המכון למחקר כלכלי בישראל
על-שם מוריס פאלק בע"מ (חל"צ)
The Maurice Falk Institute for
Economic Research in Israel Ltd.

Deferred Acceptance with News Utility

By

Bnaya Dreyfuss, Ofer Glicksohn,
Ori Heffetz and Assaf Romm

Discussion Paper No. 22.02

July 2022

בנין פרץ נפתלי, קמפוס האוניברסיטה העברית, הר הצופים, ירושלים 9190501
The Hebrew University Campus, MT. Scopus, 9190501 Jerusalem, Israel
www.falk.huji.ac.il

Deferred Acceptance with News Utility*

Bnaya Dreyfuss Ofer Glicksohn Ori Heffetz Assaf Romm

Monday 11th July, 2022

Abstract

Can incorporating expectations-based-reference-dependence (EBRD) considerations help reduce seemingly dominated choices in the Deferred Acceptance mechanism? We run an experiment ($N = 500$) where each participant is randomly assigned into one of four different variants of Deferred Acceptance—{static vs. dynamic} $\times$ {student proposing vs. student receiving}—and play ten different problems of a simulated school-assignment game under a large-market assumption. While a standard, reference-independent model predicts the same straightforward behavior across the ten problems and the four variants, an EBRD model predicts stark differences in behavior across variants and problems.

We find, as predicted by the EBRD model, that (i) across matching problems, deviations from straightforward behavior increase with the competitiveness of the setting; (ii) across variants, *dynamic student receiving* leads to significantly fewer deviations; and (iii) across both matching problems and variants, differences in *payoff-relevant* deviations are small (often non-detectable in our data).

*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 and Uber Technologies Inc. (e-mail: assaf.romm@mail.huji.ac.il). We thank Ophir Averbuch, Shengwu Li, Matthew Rabin, JESC-lab group and seminar participants at Cornell University, Hebrew University, and Harvard University for invaluable comments; Tal Asif, Ronny Gelman, Ilana Lavie and Yonatan Rahimi 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.

A growing body of empirical evidence documents puzzling behavior in strategyproof 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 strategyproof for students, i.e., submitting a straightforward ("truthful") ranking of schools is a weakly dominant strategy, a nontrivial 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; early field evidence includes Shorrer and Sóvágó (2017) from Hungary and Hassidim et al. (2021) from Israel. Similarly, in simple allocation experiments, a substantial fraction of participants rank smaller amounts of money above larger ones; see Rees-Jones and Skowronek (2018) for the first lab-in-the-field evidence from the US, and Hakimov and Kübler (2020) for a recent survey of lab evidence.

Recently, Dreyfuss et al. (forthcoming) and Meisner and von Wangenheim (2019) showed that expectations-based reference dependence (EBRD; also referred to as EBLA, for expectations-based loss aversion) could potentially explain such non-straightforward behavior. Intuitively, participants with EBRD preferences may intentionally downrank (or, in some cases, completely omit) a high-value, low-probability school to avert the likely disappointment from rejection. Both papers show how an EBRD model, as formulated by Kőszegi and Rabin (2006, 2007, 2009), is consistent with various empirical patterns from the lab and the field. The first paper also structurally analyzes existing lab data from Li (2017) and shows that the EBRD model indeed fits the data significantly better than the classical, reference-independent-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 EBRD 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

Table 1: Four Deferred Acceptance Variants

	Static: Submit 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**

schools in an order determined by rank-order lists (ROLs) they submit in advance. In the dynamic version of this variant, Dynamic student Proposing (DP), students 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. Finally, in the Dynamic student Receiving (DR) variant, students 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 (with a continuum of students and a finite number of schools), the DR variant is "EBRD-strategyproof" in that it eliminates non-straightforward behavior for any level of loss aversion. (As discussed below, in concurrent work, Meisner and von Wangenheim (2019) prove a related result.) In contrast, EBRD-driven non-straightforward behavior is predicted to be quite prevalent under the other three variants, especially in highly competitive settings where the chances of admission into top schools are low.

Second, in what we view as our main contribution, we experimentally test the model's predictions, both (a) across the four algorithm variants, by randomly assigning different participants to different variants, and (b) across matching settings, by exogenously varying the degree of competitiveness across ten matching problems that each participant faces.[1] In addition to providing strong empirical support to our DA-specific theoretical results above, our experiment is also unique in providing one of the sharpest tests to date of the

[1] In related work, Klijn et al. (2019) run a similar four-treatment experiment with the four DA variants. In contrast to our large-market framework, their experiment has four student subjects with full information about others' preferences, and four (non-strategic) schools. This design does not allow for comparisons across role designation (proposing/receiving), and does not lend itself to testing our EBRD model predictions for two main reasons. First, in non-large markets, switching role designation drastically changes the (classical-preferences) incentive structure. Second, under full information, the uncertainty—the essential part of the EBRD theoretical framework—is about other subjects' strategies, and is hence not easily estimated or modeled. Other related work with similar caveats includes Echenique et al. (2016), who run full-information, dynamic one-to-one DA. See Hakimov and Kübler (2020)'s review for further details.

EBRD model more generally. In particular, by holding everything fixed other than the timing of commitments and uncertainty resolution, (a) above provides a clean test of the model's timing predictions; and by holding everything fixed other than the competitiveness of the setting, (b) above directly tests a central prediction of the model regarding how changes in probability beliefs affect behavior.

Finally, for practitioners, who may care less about theoretical results and their empirical tests and more about pragmatics, our finding that of the four mechanism variants we examine, DR in fact stands out in minimizing (though far from completely eliminating) non-straightforward behavior has practical implications. As discussed below, we view this result as a valuable, "bottom-line" policy-relevant empirical finding regardless of its theoretical drivers.

We report our theoretical analysis in section 1. 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) and Abdulkadiroğlu et al. (2015). We show that in such economies, DR is EBRD-strategyproof: for any degree of loss aversion and any belief distribution, the EBRD 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 the student always picks the highest-value offer from the set of available offers in each step.

In contrast with DR, the other three variants are not EBRD-strategyproof: 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 1 thus leads to three sets of testable predictions under the EBRD model. First, *within* each variant, non-straightforward behavior should increase with the competitiveness of the matching setting in the three non-EBRD-strategyproof variants, but not in DR. Second, *between* the four variants, DR should stand out in having the lowest share of non-straightforward behavior in competitive environments, but not in noncompetitive ones (where shares should be similarly low in all variants). Third, non-straightforward behavior should consist of *specific strategies in specific environments* (precisely identified by the EBRD model).

These three sets of qualitative predictions, together with a given distribution of the degree of loss aversion in a population—which we plug in from previous work—yield specific quantitative predictions both across matching settings and across DA variants.

Our experiment is designed to evaluate them.

We outline our experimental design in section 2. Subjects are randomly assigned to one of the four variants, and play the same ten matching problems (but in different random orders). Each problem simulates a different large-market matching economy, with five schools whose values to participants are given in dollar amounts of takeaway money. Using the population distribution of loss aversion from a previous study, the ten matching problems are designed to create variation in the predicted prevalence, according to the EBRD model, of straightforward behavior in the three non-DR treatments. Three of the ten problems ("weak student" problems) are designed to simulate highly competitive settings, and predict a moderately high fraction of non-straightforward behavior (24–31 percent); another four ("medium student") problems predict a lower fraction of such behavior (12–24 percent); and the remaining three ("strong student") problems predict no non-straightforward behavior (0 percent). (Recall that in comparison, the standard, no-EBRD model predicts 0% of non-straightforward behavior under all four treatments and all ten matching problems.)

In section 3 we report experimental results from two independent samples, pooled together ($N = 500$): Cornell students ($N = 196$) and Israeli psychology graduate-school candidates ($N = 304$) who participated in the (DA-based) Israeli Psychology Master's Match (IPMM) and are therefore a highly relevant population. Our main predictions and results are summarized in figure 3 (page 24).

