and findings. Panel (a) presents NSF predictions by treatment and problem type under classical preferences (downwards triangle) and EBLA preferences (upwards triangle). The predictions are based on population distribution estimate from Dreyfuss et al. (forthcoming). Panel (b) presents the empirical shares. Error bars represent standard errors calculated from regressing NSF behavior on treatment, clustering at the individual level.

First, the data seem to support our within-treatments prediction: there is a clear downward trend in the fraction of observed NSF behavior in all non-DR treatments when moving from weak- to medium- to strong student problem types, from 33–48%, to 30–37% to 19–22%. This trend matches our EBLA-model predicted trend quite well, relative to the flat, zero-NSF predictions of classical preferences.⁷ In contrast, the trend in DR is much flatter, consistent with both classical and EBLA models. However, neither model can explain the prevalence of observed NSF behavior in DR.

⁵Preregistered separately at <https://aspredicted.org/rd2y5.pdf> and <https://aspredicted.org/4qc8q.pdf>, respectively.

⁶Appendix D replicates the main analysis for each sample separately. It shows that while the general level of observed NSF behavior is markedly lower in the Israeli sample compared to the Cornell sample, the trend remains the same across samples.

⁷Notice that we are not allowing ourselves any degrees of freedom in our predictions: they are generated independently from the data—in fact, they were generated before we even collected the data—and therefore we focus on trend/direction predictions rather than level prediction.

Moving to between-treatments predictions, these also seem to be supported by the data: the fraction of observed NSF behavior under DR is indeed substantially lower than in the rest of the treatments in problems where the model predicts a difference (weak- and medium-student problems), but *not* in problems where the model does not predict a difference (strong-student problems).

To summarize, the trends in the data, both between problem types and between treatments seem to match those predicted by our model, whereas classical preferences cannot easily explain either. However, we cannot explain a baseline level of observed NSF behavior that is prevalent (around 20%) even in problem types and treatments where our model coincide with classical preferences in predicting zero NSF play.

As mentioned in the previous section, our NSF measure is not easily comparable across treatments.⁸ Our analysis in appendix E suggests that this bias is not substantial and does not drive our findings. Specifically, we show that observing lower share of NSF in DR is unlikely driven by having fewer opportunities to exhibit NSF behavior.

Moving beyond aggregate NSF share predictions, our EBLA model makes predictions on the specific NSF ROLs subjects will submit in both static treatments.⁹ As an additional test for our predictions, we compare the distribution of predicted ROLs to their frequency in the data pooling together both static treatments. Figure 3 presents the aggregate distribution across all ten matching problems, in logarithmic scale.




Figure 3: Theoretical EBLA predictions and empirical shares of ROLs in static treatments, plotted on logarithmic scale. ROLs shown are ones with at least 1% predicted share or empirical share. These ROLs constitute 98% of theoretically predicted ROLs and 86% of empirically observed ROLs.

Figure 3 illustrates how well the empirical frequency of ROLs fits the EBLA model predictions. Specifically, the two highest predicted NSF ROLs (21345 and 23145) are the most empirically prevalent NSF ROLs. Furthermore, all six most empirically prevalent NSF ROLs are consistent with the structure of a potential optimal play, as characterized in [Meisner and von Wangenheim \(2019\)](#).¹⁰

Another way to examine specific predictions compared to observed behavior patterns is to look at differences between problems. We focus on weak problems, and characterize the most prevalent

⁸As explained in C, our experiment is under-powered to detect differences in alternative, more comparable measures, which is indeed what we find in the data.

⁹This also extends to DP treatment, but in this treatment we do not observe the subject's full strategy but rather only an incomplete sequence of applications.

¹⁰[Meisner and von Wangenheim \(2019\)](#) show that an optimal role has a transitive characteristic in that, for example, if optimal play ranks 3 above 2, then it must also rank 2 above 1.

ROL both in the data and in our theoretical predictions. As we can see in panels (a) and (c) in figure 4, in problems 1 and 2 the theory predict that the most prevalent ROL will rank the second-highest value school on top. In contrast, problem 3 was chosen so that the the most prevalent NSF ROL will rank the third-highest value school on top.¹¹

Panels (b) and (d) in figure 4 show the empirical share of NSF ROLs that with the second- and third-highest value school, respectively, for non-DR treatments. As we can see, the empirical data closely resembles the pattern predicted by our model predictions when examining within the *weak* student problems.

NSF by ROL Type




Figure 4: Fraction of NSF ROLs where the 2nd (3rd) highest-valued school is the one ranked first (*weak* problems only). Panels (a) and (c) present chosen NSF strategy predictions by problem type under Classical (downward triangles) and EBLA (upward triangles) preferences. The predictions are based on population distribution estimate from Dreyfuss et al. (forthcoming). Panels (b) and (d) present the empirical shares. Error bars represent standard errors calculated from regressing non-straightforward behavior on *matching problem*, clustering at the individual level.

¹¹The exact parameters of all problems are in a table in appendix C)

5 Discussion

The results from a controlled experiment appear to lead to several conclusions regarding the incorporation of EBLA preferences in deferred acceptance matching mechanisms. First, we replicate previous findings that a substantial fraction of participants in simple allocation games choose seemingly dominated actions, a behavior that Classical preferences cannot explain. The experiment was designed to test our model’s prediction and indeed the data seem to strongly support them. We conclude that based on our results, the inclusion of EBLA preferences can explain a substantial fraction of what appears to be a puzzling behavior from a classical-preferences perspective.

Specifically, we find strong support for our main within-treatment prediction of a downward trend of chosen NSF as a function of the degree of competitiveness under *non-DR* treatments. We also find that the distribution of ROLs chosen by subjects in the static treatments aligns with the EBLA model predictions, both in the aggregate and by matching problem. This implies that our conclusion does not trivially result from the EBLA model allowing for more optimal strategies as the EBLA model predicts specific types of strategies for every matching problem in each of the three the non-DR treatment.

Our main between-treatments prediction of fewer chosen NSF strategies under DR treatment is also supported the data as we do find a substantially lower level of observed NSF behavior in this treatment, in exactly the problems that we expected to find such difference. Even though this measure is potentially biased for the chosen NSF strategy, we we show evidence that suggests that this does not confound our results.

Taken altogether, it seems that while the DR treatment is EBLA-proof, implementing it does not completely eliminates NSF play, and there is still a non-negligible fraction of participants that choose dominated actions. The theoretical drivers behind this behavior remains an open question. Our results suggests that the question of which variant is better depends on policy makers’ objective function and on the type of non-straightforward behavior they are willing to tolerate.

References

- Atila Abdulkadiroğlu, Yeon-Koo Che, and Yosuke Yasuda. Expanding” choice” in school choice. *American Economic Journal: Microeconomics*, 7(1):1–42, 2015. <https://doi.org/10.1257/mic.20120027>.
- Eduardo M Azevedo and Jacob D Leshno. A supply and demand framework for two-sided matching markets. *Journal of Political Economy*, 124(5):1235–1268, 2016. <https://doi.org/10.1086/687476>.
- Bnaya Dreyfuss, Ori Heffetz, and Matthew Rabin. Expectations-based loss aversion may help explain seemingly dominated choices in strategy-proof mechanisms. Working Paper 26394, National Bureau of Economic Research, forthcoming. <https://doi.org/10.3386/w26394>.
- David Gale and Lloyd S Shapley. College admissions and the stability of marriage. *The American Mathematical Monthly*, 69(1):9–15, 1962. <https://doi.org/10.1080/00029890.1962.11989827>.
- Rustamdjan Hakimov and Dorothea Kübler. Experiments on matching markets: A survey. Technical report, WZB Discussion Paper, 2019. <http://hdl.handle.net/10419/196908>.
- Botond Köszegi and Matthew Rabin. A Model of Reference-Dependent Preferences. *The Quarterly Journal of Economics*, 121(4):1133–1165, 2006. <https://doi.org/10.1093/qje/121.4.1133>.
- Botond Köszegi and Matthew Rabin. Reference-dependent risk attitudes. *American Economic Review*, 97(4):1047–1073, 2007. <https://doi.org/10.1257/aer.97.4.1047>.
- Botond Köszegi and Matthew Rabin. Reference-dependent consumption plans. *American Economic Review*, 99(3):909–36, 2009. <https://doi.org/10.1257/aer.99.3.909>.

Shengwu Li. Obviously strategy-proof mechanisms. *American Economic Review*, 107(11):3257–87, 2017. <https://doi.org/10.1257/aer.20160425>.

Vincent Meisner and Jonas von Wangenheim. School choice and loss aversion. *Available at SSRN* 3496769, 2019. <http://dx.doi.org/10.2139/ssrn.3496769>.

A Definitions

A.1 DA Variants as Games

A history $h \in \mathcal{H}$ is a sequence of actions by the student and Nature prescribed to every period t it includes. We define $G(\cdot|h)$ as the student's joint belief over priority scores given history h , with G being her initial belief at $h = \emptyset$. A terminal history (in the set of terminal histories) $z \in \mathcal{Z}$ is a history that includes a *terminal period* i.e., the period in which no more actions are taken, and the student is informed about her final match from S . The student's final match under the terminal history z is captured by the outcome function $O(z): \mathcal{Z} \rightarrow S$, and the terminal period in the terminal history z is denoted by $\bar{t}_z$. Subhistories are denoted using subscripts, i.e., z_t $t \leq \bar{t}_z$ is the realization of the terminal history z up to period t .

A strategy $l \in \mathcal{L}$ for the student assigns an action to every period in which the student is called to take an action, for every possible history. We denote the set of continuation strategies consistent with history h by $\mathcal{L}_h$. A strategy l , a history h and a belief $G(\cdot|h)$ jointly induce a belief over terminal histories, which we denote by $\tilde{F}_{l|h}$, and a CDF over final payoffs (i.e., realizations of m_i) which we denote by $F_{l|h}$.

There are only two periods in both *static* variants: at $t = 1$, the student chooses a ROL, and $t = 2$ is the terminal period in which she learns about her final match. In both of these variants, the set of strategies $\mathcal{L}$ is the set of possible ROLs, i.e., all permutations of all sets in the power set of S . Under our large-markets assumption, both variants are outcome equivalent: for every ROL the student submits at $t = 1$, at $t = 2$ she is matched with the highest-ranked school s_j such that $\rho_j \geq T_j$.

In the dynamic student-proposing (DP) variant, the student sends an application in every odd period, and Nature sends a response in the following even period until an application is accepted. Specifically, at $t = 1$, the student applies to some school $s_j \in S$, and at $t = 2$ learns whether the school has accepted her or not. In the case of acceptance, the game ends at $t = 2$. In the case of a rejection, at $t = 3$ the student applies to a different school that has not previously rejected her (i.e., some $s_{j'} \in S/s_j$) and learns about the result at $t = 4$. The game continues in the same alternating fashion until an application is accepted.¹² Under our assumptions, an application sent to s_j is accepted iff $\rho_j \geq T_j$. Notice that any set of on-path equivalent strategies in DP can also be represented by a ROL, which determines the order of (on-path) applications.

In the dynamic student-receiving (DR) variant, the student receives offers during odd periods and responds to offers during even periods. Nature's actions are determined by a set of (known) functions $\phi_j(\cdot): [0, 1] \rightarrow 2\mathbb{N} + 1$ which map priority scores at $s_j \in S$ to (odd) time periods in which the student receives an offer from s_j . The student receives an offer from s_j iff $\rho_j > T_j$. Otherwise, we have $\phi_j(\rho_j) = \infty$, which is interpreted as not receiving an offer from s_j during the matching process. While these functions depend on priority scores and other students' actions, it is always the case that all else equal, a higher priority score at a school is associated with an offer that arrives at a (weakly) earlier period. We denote the set of offers the student receives in each period by

$$\zeta_t = \{s_j : \phi_j(\rho_j) = t\}$$

and assume $\phi_{n+1}(\cdot) = 1$, i.e, the outside option is always available on the first period.

¹²Notice that since the outside option, where admission is guaranteed, is in S , one application will always eventually be accepted.

At $t = 1$, the student receives the set of offers ζ_1 (which always includes the outside option), then at $t = 2$, she must keep exactly one offer from ζ_1 . We denote the kept offer by $\kappa_2 \in \zeta_1$. At $t = 3$, she might receive offers from a disjoint set of schools, $\zeta_3 \subseteq S/\zeta_1$, and then at $t = 4$ she must keep exactly one offer out of $\kappa_2 \cup \zeta_3$ (i.e., the currently-held offer and the set of new offers). The game continues in this alternating fashion until the student receives no more offers, i.e., in the first period such that $\zeta_t = \emptyset$.¹³

A strategy l in the DR variant determines the offer the student retains out of $\kappa_t \cup \zeta_{t+1}$, given any history of prior offers and decisions. Notice that the strategy space in DR contains, but is not limited to, all possible ROLs, as not all strategies conform to a ROL.¹⁴

A.1.1 Preferred Personal Equilibrium

We adapt the backward-recursive definition of Kőszegi and Rabin (2009) of the PPE consumption plan to our setup:

Definition 3. Let z be a terminal history. We define the set of z -credible strategies in the following backward-recursive way: the credible set $\mathcal{L}_{z_t}^*$ includes all strategies that maximize utility at the terminal period given the expectation they induce, i.e., all strategies that satisfy $U_z(l|l) \geq U_z(l'|l)$ for all $l' \in \mathcal{L}$. Then the set $\mathcal{L}_{z_t}^*$ contains all strategies that satisfy $U_{z_t}(l|l) \geq U_{z_t}(l'|l)$ for all $l' \in \mathcal{L}_{z_{t+1}}^*$. I.e., all strategies that (i) maximize utility given the expectation they induce and (ii) prescribe a credible continuation plan.

A z -credible strategy maximizes utility given the expectation induced by the strategy: a student that expects to play l must find it optimal to follow through with l at every decision node in z , assuming any future self also plays optimally.¹⁵ The definition of a z -credible strategy allows us to define our solution concept:

Definition 4. Let $\mathcal{L}^* = \bigcap_{z \in \mathcal{Z}} \mathcal{L}_{z_0}^*$, i.e., the set of strategies that are credible in all terminal histories. A PPE strategy l^* is a strategy that satisfies

$$l^* \in \operatorname{argmax}_{l \in \mathcal{L}^*} U(l).$$

The core difference in this definition between static and dynamic variants in PPE is the possibility of commitment: The submission of a ROL in period 1 in a static variant can be seen as committing in advance to a strategy in its dynamic counterpart.¹⁶ In the dynamic variants, deviations are possible in each period, and therefore a strategy has to be optimal at all decision nodes (not just the first one). For example, under DP, a set of on-path equivalent strategies can be described as a ROL. However, for a ROL to be a PPE, there can be no profitable deviations at any possible point where a decision can be made: conditional on the first application prescribed by l (and given the belief induced by the continuation of l), the student must find it optimal to follow through and send the second application prescribed by l , and so on.