Our findings support all three sets of theoretical predictions. First, within the three non-DR treatments, the observed shares of non-straightforward behavior monotonically decrease from weak- (33–48 percent) to medium- (30–37 percent) to strong-student problems (19–22 percent). In contrast, within the (EBRD-strategyproof) DR treatment, the shares are essentially flat (19 to 16 to 16 percent). Second, these shares also imply that between the variants, we find substantially lower levels of observed non-straightforward behavior in DR compared to the other three treatments in competitive environments, and essentially no difference in noncompetitive environments. Third, moving beyond mere shares of non-straightforward behavior, we further show that (i) the specific non-straightforward strategies that the model predicts and does not predict, respectively, are indeed those we commonly observe and do not observe in the data; and (ii) in specific *problems* in which the model predicts specific non-straightforward strategies (e.g., ranking the third-highest-value school on the top of the list), these strategies are indeed the most common.

A potential worry when comparing behavior across static and dynamic variants (in our case, the static variants vs. DR) is that while in static variants we observe submitted rank-order lists—essentially, the complete strategy—in dynamic variants we only observe *actions*. In theory, this could mechanically downward-bias the observed share of non-straightforward behavior in DR. In practice, however, our evidence suggests that such potential bias is unlikely. Indeed, we find that in DR, (a) observed non-straightforward behavior primarily consists of rejecting offers from low-value schools—offers that we typically observe; and (b) observed non-straightforward behavior is no more prevalent in rounds where subjects receive more offers.

We close section 3 by discussing other theories that might explain our data. We make two observations. First, we find across all problems and variants a baseline share of 16–22 percent observed non-straightforward behavior that the model does not predict and cannot explain, and that is likely explained by other models of strategic confusion, misunderstanding, and noisy decision-making. Second, no model of misunderstanding or noise that we are aware of predicts or explains the three sets of qualitative predictions that our model predicts and that are borne out in the data: *variation* within and across variants, and prevalence of *specific* types of non-straightforward behavior.

We conclude in section 4, where we discuss the potential implications, as well as limitations of our findings. From a theoretical point of view, our findings are unambiguously more consistent with the EBRD than with a no-EBRD model—importantly, without adding degrees of freedom in the EBRD model (as we fixed parameter values in advance at a previously estimated distribution). At the same time, we always find 10–20% more non-straightforward behavior than the (parameter-constrained) EBRD model predicts—strongly suggesting that loss aversion cannot account for all such observed behavior. From a policy-maker's point of view, our findings suggest that one of the four implementations of DA that we examine—dynamic student receiving (DR)—minimizes such behavior. Of course, our study leaves open the pragmatic question of whether DR is feasible to implement in practice.

1 Theoretical Analysis

This section is notation-heavy and technical. Readers with powerful intuition may wish to skip it and move directly to the mostly self-contained section 2 (p. 15).

1.1 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 schools. Throughout this section, we rely on the large-market characterization of Azevedo and Leshno (2016) and Abdulkadiroğlu et al. (2015) and their adaptation of the DA mechanism to a continuum economy.[2]

The set of schools is denoted by $S = \{s_1, \ldots, s_{n+1}\}$, where s_{n+1} represents the outside option. Each student i has strict (reference-independent) preference over schools, represented by the vector of utilities $\mathbf{m}_i = (m_{i,1}, \ldots, m_{i,n+1})$. We call this utility component *consumption utility*. Later in this section we add a *news-utility* (Kőszegi and Rabin, 2009) component. Each school j has a capacity $c_j \in (0, 1]$ that represents the share of students it can admit and assigns each student i with a relative rank $\rho_{ij} \in [0, 1]$, which we call a *priority score*.

As shown in Azevedo and Leshno (2016), given students' and schools' submitted rankings, the match result from running DA in this economy is characterized by a vector of thresholds $\mathbf{T} = (T_1, \ldots, T_{n+1})$, one for each school, where each student i is assigned to her highest-ranked school such that $\rho_{ij} \geq T_j$. Moreover, student-proposing and student-receiving DA result in the same match (and associated thresholds vector $\mathbf{T}$). 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 $\mathbf{T}$ 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 assumption captures the idea that while students know how selective a school is, they do not necessarily know their exact ranking relative to other students.[3] 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 others' (i.e., schools' and other students') behavior and therefore we often suppress the index i. Moreover, WLOG we index schools by student i's preferences so that $m_{i,j} > m_{i,k}$ for all $j < k$ and we normalize $m_{i,n+1} = 0$.

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

[3] Notice that this framework can easily capture uncertainty about the cutoffs by simply incorporating it into G_{ij}.

As explained in the introduction we focus on four different *variants* of the DA algorithm, resulting from the product {proposing, receiving} $\times$ {static, dynamic}. We denote the DA variant governing the market by M. A *matching problem* from a student perspective can therefore be summarized by $\langle M, G, \mathbf{T}, \mathbf{m}\rangle$, where M defines the matching process, G is the joint distribution over priority scores, $\mathbf{T}$ is a vector of thresholds, and $\mathbf{m}$ is a vector of school consumption utilities.

1.2 Timing, Beliefs, and Strategies

Fixing all other players' actions, we can treat each DA variant M as a game student i plays against Nature, where G governs Nature's moves. We define periods for each variant below, but generally, we think of periods as decision nodes where either the student or Nature makes a move.

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 $G(\cdot \mid h)$ 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 index of that final match is represented by the outcome function $O(z) : \mathcal{Z} \to \{1, \ldots, n+1\}$, and the terminal period is denoted by $\bar{t}_z$. Subhistories are denoted using subscripts: z_t (with $t \leq \bar{t}_z$) 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 act, for every possible history. We denote the set of continuation strategies consistent with (i.e., do not preclude reaching) history h by $\mathcal{L}_h$. A history h induces belief $G(\cdot \mid h)$. This belief, combined with a variant M, thresholds $\mathbf{T}$, and a strategy l jointly induce a belief over terminal histories, which we denote by $\tilde{F}_{l|h}$, and also a belief over final payoffs (i.e., over the support $\mathbf{m_i}$) which we denote by $F_{l|h}$.

In appendix A we formally present the game our large-market framework induces under each of the four variants. In the static variants (SP and SR), students submit a ROL in the first period and learn about the result in the next one; in DP, they apply to a school in one period and learn the result in the following period, until the first acceptance; and in DR they receive a set of offers in one period, and can keep one of them in the following period, until no more offers are received.

We now introduce our definition of non-straightforward behavior ("misrepresentation" or "non-truthfulness"), which applies to all four variants.

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

In the static variants SP and SR, the SF strategy is simply the ROL that ranks schools by their values. In DP, the SF strategy is to sequentially apply to schools by order of their values. In DR, it is retaining the highest-valued offer at any given decision node for any given history.

1.3 Expectations-Based-Reference-Dependent (EBRD) Preferences

As mentioned earlier, in addition to classical consumption utility, students' utility functions have a news-utility component. Absent this additional news component, given any history h, the student's objective function is simply $\mathbb{E}_{F_{l|h}}[m]$, i.e., the expected consumption utility under strategy l given history h. The additional news-utility component is aimed to capture a basic feature in people's preferences: the utility and disutility of belief updating. In particular, learning about a higher likelihood of a good outcome—a pleasant surprise relative to previously held beliefs—in itself entails positive utility, while learning about a lower likelihood—a disappointment—in itself entails disutility. *Loss aversion* in this model, the disutility from a downward update of one's beliefs is larger than the utility from an equally sized upward update.

News utility is therefore 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 $F_{l|h_t}$ to $F_{l|h_{t+1}}$. Let F and F' be, respectively, previously held and updated belief distributions over outcomes. The news utility function is given by:

$$N(F' \mid F) = \int_0^1 \mu\left(F'^p - F^p\right) dp, \tag{1}$$

where F^p denotes the consumption level at percentile p of F, and $\mu(\cdot)$ is defined as:

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

with $\lambda \geq 1$ representing an individual's loss-aversion parameter. The function N describes how updated beliefs are compared with previously held beliefs, percentile by percentile:

in each period t, the student compares, for each percentile, the outcome under F' to the outcome under F, with a higher weight λ on negative surprises.