B Proof of proposition 1

Proposition 2. *proposition 1 Dynamic Student-Receiving (DR) DA variant is EBLA strategy-proof while the other three DA variants (SP, SR, DP) are not EBLA strategy-proof.*

¹³We assume here WLOG that in all terminal histories, the student receives offers in all odd periods preceding $\bar{t}_z$. If that is not the case, we can delete intermediate “waiting” periods and relabel the periods following them.

¹⁴In particular, treating the available offers $\kappa_t \cup \zeta_{t+1}$ as the state variable, a ROL cannot describe non-Markovian strategies.

¹⁵The forward-looking definition implies that when evaluating a strategy’s credibility, only credible deviations are considered, i.e., potential deviations that prescribe non-credible continuation strategies are not considered.

¹⁶In these variants, a PPE strategy coincides with the definition of *Choice-Acclimating Personal Equilibrium* (Kőszegi and Rabin, 2007).

Proof. We prove here the the DR variant is EBLA strategy proof. The proof for the other variants is a simple counter example that can be found in [Dreyfuss et al. \(forthcoming\)](#) and [Meisner and von Wangenheim \(2019\)](#).

Consider the DR implementation and denote the game implied by the *Matching problem* as $\Gamma(G, \mathbf{m})$, where $\mathbf{s}$ is the vector of schools. We let $\mathbf{q}$ denote the vector that contains, for each school, the unconditional probability of being above the cutoff. To simplify the presentation, we assume that these events are independent. It will become clear in the proof that introducing correlation has not effect on our result.

We will prove the proposition by showing, using backward induction, that for every sub-game Γ^h (defined by history h), SF is the only consistent plan implying that it is also the Preferred Personal Equilibrium in the Game $\Gamma(G, \mathbf{m})$.

Looking for an optimal strategy in the finite game $\Gamma(G, \mathbf{m})$, such a strategy has to assign an action to a history where the student receives her last offers so that there are no more schools that have not yet approaches the student. That is, the decision of the student in such history immediately decides the outcome of the mechanism. Consider the case, for example, where the student has already reacted to $n-1$ schools and now receives an offer from the last remaining school. Denote such time period as $t = \bar{t}$ as there is certainty that after this decision the game ends. We claim that any optimal strategy must assign a SF strategy for such a history.

When entering such a history, the student holds an offer from some school, denote it as k (or of course she does not hold any offer at all). If the student receives multiple offers in such a history, then taking the highest-valued school of these yields higher consumption utility as well as higher news-utility. As such, WLOG let us look at the case where the student only receives an offer from one school (her last remaining option), denoted as j . We split into two possibilities: $j < k$ or $j > k$ and show that SF must be optimal

In the first case the school j is a higher-valued school. A SF strategy implies choosing the offer from school j so this creates a belief $(m(s_j), 1)$ and actually implementing this yields utility of $m(s_j)$ ¹⁷. If indeed the student deviates and rejects such offer then the utility will be $m(s_k) - \lambda(m(s_j) - m(s_k))$, where the last term is the news-utility from a decrease in consumption utility. Clearly this means that SF strategy is a *personal equilibrium* in the subgame defined at such history where the student receives her last offer. Consider now a possibility of an optimal NSF in such subgame by staying with offer k , creating a belief $(m(s_k), 1)$. Actually rejecting the offer from school j yields $m(s_k)$, but accepting it yields $m(s_j) + (m(s_j) - m(s_k))$, which constitutes a profitable deviation. This implies that NSF strategy is not a *personal equilibrium* in said subgame.

In the second case the remaining school is a lowered-valued school. A SF strategy creates a belief $(m(s_k), 1)$ as the student plans to reject that offer, and following it leads to a utility of $m(s_k)$. A deviation to accepting such offer leads to a utility of $m(s_j) - \lambda(m(s_k) - m(s_j))$, so this it not a profitable deviation. Again this means that SF strategy is a *personal equilibrium* in the subgame defined at history where the student receives her last offer. Once again consider a possibility of an playing NSF, creating a belief $(m(s_j), 1)$ and actually accepting the offer yields $m(s_j)$, but rejecting it yields $m(s_k) + (m(s_k) - m(s_j))$, which constitutes a profitable deviation. Just as in the previous case, this implies that NSF strategy is not a *personal equilibrium* in said subgame.

Taken together, we can say that any optimal strategy must specify SF at $t = \bar{t}$. We now move backwards to a history where there are still schools which have not yet approached the student, but they might also not do so at all. After making the decision at $t = \bar{t} - 1$ (for example we can consider with $t = n - 1$, in which the student just receives her possibly next-to-last offer) said future strategy is known and taken into account in the formation of beliefs. More specifically, the student believes

¹⁷Note that while there is news-utility from receiving such an offer, we are focusing on the decision after receiving the offer.

that with probability q_j^h ¹⁸ an offer from school j will arrive so her reference point is either $(m(s_k), 1)$, if $j > k$, or $(m(s_j), q_j^h; m(s_k), 1 - q_j^h)$, if $j < k$. The proposition will immediately be proved from the following lemma:

Lemma 1. *Given any history h in which the student is called to act, if the remainder of the game follows a SF strategy (so that her choice now does not effect her future strategy), then SF is the unique personal equilibrium in the subgame Γ^h .*

Lemma 1 basically states that if the student knows that she will follow a SF strategy no matter what she does now, then it is best to also act SF at this point. Since we have proven that an optimal strategy must imply SF in a history where the student knows there is no continuation, then step-by-step lemma 1 will guarantee that the unique Personal equilibrium consists of playing SF at every possible history of the game.

Proof of lemma 1. Consider some history h and either choosing an offer according to SF, denote this by j , or choosing an offer according to some NSF, denote this by $k > j$ (for example when receiving an offer from multiple schools or having one offer from some previous round). Let the belief distributions $F_h(NSF) = (m(s_1), \tilde{r}_{1|h}; m(s_2), \tilde{r}_{2|h}; \dots; m(s_n), \tilde{r}_{n|h})$ and $F_h(SF) = (m(s_1), \hat{r}_{1|h}; m(s_2), \hat{r}_{2|h}; \dots; m(s_n), \hat{r}_{n|h})$, represent both the actual lotteries as well as the (consistent) reference points.

If the remainder of the game follows a SF strategy, this means of course that the choice at this juncture only effects what is the *kept* offer for the beginning of the rest of the game. The first implication of this is that since the continuation of the game follows SF, then the final outcome is guaranteed to be at least as high as the offer she will take in this period. Specifically this implies that, for every history h' that is derived from h (including h itself):

$$\tilde{r}_{i|h'} = \hat{r}_{i|h'} = 0 \text{ for every } i > k$$

That is, at every future juncture the student holds some zero probability to end up with worse-than- k offers, regardless if the student chose offer j or k . Next, looking at better-than- j offers, an SF continuation implies that the student will of course upgrade when she receives such an offer. Since taking offer j or k does not effect the probability of getting (and thus accepting) such better offers, then for every history h' that is derived from h (including h itself):

$$\tilde{r}_{i|h'} = \hat{r}_{i|h'} \text{ for every } i < j$$

We note that this is true even when the student (consistently) updates the probability of better-than- j offers downward in case any such offer does not arrive (say if the student keeps on not receiving an offer from the highest-values school, making her rationally pessimistic she will get such an offer soon).

Clearly the only difference in future actions is when the student receives an offer d , with $j < d < k$. This and the fact that $\tilde{r}_{i|h'} = \hat{r}_{i|h'} = 0$ for all $i > k$ means that the belief distributions under SF and NSF behavior at this point will only diverge on a closed range of percentiles $[0, \bar{p}_{h'}]$ on the cumulative distribution, where $\bar{p}_{h'} = 1 - \sum_{i=1}^{j-1} \tilde{r}_{i|h'} = 1 - \sum_{i=1}^{j-1} \hat{r}_{i|h'}$

For example consider the following figure, in which the student first receives offers from schools 2 and 4 and has to choose which one to keep. In any case taking each guarantees acceptance to such school, but it does not change to probability of receiving an offer from school 1.

We now claim that at such history h the sum of news utility derived from the continuation game is strictly higher when choosing SF than NSF, whether the student planned to play SF or NSF. Consider first a reference-point of playing SF at history h (for example the student holds an offer from school k but at history h she receives an offer from school $j < k$). Playing NSF at history h (by choosing offer k) results in immediate news disutility of

¹⁸For the purpose of the proof we only require that $q_j^h > 0$.




Figure 5: Example of a change in belief distribution between SF and NSF.

$$-\lambda \cdot \sum_{i=j+1}^k (m(s_j) - m(s_i)) \cdot \tilde{r}_{i|h}$$

This corresponds exactly to the shaded in area in figure 5 multiplied by the loss-aversion parameter λ . Looking at the continuation of the game and any possible stream of news utility, then we can think about four types of possible news:

1. *Bad news* - resulting from not receiving offers from better-than-school k but still worse-than-school- j schools. In our example, given NSF at history h (by keeping the offer from school 4) this is not receiving an offer from school 3 and so updating downward the probability for school 3: $\tilde{r}_{3|h'} < \tilde{r}_{3|h}$ so that $\tilde{r}_{4|h'} > \tilde{r}_{4|h}$ (by the exact same amount in this example). The following figure illustrates such change of “area”:




Figure 6: Example of a change in belief distribution after not receiving offers from better-than-school k but still worse-than-school- j schools

In such a case only under NSF this entails disutility. It should be noted that this can be counter-balanced by later positive news (such as getting a better-than-school- j offer), but since losses are counted with a loss-aversion parameter λ , any “area” that was first counted as a loss and then counted as a gain still has a negative sum of news-utility stream compared to the SF strategy.

2. *Bad news* - resulting from not receiving offers from better-than-school- j schools. For example, when at $t = \bar{t}$ the student learns that she officially has not received and will not receive such offers.

In this case the student updates downwards the probability to get better-than-school- j offers so that $\bar{p}_{h'} > \bar{p}_h$. Clearly playing NSF creates a larger amount of disutility as for every percentile that experience such a change (every percentile between $\bar{p}_{h'}$ and $\bar{p}_h$), the drop is towards a school that is worse than school j , invoking a higher amount of disutility. In our example, figure 6 shows such a drop and while the upper diagonally-shaded area captures the disutility from playing SF and not getting (yet) an offer from school 1, playing NSF adds the lower diagonally-shaded area as well. As before, such an added area may be later compensated by actually receiving the offer from school 1, but again, this area will be counted as a loss and a gain, which together is negative since $\lambda > 1$.




Figure 7: Example of a change in belief distribution after not receiving from better-than-school- j offers

3. *Good news* - resulting from receiving a better-than-school- k , but still worse-than-school- j offers. In such a case, only if the student chose NSF at history h she gains positive news utility. But, such good news only compensates for part of the loss at history h when she chose k , and because $\lambda > 1$ the total effect is negative. In our example, after choosing school 4 over school 3, it is receiving an offer from school 3 (and taking it). This compensates part of the shaded area in figure 5 but not all of it. Taking into account the possibilities of bad news covered just now, it is clear that the “best case scenario” from the perspective of playing NSF is when the student (immediately) receives good news in the form of a better-than-school- j offer¹⁹. For example, right after taking the offer from school 4 the student receives an offer from school 1. This is covered next.

¹⁹We wish to highlight that with no downward updating (the probability of receiving high-valued offers), news utility has an aggregate nature so that continuously receiving good news yields the same news utility as just receiving the latest good news. With downward updating it is strictly preferable to immediately receive the latest good news, so that there is only a single update.

4. *Good news* - resulting from receiving a better-than-school- j offers. In such a case the student will of course accept such an offer. This type of offer does indeed invoke higher news-utility when the student accepts such an offer, but only for the first better-than-school- j offer received. The maximal difference in news utility, is unsurprisingly:

$$\sum_{i=j+1}^k (m(s_j) - m(s_i)) \cdot \tilde{r}_{i|h}$$

Which in our example corresponds to compensating the entire shaded area in 5. Because of loss-aversion this is of course lower than the negative news-utility suffered from choosing to play NSF. Intuitively, the idea is that if the student expects to play SF, then any future news-utility gains are negated by the immediate drop in news utility.

Thus, given that choosing SF at history h yields a higher expected consumption value, as well as higher expected sum of news utility, then choosing SF is a *personal equilibrium* at history h .

Looking at a reference point of NSF, then the immediate news-utility gain from deviating to SF is given by:

$$\sum_{i=j+1}^k (m(s_j) - m(s_i)) \cdot \tilde{r}_{i|h}$$

The analysis of this is very much similar to the previous case. More specifically, the “best case scenario” from the NSF standpoint is again receiving a better-than-school- j offer yielding a maximal difference in future news-utility of exactly $\sum_{i=j+1}^k (m(s_j) - m(s_i)) \cdot \tilde{r}_{i|h}$. It should be noted that choosing SF and taking the higher-valued offer (in our example school 2’s offer) does not create more uncertainty as the probability to get better-than-school- j offers is the same regardless. But, since the probability to actually get better-than-school- j offers is not 1, then deviating to SF yields both higher expected consumption utility as well as higher expected news-utility.

For example, if the student knows for certainty that school 1 will indeed give her an offer sometime in the future, then it does not matter if she chooses school 2 over school 4 or vice versa. But, since the probability of an offer from school 1 is never 1, then even with a reference point of NSF, it is best to take school 2’s offer (and gain some utility for sure) rather than taking school 2’s offer.

Concluding, this implies that any NSF is not a *personal equilibrium* in subgame Γ^h (given of course that the rest of the game is SF), which implies that at history h SF is the unique a *personal equilibrium* at history h and as such it is also the *preferred personal equilibrium* at the subgame Γ^h . QED

Given this lemma and that in a history where the student knows she is at $\bar{t}$ by having no more schools left, any optimal strategy must be SF, then we conclude that the only *preferred personal equilibrium* of the entire game is SF. Specifically, SF strategy is the only *personal equilibrium* of the game $\Gamma(G, \mathbf{m})$.

QED

C Experimental Design: Further Details

Table 2 presents the ten matching problems given to each subject, held fixed across treatments. Each problem is defined as five school values and (unconditional) probabilities of acceptance. Columns 8–10 describe the population prediction for the prevalence of NSF strategies under all treatments (recall that our large market property implies that the two static treatments are equivalent, so their predictions coincide). For example, in problem 1, 30.56% of the population is estimated to have λ greater than 3.1 so that the optimal strategy for them in the *static* treatments is not the ROL 12345. We now explain our empirical predictions and use of alternative measures (columns 11–15).