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

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

The first term is expected consumption utility. The last two terms are the expected sum of news utility streams from the terminal history, given a probability distribution over possible terminal histories $\tilde{F}_l$.[4] Note that beliefs over payoff prior to period 1, i.e., prior to choosing a strategy, $F_{l|\emptyset}$, are not yet defined. 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 $x \geq m_{n+1} = 0$ and 0 otherwise. This assumption is calibrationally (i.e., quantitatively), but not qualitatively, substantive. In particular, it does not affect our between-variant results nor change our qualitative predictions within each mechanism. However, it is calibrationally substantive in that assuming more optimistic initial beliefs will require a higher value of λ for the optimal strategy to be NSF.

Notice that this formulation implicitly assumes that utility from attending different schools belongs to the same consumption dimension, which captures situations where schools are differentiated vertically (rather than horizontally). Our between-variant result (proposition 1) does not hold if we relax it. Dreyfuss et al. (forthcoming) shows a within-variant result (similar to proposition 2) under complete horizontal differentiation.

To summarize, given a realization of a terminal history z and a chosen strategy l, timing in our model is defined as follows:

- **Period 0:** The student believes that she will consume the outside option with certainty. No action is taken.

- **Period 1:** The student learns about the mechanism and chooses a strategy l. If she is called to act, she takes the action prescribed to z_1 by l.[5] She updates her beliefs from

[4] In the original version of the Kőszegi and Rabin (2009) model there are two additional parameters: η, which captures the weight of news utility compared to consumption utility, and γ, which discounts news on future consumption. We assume that $\eta = 1$ as well as $\gamma = 1$. The first assumption is simply a normalization. The second assumption is more substantive and implies that the weight of news utility does not depend on when in the future said consumption occurs. For further details on both assumptions, see Dreyfuss et al. (forthcoming).

[5] In the DR variant, Nature is the first to take an action, whereas in the other three variants, the student takes the first action.

F_0 to $F_{l|z_1}$ and receives $N(F_{l|z_1} \mid F_0)$.

- **Period $\mathbf{1 < t < \bar{t}_z}$**: If called to act, the student takes the action prescribed to z_t by l. She updates from $F_{l|z_{t-1}}$ to $F_{l|z_t}$ and receives $N(F_{l|z_t} \mid F_{l|z_{t-1}})$.[6]

- **Period $\mathbf{\bar{t}_z}$**: The student learns about her final match and updates to the degenerate CDF $F_{l|z}$. She receives $m_{O(z)} + N\left(F_{l|z} \mid F_{l|z_{\bar{t}-1}}\right)$.[7]

In both static variants, the news utility stream reduces to $N(F_l \mid F_0) + \mathbb{E}_{\tilde{F}_l}\left[N\left(F_{l|z} \mid F_l\right)\right]$. The first term compares the belief over final matches induced by the submission of the ROL l in period 1 to an initial belief F_0 (consuming the outside option with certainty). The second term compares the result of the matching process with the belief over outcomes induced by the submitted ROL l (and recall that $F_{l|z}$ is degenerate). In the two dynamic variants, after the initial period-1 update, beliefs potentially update after each time Nature takes an action (or when the student changes her strategy, which does not happen in equilibrium). For example, in DP, the student updates after each rejection or acceptance, and in DR, she potentially updates after receiving each new set of offers.

1.4 Optimal Strategies

To derive predictions, we use Preferred Personal Equilibrium (PPE; Kőszegi and Rabin, 2009) as our solution concept for all variants.[8] First, given a strategy l and a non-terminal, t-periods-long history h, utility from deviating to some strategy $l' \in \mathcal{L}_h$ is given by

$$U_h(l' \mid l) = \mathbb{E}_{F_{l'|h}}[m] + N\left(F_{l'|h} \mid F_{l|h_{t-1}}\right) + \mathbb{E}_{\tilde{F}_{l'|h}}\left[\sum_{s=t+1}^{\bar{t}_z} N\left(F_{l'|z_s} \mid F_{l'|z_{s-1}}\right)\right], \tag{4}$$

i.e., expected consumption utility under the deviating strategy plus news utility from deviating from l to l' (at time period $t+1$) plus the expected sum of streams of news utility given the distribution of terminal histories under l'.

[6] Note, however, that in equilibrium, in all four variants new information is only learned after Nature's actions. Therefore, action taking and non-degenerate belief updating do not occur in the same period.

[7] In all four variants, Nature is always the last one to take an action: in the static variants, it responds to the submitted ROL with a final match; in DP, it accepts an application; and in DR it stops sending offers.

[8] More precisely, we use a slightly modified solution concept from Kőszegi and Rabin (2009)'s appendix, called *optimal consistent plan* (OCP). The main difference between the two solution concepts is that while in PPE the DM holds correct beliefs about all future contingencies from the moment of birth (i.e., before the first period), in OCP the DM has some initial prior beliefs when she forms her plan (which we fixed to F_0 above).

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

Definition 2. *Let z be a terminal history. We define the set of z-credible strategies in the following backward-recursive way: the credible set $\mathcal{L}^*_{z_{\bar{t}_z-1}}$ includes all strategies that maximize utility at the last action period given the expectation they induce, i.e., all strategies $l \in \mathcal{L}$ that satisfy $U_{z_{\bar{t}_z-1}}(l \mid l) \geq U_{z_{\bar{t}_z-1}}(l' \mid l)$ for all $l' \in \mathcal{L}$. Then the set $\mathcal{L}^*_{z_t}$ contains all strategies $l \in \cap_{h\in\mathcal{H}_{z_t}} \mathcal{L}^*_h$ that satisfy $U_{z_t}(l \mid l) \geq U_{z_t}(l' \mid l)$ for all $l' \in \cap_{h\in\mathcal{H}_{z_t}} \mathcal{L}^*_h$, where $\mathcal{H}_{z_t}$ is the set of histories that contain the subhistory z_t. 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.[9] The definition of a z-credible strategy allows us to define our solution concept :

Definition 3. *Let $\mathcal{L}^* \equiv \mathcal{L}^*_\emptyset$, i.e., the set of strategies that are credible in the empty history. A PPE strategy l^* is a strategy that satisfies*

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

In PPE, the core difference between static and dynamic variants 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.[10] 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.

[9]The 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.

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

1.5 Strategyproofness

As mentioned above, under our large markets assumption, all four variants are strategyproof for classical-preferences students and equivalently to loss-neutral students ($\lambda = 1$). The following definition extends strategyproofness to markets with agents who have EBRD preferences.

Definition 4. *A DA variant is* ***EBRD-strategyproof*** *if for any degree of loss aversion* λ*, and any belief* G*, the optimal (PPE) strategy is SF.*

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

Proposition 1. *The Dynamic student-Receiving (DR) variant is EBRD-strategyproof, while the other three variants (SP, SR, DP) are not EBRD-strategyproof.*

The proof is relegated to the online appendix. The second part of the proposition has been extensively discussed in Dreyfuss et al. (forthcoming) and Meisner and von Wangenheim (2019) and is easily proved by a counterexample. (We provide examples in section 2.) 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 only between schools that have sent her an offer, i.e., only between "sure things," and hence cannot do better than keeping the highest-value one. Applying this argument iteratively shows that in DR, the straightforward strategy is indeed the only one that is optimal in any decision node, in any possible history. This highlights the importance of dynamic implementation: In (the static) SR, 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 forced to choose a SF strategy.[11]

We note that in concurrent work, Meisner and von Wangenheim (2019) (proposition 4) show a similar, albeit different result about DR, using Rosato (2014)'s model of dynamic EBRD preferences. Assuming a unique stable matching for any realization of preferences,

[11] This 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 suspect EBRD-driven behavior in these settings may be a mistake. For further discussion, see DHR.

they show the existence of a Bayesian equilibrium in which all players play the SF strategy. A unique stable matching always exists in large markets; see Karpov (2019) for necessary and sufficient conditions in non-large markets. We assume large markets and show EBRD-strategyproofness, i.e., for any profile of other players' strategies, the player picks the SF strategy. We view the two results as complementary: assuming large markets implies (EBRD) strategyproofness, generalizing to all markets in which a unique stable matching is guaranteed implies existence.