		Matching Problem					Predictions						
							NSF (%)		Costly NSF (%)		DO NSF (%)		
#		s_1	s_2	s_3	s_4	s_5	DR	SP/SR	DP	SP/SR	DP	SP	SR
1	Values	\$1.5	\$1	\$0.75	\$0.5	\$0.25	0%	30.56%	26.39%	7.25%	8.34%	30.56%	7.25%
	Probabilities	0.05	0.2	0.25	0.9	1							
2	Values	\$1.5	\$1	\$0.75	\$0.5	\$0.25	0%	30.56%	23.61%	7.05%	9.13%	30.56%	7.05%
	Probabilities	0.05	0.2	0.3	0.4	1							
3	Values	\$1.5	\$1	\$0.75	\$0.5	\$0.25	0%	29.17%	29.17%	3.69%	3.7%	29.17%	3.69%
	Probabilities	0.05	0.1	0.85	0.9	1							
4	Values	\$1.25	\$1	\$0.75	\$0.5	\$0.25	0%	23.61%	17.5%	3.78%	3.5%	23.61%	3.78%
	Probabilities	0.2	0.8	0.9	0.95	1							
5	Values	\$1.5	\$1	\$0.75	\$0.5	\$0.25	0%	20.83%	18.47%	4.47%	4.62%	20.83%	4.47%
	Probabilities	0.25	0.8	0.85	0.95	1							
6	Values	\$1.25	\$1	\$0.75	\$0.5	\$0.25	0%	19.44%	11.67%	3.65%	2.92%	19.44%	3.65%
	Probabilities	0.25	0.75	0.8	0.85	1							
7	Values	\$1.25	\$1	\$0.75	\$0.5	\$0.25	0%	18.47%	13.61%	4.71%	4.08%	18.47%	4.71%
	Probabilities	0.3	0.85	0.9	0.95	1							
8	Values	\$1.25	\$1	\$0.75	\$0.5	\$0.25	0%	0%	0%	0%	0%	0%	0%
	Probabilities	0.65	0.7	0.85	0.95	1							
9	Values	\$1.5	\$1.25	\$1	\$0.5	\$0.25	0%	0%	0%	0%	0%	0%	0%
	Probabilities	0.65	0.75	0.8	0.85	1							
10	Values	\$1.25	\$1	\$0.75	\$0.5	\$0.25	0%	0%	0%	0%	0%	0%	0%
	Probabilities	0.6	0.65	0.8	0.9	1							

Table 2: Population theoretical predictions are based on empirically estimated distribution of λ from [Dreyfuss et al. \(forthcoming\)](#). Each of the columns s_1 – s_5 describes the money value of the school (top row) and the unconditional probability of acceptance (bottom row) in each of the ten matching problems.

We label matching problems 1–3 “weak-student” problems, where the probability of acceptance at the highest-valued school is small (5%), which implies a large predicted share of NSF strategies (24%–31% in the non-DR “NSF (%)” columns). We label problems 4–7 “medium-student” problems, where the probability of acceptance at the highest-value school is larger (20%–30%), with a lower predicted share of NSF strategies (12%–24%). Finally, problems 8–10 are labeled “strong-student” where the probability of acceptance at the highest-value school is high, which implies no predicted NSF strategies (under the past empirically estimated distribution of λ).

C.1 Alternative Measures

Table 2’s four rightmost columns report theoretical predictions using two alternatives to our primary (“NSF (%)”) measure. These alternative measures allow for a more direct comparison across treatments. However, as we discuss below, the theory predicts only relatively small cross-treatment differences in these measures—differences that our experiment does not have power to detect.

The first measure is *costly NSF* behavior, which measures NSF behavior that is also payoff relevant (i.e., SF-inconsistent behavior that had an effect on the subject’s final match). This measure allows for direct comparisons between all four treatments. Clearly, under the DR treatment, we predict no such mistakes; however, as columns 11–12 in table 2 show, the predicted share for this measure varies little across both matching problems and treatments (and remains always below 10%).²⁰

The second measure is *Dynamically-observable (DO) NSF* behavior, which counts NSF ROLs in the static treatments that would have been observed as SF-inconsistent in their dynamic counterpart

²⁰A potential alternative to costly NSF is actual earnings; this measure is omitted from table 2 but is presented in the online appendix.

treatment. For SP, these are ROLs that, if implemented as a sequence of applications to schools, would result in SF-inconsistent behavior. For SR, these are ROLs that, if implemented as keep/reject responses to sequentially arriving school offers, would have been classified as SF-inconsistent.

As seen in columns 14-15 of table 2, the theory predicts that every NSF strategy under SP is also dynamically observable. In contrast, every NSF strategy under SR that is not costly is not dynamically observable. This implies that using this measure for comparing SP and DP will have little information for testing the theory (as the model predicts no difference). Similarly, just like costly NSF, our experiment is under powered for using this measure to compare SR and DR.²¹

D Results: Further Details

D.1 Samples

Sample 1: Cornell BSL We run the experiment using the Cornell BSL (Business Simulation Lab) SONA system, an online pool of Cornell students. We planned to run 200 subjects (80 in the DR treatment and 40 in each of the other three treatments²²). We closed the experiment after 206 subjects completed it (out of 223 who clicked on the experiment link). Out of these 206 subjects, 198 answered the screening question correctly. The average total payoff (for all 206 subjects), including the participation fee, was \$14.23.

Each subject who clicked on the link to the experiment could participate and finish the experiment with no time constraints. The median completion time for the 206 subjects was 15.61 minutes (interquartile range 12.82-21.30 minutes). Since there were no time constraints and subjects could revisit the experiment, 13 subjects took more than an hour to finish the experiment. We omitted two subjects who took longer than an hour to finish the experiment from the end of the tutorial²³, as well as one subject who finished the entire experiment in less than 5 minutes (and who unsurprisingly failed the attention question). Altogether, we have an analysis sample of 196 subjects (77 in the DR treatment, 39 in the SR treatment, and 40 in each of the SP and DP treatments).

Sample 2: Previous-IPMM Participants We run the experiment on 2021 and 2020 Israeli Psychology Matching Mechanism (hereafter IPMM) participants, having already experienced a deferred-acceptance matching mechanism when assigned to a graduate program in Psychology. We approached the pool of the 2021 and 2020 IPMM participants through individual emails inviting them to also participate in our experiment.

In the initial wave, we sent 1096 invites to the 2021 pool, and of these, 226 clicked on the experiment link, and 171 finished the experiment. We sent 1206 invites to the 2020 IPMM participant's pool (excluding those who participated in the 2021 IPMM). Of these, 215 clicked on the experiment link, and 138 finished the experiment. Out of the 309 subjects who finished the experiment, 304 answered the screening question correctly. The average total payoff (for all 309 subjects), including the participation fee, was approximate \$16.43.²⁴

²¹The driver behind this predicted pattern is the fact that the NSF strategies predicted by the model in the non-DR treatments always involve the highest-value schools. See [Meisner and von Wangenheim \(2019\)](#)'s characterization.

²²In order to minimize the extent of possible imbalanced treatments, for every five-subject intervals we randomized the order using of two DR and one of each of the other treatments.