1.6 Varying Competitiveness

Proposition 1 makes a stark prediction about observed behavior in DR vs. the other variants. We now formalize the model's second prediction, which relates observed behavior in the static variants to competitiveness.

In each non-DR variant, the threshold λ above which NSF behavior is optimal depends on the matching environment. In particular, since NSF behavior reflects an attempt to avoid disappointment from future rejections, this threshold λ increases as disappointment becomes less likely, all else equal. This makes an additional testable comparative static: NSF behavior increases with competitiveness.

Formally, denote the admission probability to school j by $\Pr(\rho_j > T_j) \equiv q_j$. Competitiveness is defined by the probability of acceptance at the highest-valued school, q_1. The next proposition is an extension of Meisner and von Wangenheim (2019)'s proposition 2. Before stating it, we impose the following assumption:

Assumption 1.

1. *$\rho_j \perp \rho_k$ for all j, k.*
2. *$q_j < q_k$ for all $j < k$.*

In words, we assume that (1) priority score is independent across schools, and (2) admission probability is decreasing with school value. While its part (1) in particular may not hold in important real-life deployments of DA, assumption 1 yields sharper theoretical predictions that our experiment—which satisfied it by design—can cleanly test. In particular, it generates a tighter upper bound on the threshold λ, and it generates an upper bound in the case of $q_1 \geq 0.5$.[12]

[12] In the parameters of our model, Meisner and von Wangenheim (2019)'s upper bound $\overline{\lambda}$ translates to $1 + \frac{1}{1-2q_1}$, and does not exist for $q_1 \geq 0.5$.

Proposition 2 (Based on Meisner and von Wangenheim 2019)**.**

1. *If* $\lambda < 1 + \frac{2}{1-q_1} \equiv \underline{\lambda}$*, the SF ROL is strictly optimal in SP and SR.*

2. *Suppose that assumption 1 holds. Then:*
   If $\lambda > 1 + \frac{2}{(1-q_1)^2} \equiv \overline{\lambda}$*, the SF ROL is strictly suboptimal in SP and SR.*

In the following sections, we test specific predictions implied by this proposition (see, e.g., figure 2 and the related explanation in the next section). We do so by experimentally varying q_1. Specifically, subjects in our experiment face matching environments (i.e., rounds) with three levels of competitiveness: high, medium, and low, with $q_1 = 0.05$, $q_1 = 0.2$–0.3, and $q_1 = 0.6$–0.65, respectively. In proposition 2 these q_1 values translate to the following non-overlapping bounds around λ: $\underline{\lambda} = 3.11$, $\overline{\lambda} = 3.22$ (high competitiveness); $\underline{\lambda} = 3.50$–$3.85$, $\overline{\lambda} = 4.13$–$5.08$ (medium); and $\underline{\lambda} = 6.00$–$6.71$, $\overline{\lambda} = 13.50$–$17.32$ (low competitiveness). As long as the population distribution of λ has a sufficiently large mass outside the bounded intervals (i.e., inside the ranges 3.22–3.5 and 5.08–6.00), the proposition implies that the prevalence of NSF behavior should increase with round competitiveness.[13]

2 Experimental Design and Predictions

2.1 The Experiment

The experiment consists of simulating a large matching market. Each subject is randomly assigned to one of 2×2 ({Static vs. Dynamic} $\times$ {Proposing vs. Receiving}) treatments denoted SP, SR, DP, and DR. This 2×2 design allows us to independently explore the difference both between proposing and receiving mechanisms and between static and dynamic implementations, by varying each of the two features while holding the other one fixed at both states. However, since the theory sets the EBRD-proof DR variant apart from the other three, it is assigned 40% of the subjects, while the other three are each assigned 20% of subjects.

Our goal was to design four treatments as similar to each other as reasonably possible in terms of instructions structure, length and language, and, more generally, all aspects of the

[13] Specifically, these analytical bounds imply that the increase in NSF share between high vs. medium round competitiveness is determined by the mass inside the range 3.22–3.5 and the increase in NSF share between medium vs. low competitiveness is determined by the mass inside the range 5.08–6.00. Notice that in the next section, we use specific matching problems and therefore get specific thresholds (rather than bounds).

user-interface look and feel. (Appendix C uses a four-color-coding scheme to indicate all cross-treatment differences.) In all treatments, following a tutorial and an attention check, each subject participates in ten order-randomized incentivized *matching problems*, each simulating a different large matching market. We then collect demographic variables and feedback. Each subject who successfully finishes the experiment receives a participation fee and the sum of matched-school values from every matching problem.

2.1.1 A Matching Problem

A matching problem consists of five schools with which the subject can be matched 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: The subject is assigned a *priority score* at every school, which is a random, uniformly and independently distributed integer between 0 and 99. 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 each problem, the subject learns each school's threshold but not her own priority scores at the different schools. This creates an *unconditional* probability of acceptance at each school ($= 1 - \frac{\text{threshold}}{100}$). Figure 1 reproduces the example matching problem given in the tutorial, the way it appears on subjects' screen.

Figure 1: Screenshot of an example matching problem

School	Threshold	Value	Your Priority Score	Chance that Your Priority Score ≥ 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%

Notes: Example screenshot of the table describing a matching problem. The table's structure is identical in all treatments and rounds. School names and parameters differ by round. See table 2 for the parameters used in the ten paying rounds (the shown screenshot is taken from the tutorial round).

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

2.1.2 The Matching Process

The matching process varies by treatment, closely resembling the structure outlined in the theory section. In both *static* treatments (SP and SR), subjects submit a rank-order list (ROL) of schools in advance and are matched with the highest-ranked school in which their priority score exceeds its threshold. In DP, subjects sequentially apply to schools, and get accepted to the first school in which their priority score exceeds its threshold. In DR, subjects sequentially receive offers from schools in which they exceed the threshold, with scores higher above the threshold resulting in earlier-arriving offers, and can keep at most one offer at any point. Appendix C provides screenshots for all treatments.[14]

2.2 Theoretical Predictions

2.2.1 Matching Problems

As explained in the theory section, a subject with a coefficient of loss aversion λ is faced with a matching problem $\langle M, G, \mathbf{T}, \mathbf{m} \rangle$, where M is the matching process, G is the joint distribution over priority scores (recall that priority scores in the experiment are uniform and independent), $\mathbf{T}$ is a vector of thresholds and $\mathbf{m}$ is a vector of school consumption utilities. We assume a linear consumption utility, i.e., if s_j is worth v dollars, then $m_j = v$. Last, schools are denoted by their indices (i.e., s_1 is denoted by 1), and ROLs are represented by a sequences of numbers denoting the order in the list (i.e., the ROL ranking five schools in descending order of value is 12345).

When designing our experiment, our goal was to create variation in NSF predictions not only across treatments—DR vs. the other three variants, testing proposition 1—but also across problems—more vs. less competitive, testing proposition 2. To create variation in competitiveness, we searched for settings we could classify into three levels of predicted NSF: high, medium, and low. We briefly describe our search process; for further details, see the appendix.

We calculated optimal-strategy predictions for a range of λs, for each of the non-DR

[14] The stream of offers a subject i receives in DR contains at most five periods and is determined as follows. For each school s_j in which the subject exceeds the threshold, the segment [threshold, 99] is divided into five quintiles, $Q_1, \ldots, Q_5$ (where Q_1 is the bottom quintile). The subject then receives an offer from s_j in period $6 - Q_{ij}$, where Q_{ij} is the quintile in which her priority score lies. This induces a stream of offers from different sets of schools at different periods, where periods with no offers are eliminated.

variants, in each matching problem, from a large pool of candidate problems. Figure 2 illustrates, for two example matching problems (#1 and #7), the prediction over the range $\lambda \in [2.5, 5.5]$. For each possible strategy—which, for these three variants, can be represented as a ROL—we calculated the resulting expected overall (consumption + news) utility at a range of λs. Each subfigure shows all the ROLs that maximize utility for some λ in the range for a given problem and variant. For the dynamic variant DP, we also verified that any point on the envelope represents a consistent strategy, i.e., it is immune to profitable surprise deviations, and is thus a PPE. Therefore, the envelope in each subfigure represents the model's optimal-strategy predictions by λ.