²³We included 11 of these 13 subjects as the time per page information suggests that they presumably read the first page but went back to experiment only later.

²⁴The main driving force behind the increase in the average payoff was the exchange rate. In order to use whole numbers to be presented to subjects, we chose to use an exchange rate of 4 NIS to a dollar, though the real exchange rate was around 3.2. We also note that we wanted the payment method to be as private as possible because we approached these subjects directly. Each subject was given a unique identifier to match her payoff upon starting the experiment. Subjects were told that upon finishing the experiment, they had to send this identifier as a message from any phone they chose to a specific phone (a research assistant's phone) who would match this identifier to the payoff and transfer it amount using a money transfer app. Out of the 309 subjects, 300 followed these orders and were paid without any intervention. For the nine remaining subjects, we chose to identify them and send them private emails reminding them of

The median completion time for 308 out of the 309 subjects was 17.80 minutes²⁵ (interquartile range 15.30-21.87 minutes). Nine subjects took more than an hour to finish the experiment, but since none took more than an hour starting from the end of the tutorial, we included them all. This resulted in an analysis sample of 304 subjects (126 in the DR treatment, 60 in the SR treatment, and 59 in each of the SP and DP treatments). All instructions and payoffs for the second sample were translated to Hebrew as equivalent as possible.

D.2 By-sample Analysis

Non-straightforward (NSF) behavior by Sample




Figure 8: Empirical NSF shares by sample. Error bars represent standard errors calculated from regressing NSF behavior on treatment, clustering at the individual level.

D.3 Alternative Measures

E Discussion of Bias in DR

Examining the DR treatment by itself, 338 out of 2030 observations (16.7%) were classified as observed NSF behavior, suggesting that both the classical and the EBLA model fail to characterize a significant portion of subjects. Moreover, this measure of observed NSF behavior is only a lower bound for the chosen NSF strategies in this treatment, and this is suspected to be more relevant as the student problem type is weaker (since in these problems, we observe fewer decisions, and so presumably, fewer ways in which to rule out chosen SF). This raises the concern that not only do our main between-treatments findings confounded but also that our within-treatments finding in the other treatment can perhaps be explained by other means as they might show up in the treatment in which we do not predict it.

how to contact us and receive their payment, and six of them cooperated. Three subjects did not comply despite repeated emails and were not paid for the experiment. We did, however, use their observations.

²⁵A technical error caused us to lose time of completEachsince almost no observation was excluded by time restriction, we did not exclude this observation.

We argue that while the results for the DR treatment shown in panel (b) of figure 8 are downward biased, there is some evidence that this is not biased to the degree that should negate our main findings. Our main argument roughly goes as follows: if the chosen strategies under the DR treatment are such that they negate our main findings, then it would create ex-post findings we do not observe.

First, we note that in the DR treatment, each school might send the subject an offer in any of five application rounds, if at all. Looking at the 338 observed NSF behaviors classified in the DR treatment, 249 of them (75%) are those in which the subject declined to keep any offer from those she initially received. Typically, these mistakes are passing on an offer from a lower-valued school.

This behavior, roughly corresponding to a ROL of sort 12XXX, is not predicted by the EBLA model, not even for the other treatments. Moreover, such omitting of lower-valued alternatives is barely evident in the other treatments, suggesting that these types of mistakes are unique to the DR treatment.²⁶ As the probability of first receiving only offers from higher-valued schools is minimal, we view these mistakes as unavoidable in the sense that a 12XXX chosen strategy would, with a high probability, be observed and classified as an NSF behavior.

Next, one should be worried that the rest of the mistakes result from a chosen NSF strategy similar to those in the other treatment, namely a strategy consistent with a ROL of sort 21345. This might imply that, for example, an observation where the subject did not receive an offer from the highest-valued school would not be classified as an observed NSF behavior. Simply put, we do not give the subjects enough options to “make a mistake”.

In order to examine this, we classified all 2030 DR observations by the number of offers received. As table 3 indicates, we do not find that more offers lead to more observed NSF behavior. We note that the lowest-valued school has an acceptance probability of one, implying that if the subject received only one offer, it was from that school. The higher fraction of NSF behavior when the subject only received one offer can be attributed to the above mistake.

Number of offers	Number of observations	Number of NSF-behavior observations	Observed NSF behavior fraction
1	95	25	0.26
2	294	51	0.17
3	592	94	0.16
4	743	122	0.16
5	306	46	0.15

Table 3: Distribution of DR observations by the number of offers the subject received

Lastly, we classified all 2030 DR observations into those in which the subject received an offer from the highest-valued school and those in which she did not. Out of 621 observations of the former kind, 98 (16%) were classified as observed NSF behavior, while out of 1409 of the latter kind, 240 (17%) were classified as observed NSF behavior. If subjects followed a 21345 strategy for some reason, we would have observed a higher fraction of observed NSF behavior conditional on receiving an offer from the highest-valued school.²⁷

Taken altogether, it seems that while observed NSF behavior measure is biased in the DR treatment²⁸, it does not seem likely to undermine our main findings.

²⁶Though unexpected and unpredicted, we think that these mistakes do capture important cognitive reasoning. Consistent with some feedback we got from subjects, we suspect that accepting a lower-valued offer signals that the subject is willing to compromise.

²⁷Clearly, there is no reason for a subject in the DR to choose the strategy consistent with the ROL 21345. It is even a strategy-proof Li (2017) strategy to take such an offer.

²⁸It should also be noted that the dynamic feature of DR might also lead to an opposite bias. Specifically, the DR treatment requires the subjects to make multiple decisions instead of submitting a ROL. If subjects exhibit trembling hand behavior wherein each node they make an error with some probability, then this might lead to more observed NSF behavior than expected.

E.1 Alternative Measures: Results

While our primary measure is the fraction of observed NSF behavior, we are also interested in other, readily comparable measures. As shown in figure 9, we predict a small percentage of such costly NSF behavior with slight variance between and within treatments. Empirically, we find a slightly higher-than-expected level of costly NSF behavior, with no substantial difference across any treatment and problem types.




Figure 9: Summary of the predictions and results for the costly NSF measure. Panel (a) presents costly NSF predictions by problem type and treatment under classical preferences (downwards triangle) as well as EBLA preferences (upwards triangle), using [Dreyfuss et al. \(forthcoming\)](#) calibration. Panel (b) presents the empirical shares by problem type and treatment. Error bars represent standard errors calculated from regressing costly NSF behavior on treatment, clustering at the individual level.


It is important to note that while in the *DR* treatment we observe less NSF behavior, each of these has, on average, a higher probability of being a costly mistake compared to the rest of the treatments.²⁹

Another measure of interest is Dynamically-observable NSF, which is defined as mistakes in the *static* treatments, that would have been observable in the corresponding *dynamic* treatment. As figure 10 shows, the pattern of behavior in the *static* treatment aligns with the EBLA predictions. Specifically, the NSF ROLs subjects submitted are dynamically-observable in the *SP* treatment and not in the *SR* treatment. A ROL of 21345 is observable NSF behavior if applied to the *DP* treatment but most likely not when applied to the *DR* treatment.