As the figure demonstrates, an optimal strategy in the non-DR variants depends on λ. For low levels of loss aversion, the SF strategy (i.e., the ROL 12345) is optimal. However, a threshold λ exists above which NSF strategies are optimal. Because the gradual resolution of uncertainty in DP yields a different expected news utility, said threshold differs across the variants—compare subfigures (a) and (c) with subfigures (b) and (d). Moreover, because the environment in problem #7 is less competitive than in problem #1, the threshold is higher in problem #7, implying less prevalent NSF behavior—compare subfigures (a) and (b) with subfigures (c) and (d). Notably, the envelopes consist of *specific* NSF strategies that the model predicts as optimal under the different problems and variants, yielding an additional prediction that we assess empirically below.

Having created, for each candidate matching problem, optimal-strategy predictions as a function of an individual's loss aversion λ, we generated population predictions. We wanted to create ex-ante, no-degrees-of-freedom predictions but, naturally, did not have a previously estimated distribution of λ in our subject populations. We opted for basing our population predictions on past estimates from a somewhat similar population: estimates from (Dreyfuss et al., forthcoming) among lab-experiment participants in a related context (Li, 2017). In those estimates, 67 percent have $1 \leq \lambda \leq 3$, and are therefore never predicted to deviate from SF behavior; 24 percent have $3 < \lambda \leq 5$; 7 percent have $5 < \lambda \leq 7$; and 2 percent have $7 < \lambda \leq 10$; for details, see the online appendix. We note that any distribution with a tail of $\lambda > 3$ will result in qualitatively similar directional predictions.

Table 2 presents the ten matching problems we selected for our experiment. Each problem is defined as five school values and (unconditional) probabilities of acceptance. The three columns under "NSF (%)" report the population prediction for the prevalence of NSF strategies under each of the four treatments (recall that our large market property implies that the two static treatments are equivalent, so their predictions coincide). For

Figure 2: Candidate optimal strategies in example matching problems

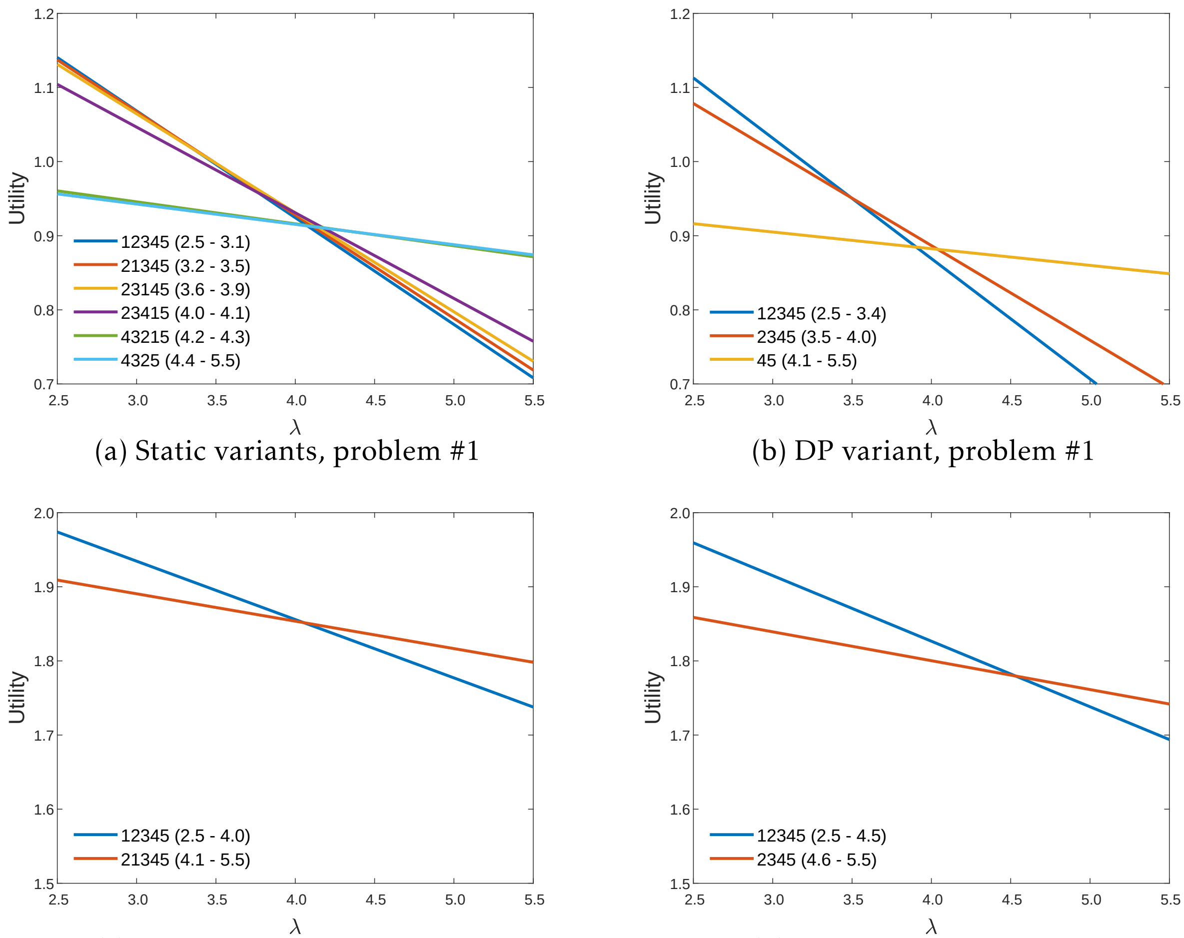

(a) Static variants, problem #1

(b) DP variant, problem #1

(c) Static variants, problem #7

(d) DP variant, problem #7

Notes: Candidate optimal strategies for matching problems #1 and #7. Schools are denoted by their indices (i.e., s_1 is denoted by 1), and ROLs are represented by a sequences of numbers denoting the order in the list. Matching problem #1 is defined by the vectors of dollar school values (1.5, 1, 0.75, 0.5 0.25) and thresholds (0.95,0.8,0.75,0.1,0), and problem #7 is defined by school values (1.25, 1, 0.75, 0.5 0.25) and thresholds (0.7,0.15,0.1,0.05,0). (See table 2 for school values and thresholds for all problems #1–#10.) Subfigures (a) and (c) plot, for the static variants, utilities from all strategies that are optimal for some loss-aversion parameter in the range $2.5 \leq \lambda \leq 5.5$ (grid step $= 0.1$). Panels (b) and (d) do the same for DP.

example, the 31% prediction in problem #1 in the SP/SR column reflects the estimated population share with $\lambda > 3.1$, which in that problem implies that the ROL 12345 is no longer an optimal strategy (see figure 2(a)). We chose these ten matching problems, which each subject in our experiment faces, based on the population predictions in these "NSF (%)" columns. (The four rightmost columns report population predictions using alternative measures; we discuss them below.)