²⁹For comparison, the most common type of mistake under the static treatments is a ROL of 21345 which means that ex-ante, it is a costly mistake with probability $p_1 \cdot p_2$, where these probabilities are the unconditional probabilities of acceptance by the highest-valued and second highest-valued, as shown in the table (for *matching problem 1* it is simply $0.05 \cdot 0.2 = 0.01$). Conversely, the most common type of mistake under the *DR* treatment is declining an offer at the first round from the lowest-valued school, which is, ex-ante, costly with probability $\prod_{i=1}^4 (1 - p_i)$, which for *matching problem 1* is $0.95 \cdot 0.8 \cdot 0.75 \cdot 0.1 = 0.057$.

Dynamically-observable (DO) NSF behavior

(a) Theoretical Predictions: Classical (▼) & EBLA (▲) Preferences



(b) Empirical Shares (■)




Figure 10: Summary of the predictions and results for the dynamically-observable NSF measure. Panel (a) presents dynamically-observable NSF predictions by treatment and problem type under classical preferences (downwards triangle) as well as EBLA preferences (upwards triangle), using [Dreyfuss et al. \(forthcoming\)](#) calibration. Panel (b) presents the empirical frequencies by problem type and treatment. Error bars represent standard errors calculated from regressing DO NSF behavior on treatment, clustering at the individual level.

F Experiment Screenshots


The experiment consists of the following parts:

1. Consent form
2. General instructions
3. Specific instruction (for each treatment)
4. Tutorial
5. Attention question
6. Ten order-randomized *matching problems*
7. Demographic questions
8. Feedback

We provide screenshots for the **General instructions**, **Specific instructions** and an example for the interface in a **matching problem**. These screenshots are color-coded (only for illustration here) to highlight the differences between treatments. We will use the following color scheme:

All pages shown here will have a color-coded symbol at the **bottom right corner**, signifying for which treatment it is relevant. For example the symbol  represents an entire page that is the same in all treatments, while the symbol  represents an page that is **unique** to the *DP* treatment.

F1 General Instructions



Color-coding for screenshots

This is a study about decision making. Based on your decisions, you might earn a considerable amount of money **in addition** to the \$4 participation payment. In this study, we simulate a procedure to allocate students to schools. The procedure, payment rules, and student allocation method will be described in the next few screens. Everything will be exactly as specified in the instructions.

Next >



General instructions - page 1

Imagine that there are five different schools, and you are one of many candidates considering applying this year. These schools differ in size, geographic location, topics taught, and quality of instruction. Each school has a dollar **value** that reflects all these qualities. The *value* of the school you are matched with will be added to your payment.

< Back Next >




General instructions - page 2

E.2 Specific Instructions

The following four screenshots present the instruction page that differs between treatments. These pages also include color-coded symbols for specific lines (**left to the specific line**), indicating whether this line is unique for this treatment, or whether it is the same across treatments. For example the symbol  represents a specific instruction that is the same for both *SP* and *DP*.

School seats are limited, and each school can accept only a small share of the candidates applying.

Each school gives you, the candidate, a **priority score** which is chosen at random between 0 and 99 (every round number between 0 and 99 is equally likely to be chosen). *Priority scores* represent your rank in a school relative to the rest of the candidates. For example, if your *priority score* at some school is 78, it means that you are ranked higher than 78 percent of the candidates, and lower than or equal to the other 22 percent.



Please note: your *priority score* at each school is chosen randomly and separately, therefore, your *priority scores* at different schools are not connected to each other in any way.

Each school has an **acceptance threshold**. Schools are only willing to accept candidates whose *priority score* is greater than or equal to this *threshold*. For example, a school with a *threshold* of 75 is willing to accept you if your *priority score* is 80 and not willing to if your *priority score* is 70.

< Back

Next >



General instructions - page 3

E3 *Matching problem* example

We provide here an example for a *matching problem* for each of the treatments. In each treatment the participant was introduced to the problem, and then interacted with a treatment interface, and lastly got feedback regarding the outcome as well as the realization of **priority scores**.

For each participant we randomized the order of the *matching problems*. For each participant and every *matching problem* we randomized, in all tables (such as those shown in the following pages), the order of schools. We also randomized for each participant whether the **Value** column will be shown to the left or right of the **Threshold** column.

In this study you will participate in 10 rounds. Each round will simulate the matching process to 5 schools. In each round, you can be matched with at most 1 school.

At the beginning of each round, you will know the *value* of each school and its *threshold*, but you won't know your *priority score* until the end of that round. Note: since your *priority score* at each school is a random number between 0 and 99, a school's *threshold* conveys the probability it is willing to accept you (that is, the chance that your *priority score* is greater than or equal to its *threshold*). For example, if some school's *threshold* is 25, then there is a 75% chance that your *priority score* is greater than or equal to this *threshold*, meaning a 75% chance they will be willing to accept you.

In each round, before the matching process begins, you will be presented with a table similar to this one. This table will also be visible to you throughout the process.

School	Value	Threshold	Your Priority Score	Chance that Your Priority Score $\geq$ Threshold
Pine Peak	\$0.75	60	?	40%
Birch Hill	\$0.25	0	?	100%
Hickory Bridge	\$0.50	50	?	50%
Maplecrest	\$1.25	80	?	20%
Elm South	\$1.00	70	?	30%

In this example, Elm South's *threshold* is 70, which means that the chance that your *priority score* is greater than or equal to this *threshold* is 30%. If you are matched with Elm South at the end of the matching process, you will receive \$1.00 in addition to your \$4.00 participation payment.

< Back

Next >



Description of the Matching Process

- Before the matching process starts, you will see a list of schools, their *thresholds*, and their associated *values*.
- Throughout the matching process we will send applications to schools on your behalf.
- If your *priority score* at a school is greater than or equal to that school's *threshold*, then that school will accept your application.

The matching process will proceed as follows:

- You will be prompted to rank any of the five schools in any order of your choice.
- At the first stage we will send an application to the top school, according to your ranking.
- If your *priority score* is lower than the school's *threshold*, your application will be rejected, and we will continue sending applications to schools according to the order you ranked them.
- This process will continue until an application is accepted, or until there are no more schools left on your list.
- Keep in mind that we will not send an additional application to a school that has previously rejected your application.
- Your final match will be the school that accepted your application. Your earnings for that round will be the *value* of that school.

In order to familiarize you with the matching process, you will now participate in a quick tutorial. Notice that **the sums of money in the tutorial are hypothetical**. During this short tutorial, we will walk you through the matching process step by step. *Explanations in this font and color* are a part of the tutorial. You will not see them during the real matching process.

When you are ready to start the tutorial press "Start Tutorial".

< Back

Start Tutorial

SP

Description of the Matching Process

- Before the matching process starts, you will see a list of schools, their *thresholds*, and their associated *values*.
- Throughout the matching process we will respond to offers from schools on your behalf.
- If your *priority score* at a school is greater than or equal to that school's *threshold*, then you will receive an offer from that school, at some point.

The matching process will proceed as follows:

- You will be prompted to rank any of the five schools in any order of your choice.
- At the first stage you might receive offers from some schools. We will keep the offer of the highest-ranked school among these, according to your ranking, and the rest will be rejected.
- At every subsequent stage you might receive additional offers from some schools. From among these offers and the offer you are currently holding, we will keep the offer of the highest-ranked school (according to your ranking) and reject the rest.
- This process will continue until there are no more offers.
- Keep in mind that a school whose offer has been rejected will not send you any more offers.
- Your final match will be the school whose offer you are holding at the end of the round. Your earnings for that round will be the *value* of that school.