Table 2: Matching problems and EBRD predictions

#	Matching problem: \$ Values $(s_1, s_2, s_3, s_4, s_5)$ / Probs. $(q_1, q_2, q_3, q_4, q_5)$	Predictions: NSF (%) DR	NSF (%) SP/SR	NSF (%) DP	Costly NSF (%) SP/SR	Costly NSF (%) DP	DO NSF (%) SP	DO NSF (%) SR
1	(1.50, 1.00, 0.75, 0.50, 0.25) (0.05, 0.20, 0.25, 0.90, 1.00)	0%	31%	26%	7%	7%	31%	7%
2	(1.50, 1.00, 0.75, 0.50, 0.25) (0.05, 0.20, 0.30, 0.40, 1.00)	0%	31%	24%	7%	9%	31%	7%
3	(1.50, 1.00, 0.75, 0.50, 0.25) (0.05, 0.10, 0.85, 0.9, 1.00)	0%	29%	29%	4%	3%	29%	4%
4	(1.25, 1.00, 0.75, 0.50, 0.25) (0.20, 0.80, 0.90, 0.95, 1.00)	0%	24%	18%	4%	3%	24%	4%
5	(1.50, 1.00, 0.75, 0.50, 0.25) (0.25, 0.80, 0.85, 0.95, 1.00)	0%	21%	18%	4%	4%	21%	4%
6	(1.25, 1.00, 0.75, 0.50, 0.25) (0.25, 0.75, 0.80, 0.85, 1.00)	0%	19%	12%	4%	2%	19%	4%
7	(1.25, 1.00, 0.75, 0.50, 0.25) (0.30, 0.85, 0.90, 0.95, 1.00)	0%	18%	14%	5%	3%	18%	5%
8	(1.25, 1.00, 0.75, 0.50, 0.25) (0.65, 0.70, 0.85, 0.95, 1.00)	0%	0%	0%	0%	0%	0%	0%
9	(1.50, 1.25, 1.00, 0.50, 0.25) (0.65, 0.75, 0.80, 0.85, 1.00)	0%	0%	0%	0%	0%	0%	0%
10	(1.25, 1.00, 0.75, 0.50, 0.25) (0.60, 0.65, 0.80, 0.90, 1.00)	0%	0%	0%	0%	0%	0%	0%

Notes: Parameters and population predictions for each of the ten matching problems. Matching problem columns describe schools' money values (top row) and unconditional probabilities of acceptance (bottom row) in each problem. Prediction columns are based on an empirically estimated distribution of λ from Dreyfuss et al. (forthcoming).

2.2.2 Within-treatment Predictions

The most important feature for predicting NSF shares in a matching problem is the probability of acceptance at the highest-value school (see proposition 2). Specifically, this probability determines the predicted behavior of a subject with a moderately high loss-aversion parameter (say, $3 < \lambda < 6$).

We use the label "weak-student" problems for matching problems 1–3, 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 use the label "medium-student" problems for problems 4–7, where the probability of acceptance at the highest-value school is larger (20%–30%), implying a lower predicted share of NSF strategies (12%–24%). Finally, we use the label "strong-student" problems for problems 8–10, 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 we use). The online appendix further describes the process and the criteria we used for choosing the sets of values and probabilities in the ten problems.

In summary, we predict that under the three non-DR treatments, the share of NSF strategies will monotonically decrease from weak- to medium- to strong-student problems. We also predict *how* agents will depart from SF strategies. For example, as Figure 2 shows, the second and third most prevalent NSF ROLs in the static treatments are predicted to be 21345 and 23145, respectively.

2.2.3 Between-treatments Predictions

As reported in the three columns under "NSF (%)" in table 2, we predict varying shares of NSF strategies under the three non-DR treatments, and none under DR. It is important to note that these are predictions on chosen *strategies*; however, in our experiment we only collect data on observed *behavior*. Under the static treatments (SP, SR) there is no difference between strategies and behavior, as we fully observe subjects' submitted ROLs. However, under the dynamic treatments (DP, DR), we only observe actions dynamically taken by subjects, which only reveal partial information on full strategies.

Specifically, under DP, we only observe the set and order of schools a 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 sequence of offer sets a student received and the schools the student decided to keep from those sets. Such observed decisions can be either *SF-consistent* or *SF-inconsistent*.

For example, suppose a subject's priority score at the highest-value school is higher than its threshold (so once the subject applies to that school, the process terminates). In that case, a SF-consistent behavior under DP consists of the subject only applying to that school. However, such observed behavior could also result from a NSF strategy such as 14235. On the other hand, suppose a subject's priority score at the highest-value school is *lower* than the threshold. Under DR, the school never sends an offer to that subject. In that case—by far the most common case in weak-student matching problems (see table 2)—we will never observe the subject's choice regarding that highest-value school in DR. More generally, a subject's behavior may be SF-consistent merely because we had limited opportunity to observe deviations. As a result, SF-inconsistent behavior in dynamic treatments is only a lower bound on NSF strategies.

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 discussed below, the theory predicts only relatively small cross-treatment differences in these measures—differences that our experiment is not optimized to detect.

The first measure is *costly* NSF behavior, which measures NSF behavior that is also payoff relevant (i.e., SF-inconsistent behavior that affected the subject's final match). This measure allows for direct comparisons between all four treatments. Clearly, under the DR treatment, we predict no such behavior; however, as the two columns under "Costly NSF (%)" in table 2 show, the predicted share for this measure varies little across both matching problems and treatments (and always remains below 10%).[15]

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

As suggested by the two columns under "DO NSF (%)" in table 2, the theory predicts that under SP, every NSF strategy is also dynamically observable. In contrast, the theory predicts that under SR, all dynamically observable NSF strategies are costly. These two predicted identities limit the usability of this measure as an alternative for cross-treatment comparisons; however, they provide additional within-treatment predictions that we investigate in the next section.[16]

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

[16] The driver behind these predicted identities is the fact that the NSF strategies predicted by the model in the non-DR treatments always involve the highest-value school. See Meisner and von Wangenheim (2019)'s characterization.

3 Results

3.1 Sample

We ran the experiment on two different samples:[17] participants in Cornell's BSL (Business Simulation Lab) SONA-system, recruited from June 30 to July 9, 2021, and participants in the 2020 and 2021 Israeli Psychology Matching Mechanism (IPMM)—a DA-based clearinghouse that matches students and graduate programs in psychology—recruited from August 26 to September 5, 2021. We used identical experimental interfaces, except for language (English vs. Hebrew) and currency (we used 1 USD = 4 NIS in the experiment and $4 vs. 15 NIS as a show-up fee).

At Cornell, 223 subjects clicked on the experiment's link, and we closed the experiment after 206 subjects completed it. They earned on average $13.23 (including the show-up fee); the median completion time was 16.8 minutes. As preregistered, we dropped eight subjects who failed the attention check and another two who completed the experiment itself (after the tutorial) in more than an hour. The remaining 196 subjects' assignment is: 77 (39.3 percent) DR, 39 (19.9) SR, 40 (20.4) SP and 40 (20.4) DP.

In Israel, we first emailed 1,095 invites to the 2021 IPMM participant pool, of whom 225 clicked on the experiment and 171 completed it. To meet our preregistered 200-subjects target, we emailed 1,206 additional invites to the 2020 pool, of whom 215 clicked on the experiment and 138 completed it. Together, they earned an average of $16.43 (including the show-up fee) in a median completion time of 22.2 minutes. After dropping five subjects who failed the attention check, the remaining 304 subjects' assignment is: 126 (41.4 percent) DR, 60 (19.7) SR, 59 (19.4) SP, and 59 (19.4) DP.

In the rest of this section we pool the samples ($N = 500$; 203 DR and 99 each SP, SR, and DP). Online appendix D replicates the main analysis by sample. While the general level of observed NSF behavior is markedly lower in the IPMM sample, our main findings remain similar across the samples.

3.2 NSF Shares

Panel (a) of figure 3 presents the EBRD model's predictions regarding NSF shares (based on said past estimated population distribution of loss aversion; ▲) compared to classical-preferences predictions (▼), by treatment and problem type. Within treatments, in the

[17] Preregistered separately at https://aspredicted.org/rd2y5.pdf and https://aspredicted.org/4qc8q.pdf.

three non-DR treatments, the EBRD-predicted NSF-share declines as competitiveness decreases. Across treatments, in all but strong-student problem types, the EBRD-predicted lower (0 percent) NSF share in DR compared to the other three treatments. Trivially, classical preferences predict 0 percent NSF in all treatments and problems.

Panel (b) presents empirical shares of NSF behavior (■), as well as p-values from equality-of-coefficients tests.