In order to familiarize you with the matching process, you will now participate in a quick tutorial. Notice that **the sums of money in the tutorial are hypothetical**. During this short tutorial, we will walk you through the matching process step by step. *Explanations in this font and color* are a part of the tutorial. You will not see them during the real matching process.

When you are ready to start the tutorial press "Start Tutorial".

< Back

Start Tutorial

SR

Specific instructions - SP and SR treatments

Description of the Matching Process

- Before the matching process starts, you will see a list of schools, their *thresholds*, and their associated *values*.
- Throughout the matching process you will send applications to schools.
- If your *priority score* at a school is greater than or equal to that school's *threshold*, then that school will accept your application.

The matching process will proceed as follows:

- At the first stage you can send an application to any school you choose.
- If your *priority score* is lower than the school's *threshold*, your application will be rejected, and you can continue sending applications to the remaining schools in any order you choose.
- This process will continue until an application is accepted, or until there are no more schools left. You can choose to stop the process at any point (in which case, you will not be matched with any school).
- Keep in mind that you cannot send an additional application to a school that has previously rejected your application.
- Your final match will be the school that accepted your application. Your earnings for that round will be the *value* of that school.

In order to familiarize you with the matching process, you will now participate in a quick tutorial. Notice that **the sums of money in the tutorial are hypothetical**. During this short tutorial, we will walk you through the matching process step by step. **Explanations in this font and color** are a part of the tutorial. You will not see them during the real matching process.

When you are ready to start the tutorial press "Start Tutorial".

< Back

Start Tutorial

DP

Description of the Matching Process

- Before the matching process starts, you will see a list of schools, their *thresholds*, and their associated *values*.
- Throughout the matching process you will respond to offers from schools.
- If your *priority score* at a school is greater than or equal to that school's *threshold*, then you will receive an offer from that school, at some point.

The matching process will proceed as follows:

- At the first stage you might receive offers from some schools. You can keep at most one of these offers and the rest will be rejected.
- At every subsequent stage you might receive additional offers from some schools. From among these offers and the offer you are currently holding, you can then keep at most one offer, and the rest will be rejected.
- This process will continue until there are no more offers.
- Keep in mind that a school whose offer has been rejected will not send you any more offers.
- Your final match will be the school whose offer you are holding at the end of the round. Your earnings for that round will be the *value* of that school.

In order to familiarize you with the matching process, you will now participate in a quick tutorial. Notice that **the sums of money in the tutorial are hypothetical**. During this short tutorial, we will walk you through the matching process step by step. **Explanations in this font and color** are a part of the tutorial. You will not see them during the real matching process.

When you are ready to start the tutorial press "Start Tutorial".

< Back

Start Tutorial

DR

Specific instructions - DP and DR treatments

You will now participate in an additional matching process. The money you earn from holding a slot at a school at the end of this matching will be **added** to what you have earned so far in the study. The rules are identical.

Description of the Matching Process

Note

Schools' *values* and *thresholds* in this matching process are presented in the following table:

School	Value	Threshold	Chance that Your Priority Score $\geq$ Threshold
Fig Point	\$1.25	75	25%
Grapevine Pass	\$0.25	0	100%
Orange Terrace	\$1.00	25	75%
Citrus Park	\$0.75	20	80%
Juniper Square	\$0.50	15	85%

Press "Continue" to start the matching process

Continue

Round: 7/10



Figure 13: *matching problem* example

Round: 7/10

Please rank any of the five schools in any order of your choice.
(click "Submit" when you are done):

Description of the Matching Process

First-ranked School:

Second-ranked School:

Third-ranked School:

Fourth-ranked School:

Fifth-ranked School:

School	Value	Threshold	Your Priority Score	Chance that Your Priority Score $\geq$ Threshold
Fig Point	\$1.25	75	?	25%
Grapevine Pass	\$0.25	0	?	100%
Orange Terrace	\$1.00	25	?	75%
Citrus Park	\$0.75	20	?	80%
Juniper Square	\$0.50	15	?	85%

SR SP

The matching process has ended!

You ranked the schools in the following order: **Fig Point**, **Grapevine Pass**, **Orange Terrace**, **Citrus Park** and **Juniper Square**.

Your final match is: **Grapevine Pass (\$0.25)**

School	Value	Threshold	Your Priority Score
Fig Point	\$1.25	75	18 (Below threshold)
Grapevine Pass	\$0.25	0	3 (Above threshold)
Orange Terrace	\$1.00	25	23 (Below threshold)
Citrus Park	\$0.75	20	75 (Above threshold)
Juniper Square	\$0.50	15	94 (Above threshold)

Continue

Round: 7/10

SR SP

matching problem interface and results - SP&SR treatments

Please choose where to send your next application:

Fig Point

Grapevine Pass

Orange Terrace

Juniper Square

Citrus Park

Stop matching process (you will not match with any school)

Schools that have already rejected you: **Fig Point**

Round: 5/10

DP

School	Threshold	Value	Your Priority Score	Chance that Your Priority Score $\geq$ Threshold
Fig Point	75	\$1.25	?	Originally: 25%. Now: 0%
Grapevine Pass	0	\$0.25	?	Originally: 100%
Orange Terrace	25	\$1.00	?	Originally: 75%
Juniper Square	15	\$0.50	?	Originally: 85%
Citrus Park	20	\$0.75	?	Originally: 80%

The matching process has ended!

You sent applications to schools in the following order: **Fig Point** and **Grapevine Pass**.

Your final match is: **Grapevine Pass (\$0.25)**

School	Threshold	Value	Your Priority Score
Fig Point	75	\$1.25	40 (Below threshold)
Grapevine Pass	0	\$0.25	48 (Above threshold)
Orange Terrace	25	\$1.00	94 (Above threshold)
Juniper Square	15	\$0.50	61 (Above threshold)
Citrus Park	20	\$0.75	77 (Above threshold)

Continue

Round: 5/10

DP

matching problem interface and results - DP treatment

New offer received!

Description of the Matching Process

The new offer is from:

Citrus Park

Remember: this means that your Priority Score at this school is higher than this school's Threshold.

If you want to switch to this offer, click on it.

Otherwise, you can choose to reject the new offer (and keep holding the offer you currently hold): **Reject**

Offer you currently hold: **Orange Terrace**

Offers you already rejected: **Grapevine Pass**

School	Threshold	Value	Your Priority Score	Chance that Your Priority Score $\geq$ Threshold
Citrus Park	20	\$0.75	?	Originally: 80%. Now: 100%
Orange Terrace	25	\$1.00	?	Originally: 75%. Now: 100%
Juniper Square	15	\$0.50	?	Originally: 85%
Fig Point	75	\$1.25	?	Originally: 25%
Grapevine Pass	0	\$0.25	?	Originally: 100%. Now: 100%

Round: 2/10

DR

The matching process has ended!

Your Priority Scores at **Orange Terrace**, **Grapevine Pass**, **Citrus Park** and **Juniper Square** were all higher than their Thresholds, and therefore you received offers from all of them during the matching process.

Your final match is: **Orange Terrace (\$1.00)**

School	Threshold	Value	Your Priority Score
Citrus Park	20	\$0.75	55 (Above threshold)
Orange Terrace	25	\$1.00	91 (Above threshold)
Juniper Square	15	\$0.50	33 (Above threshold)
Fig Point	75	\$1.25	29 (Below threshold)
Grapevine Pass	0	\$0.25	78 (Above threshold)

Continue

Round: 2/10

DR

matching problem interface and results - DR treatment