Figure 3: Non-straightforward (NSF) behavior

(a) Theoretical Predictions: Classical & EBRD Preferences

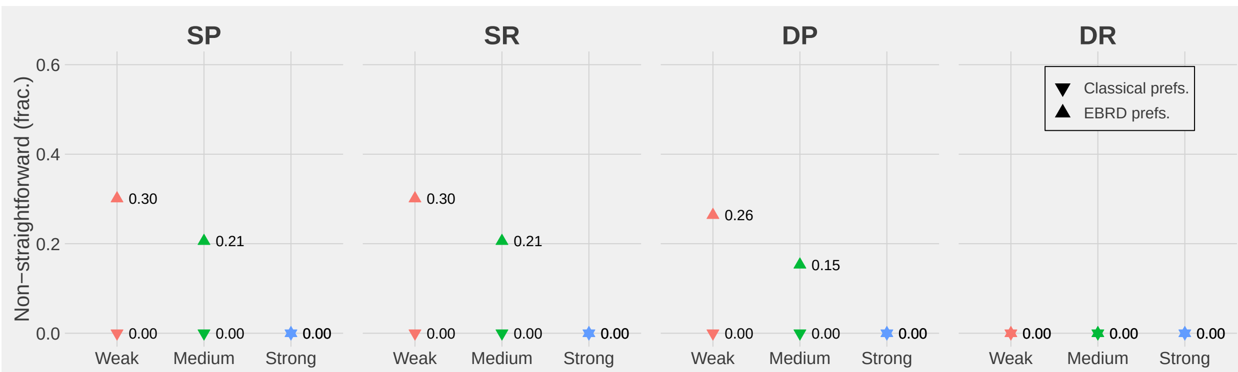

(b) Results

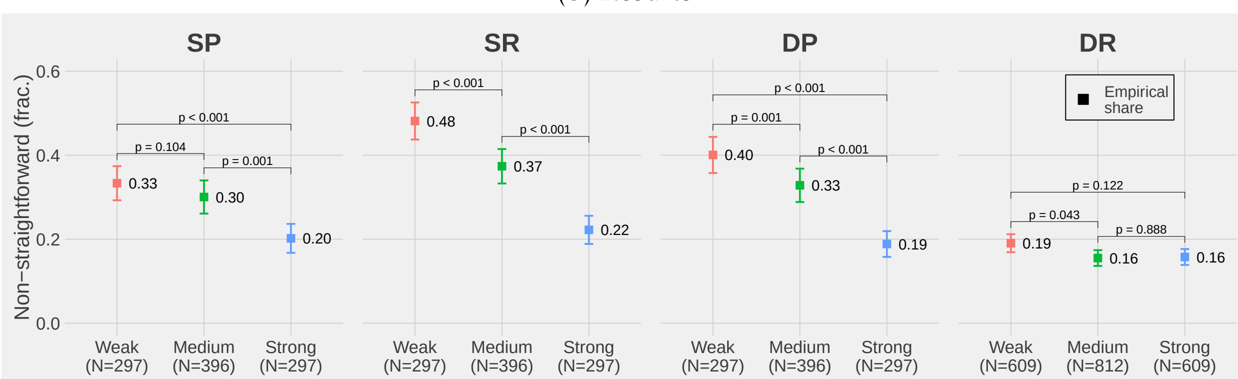

Notes: Panel (a): NSF-share predictions by treatment and problem type under classical preferences (downwards triangle) and EBRD preferences (upwards triangle). Panel (b): empirical shares. Error bars: standard errors from a regression of NSF behavior on problem type, clustering at the individual level. p-values: Wald tests.

We make three observations. First, looking at *levels*, there appears to be an across-the-board minimum "baseline" of 16–22 percent NSF behavior, in all treatments, that neither classical nor EBRD preferences can explain.

Second, looking at within-treatment *trends*, the data strongly support the EBRD prediction of a notable decrease in the share 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 percent. (The flat, zero-NSF predictions of classical preferences are easily rejected.[18]) In contrast, the trend in DR is much flatter (from 19 to 16 to 16 percent NSF), consistent with our EBRD-model predictions (which, for DR, coincide with those of classical preferences).

Third, comparing across treatments, the EBRD predictions are also supported by the data. In contrast with the constant no-behavior-difference prediction of classical preferences, NSF shares in DR are substantially lower than in other treatments in weak- and medium-student problems, where EBRD predicts a difference; but are essentially indistinguishable in strong-student problems, where the model predicts no difference.

Online appendix D reproduces panel (b) of figure 3 twice: for subjects' first vs. last five rounds. While the above three observations generally hold in each of the two subsamples, NSF shares noticeably drop from the first to last five rounds—suggesting that experience-based learning during the experiment may make our results still more consistent with the EBRD model's predictions. In particular, regarding *levels*, the minimum baseline that neither model can explain decreases in the last five rounds to 10–17 percent; and, regarding *trends*, the declines in the non-DR treatments from weak- to medium- to strong-student problems become 30–37 to 23–36 to 12–16.

3.3 ROL Types

Moving beyond mere NSF shares, the model (with our distribution of λ) predicts a specific distribution of NSF ROLs in both static treatments.[19] Figure 4 compares the predicted distribution (horizontal axis) vs. the observed distribution (vertical axis), pooling together all ten matching problems in the two static treatments.

The figure shows that NSF ROLs that are predicted to be prevalent—dots closer to the right half of the figure—are indeed roughly as empirically common—close to the 45° line. In particular, the two NSF ROLs with the first and second highest predicted shares—21345 and 23145, respectively—are also the first and second most empirically prevalent NSF ROLs. While some predicted ROLs are never observed in the data and vice versa, the differences between prediction and empirical prevalence are never much higher than one percentage point.

Finally, zooming further in, we examine specific ROL-type predictions in specific

[18] The classical-vs.-EBRD comparison is "fair" in that both models have zero degrees of freedom. (While the EBRD model has a free parameter (λ), our predictions have no free parameters, as they were generated, prior to data collection, based on a previously estimated population distribution of λ.)

[19] The same is true for DP; however, in DP we do not observe complete strategies but rather incomplete sequences of applications.

Figure 4: Predicted vs. observed frequency of ROLs

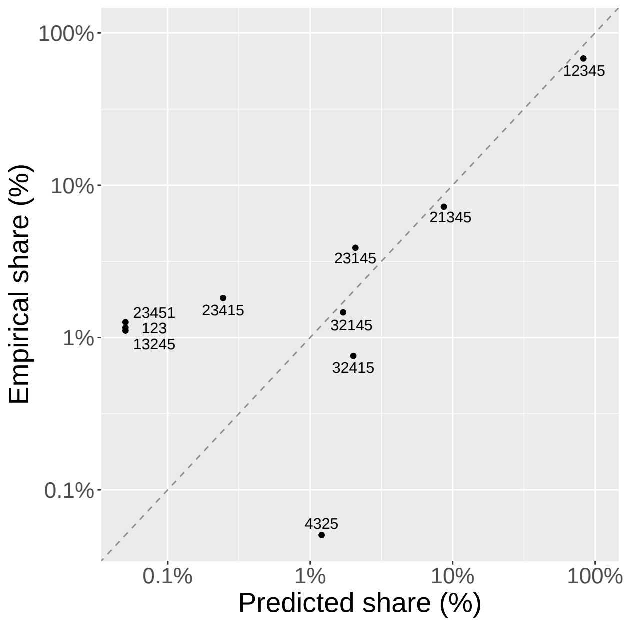

Notes: Theoretical EBRD predictions and empirical shares of ROLs in the static treatments (logarithmic scale; N =1,980). ROLs shown are all those with at least 1% predicted share or empirical share.

problems. We focus on weak problems: problems #1–3 in table 2. In addition to having a higher predicted NSF share, we intentionally designed these three problems to show variation in the top-ranked school among loss-averse subjects. Specifically, in problem #3, $q_2 = 0.1$ is relatively low, while $q_3 = 0.85$ is relatively high, yielding a prediction that the most prevalent NSF ROL will rank the third-highest-value school on top. In contrast, in problems #1 and #2, $q_2 = 0.2$ is higher, and $q_3 = 0.25$–0.3 is only slightly above q_2, yielding a prediction that the most prevalent ROL will rank the second-highest-value school on top. Since problems #1–3 are otherwise similar to each other, we view the test of these predictions as a particularly sharp test of the theory.

Figure 5 has similar structure to figure 3, but it shows predictions and results by *problem*, only for problems #1–3 and only for NSF behavior with s_2 on top in panels (a) and (b), and with s_3 on top in panels (c) and (d).

Comparing panel (a) vs. (b) and panel (c) vs. (d) suggests that the data qualitatively track most patterns predicted by the EBRD model. In contrast, the the data mostly reject non-trend predicted by classical preferences.

3.4 NSF Behavior in DR

Our main between-treatments prediction compares DR to the other three treatments. However, as discussed in section 2, such comparisons may favor DR because our main outcome, NSF shares, only counts *actions* taken by subjects, not complete strategies. The general worry is that unlike in the static treatments, in DR, some NSF strategies would not lead to observed NSF behavior unless the subject received sufficiently many offers. In particular, if NSF behavior in DR is driven by strategies consistent with ROLs prevalent in the static treatments—e.g., ROLs like 21345—then observations in which the subject does not receive an offer from the highest-value school, which are quite common given our design, would not be classified as SF-inconsistent even if the subject chose a non-straightforward strategy. In short, subjects may have fewer opportunities in DR to display SF-inconsistent behavior—especially if they play NSF strategies consistent with common NSF ROLs observed in the static treatments.

In this section, we present two types of evidence suggesting this is not likely the case. First, we show that the most prevalent NSF behavior in DR does not reflect NSF strategies that often fail to appear as SF-inconsistent. Second, we show that observed NSF does not increase with the number of offers or with the presence of high-value offers.

First, across all ten problems, 338 out of 2,030 subject-problem observations (17 percent) in DR are classified as NSF. Of these 338 NSF observations, 249 (74 percent) involve rejecting offers received in the first period. Of these first-period offer rejection sets, 215 (86 percent) contain only the lowest-, or second-to-lowest-value school, or both. In other words, most observed NSF behavior in DR involves getting offers only from low-value schools in the first period, and rejecting them. This behavior appears unique to DR and implies that NSF behavior in our DR data typically occurs in situations that subjects often face.

Second, we do not find that observed NSF behavior increases with the number of offers received. Table 3 classifies all 2,030 DR observations and 338 DR NSFs by the number of offers received and, if anything, shows the opposite trend.

Moreover, if subjects followed a strategy consistent with a ROL such as 21345, we would expect to see a higher fraction of observed NSF behavior conditional on receiving an offer from the highest-value school compared to not receiving it. However, we do not find that in the data: There are 621 observations in which subjects received an offer from the highest-value school and 1,409 in which they did not; of those, respectively, 98 (16 percent) and 240 (17 percent) are classified as NSF.

Table 3: NSF by number of received offers in DR

# offers	# obs.	# NSF behavior obs.	% NSF behavior obs.
1	95	25	26%
2	294	51	17%
3	592	94	16%
4	743	122	16%
5	306	46	15%
1–5	2,030	338	17%

Notes: Distribution of DR observations by the number of offers the subject received ($N = 2{,}030$).

3.5 Alternative Measures

As discussed in section 2, in addition to the share of all NSF behavior, we had two additional NSF measures that allow for more direct comparisons across treatments. The first, costly NSF, counts only payoff-relevant NSF behavior, and allows for direct comparisons across all four treatments. The second, dynamically observable NSF (DO NSF), counts only NSF behavior that would have been observable in a dynamic implementation, and allows for direct comparisons between SP and DP, and between SR and DR.

Panels (a) and (c) in figure 6 show our predictions (see also table 2 on page 20 and its accompanying discussion). As panel (a) shows, we predict a small share of costly NSF, with little variation across and within treatments (0–6 percent). Panel (c) shows that for DO NSF, in SP we predict significant variation—the same variation predicted for NSF—across problem types, and in SR we predict small variation—the same variation predicted for costly NSF—across problem types.

Empirically, panel (b) shows a slightly higher-than-predicted level of costly NSF behavior (6–14 percent), with no systematic differences across treatments and problem types. Panel (d) shows that for SP, when moving from weak- to medium- to strong-student problem types, DO NSF drops from 31 to 27 to 15 percent, consistent with the EBRD-predicted trend. In SR, the level of DO NSF is higher than predicted (10–16 percent), with no clear trend across problem types.

While we did not optimize our experiment to detect differences in these alternative measures, in principle we did collect enough observations to marginally detect a 4–6 percent predicted difference in them, both within each of the non-DR treatments and between DR and non-DR treatments, under ideal conditions—i.e., if the data were perfectly described by the model. However, as discussed above, we find higher NSF shares than our model predicts. These higher NSF shares muddy the picture for both alternative measures.

Start with costly NSF. We find that NSF behavior where the model does not predict it—in DR and in strong-student rounds—is costly 36–63 percent of the time, whereas in weak- and medium-student rounds in the non-DR treatments it is costly only 18–42 percent of the time (and always less, within each non-DR treatment, than in strong-student rounds). Therefore, the empirical *NSF-to-costly-NSF ratio* differs both between DR and non-DR treatments, and between problem types within the non-DR treatments, reflecting behavior that our model does not predict.

Moving to DO NSF, in SR we see a similarly high NSF-to-DO-NSF ratio in strong- relative to medium- and weak-student problems, masking the (small) predicted difference in DO NSF. In contrast, in SP, this ratio is similarly high across problem types. Both of these findings simply reflect the empirical distribution of NSF ROLs, discussed in section 3.3: since NSF ROLs typically involve downranking the highest-value school, they are typically dynamically observable under DP, but much less so under DR.

To summarize, we examine two alternative measures, both allowing direct comparisons between DR and other treatments. However, for these measures, our model predicts small variation, making it difficult to detect such differences in our (inherently noisy) data. Digging further in, we find that this variation is masked by the presence and the type of NSF behavior in the treatment (DR) and problem type (strong) in which the model predicts no NSF.

3.6 Alternative Explanations

Throughout this section, we show that the data fit our (no-degrees-of-freedom) EBRD model significantly better than they do classical preferences. Nevertheless, how well can alternative models explain the data?

Recall that we review three sets of empirical patterns, all consistent with our EBRD model: (i) variation in NSF shares within DA variants; (ii) variation in shares across variants;[20] and (iii) prevalence of specific NSF types, both pooling across matching problems, and for specific problems. At the same time, as we note at the beginning of this section, we also find a little-varying "non-pattern": an across-the-board "baseline" NSF share (of roughly 16–22 percent in figure 3) that neither classical nor EBRD preferences can explain.

In summary, the EBRD model, while explaining a lot of the observed data, appears to be an incomplete explanation. Alternative explanations—noisy decision making, cognitive

[20] We continue to assume that a mechanical bias does not drive the lower observed NSF share in DR. See the discussion in section 3.4.

limitations, strategic confusion, misunderstanding, or other explanations—likely play a role too. However, no alternative model we are aware of can replace EBRD in explaining the three patterns above.

Trembling-hand models, with an exogenous probability of making errors, are inconsistent with the observed within-treatment variation (as long as the trembles are independent of the matching problem—as is typically assumed). They are also inconsistent with the observed cross-treatment variation: In DR, subjects typically make multiple Keep/Reject decisions, and are therefore predicted by trembling-hand models to make *more* errors than in the (single-ROL-decision) static treatments. Finally, they cannot easily explain the prevalence of specific NSF strategies (unless one imposes a specific distribution over trembles), and in particular of specific NSF strategies in specific matching problems.

Random utility models (RUMs)—which in our context boil down to decision-makers randomly making errors, with the probability of making an error decreasing with its cost—can potentially explain the observed within-variant patterns, since increasing competitiveness decreases the cost of, e.g., submitting 21345 instead of 12345. However, they too are inconsistent with the observed cross-treatment variation, for the same reason above. Importantly, these models are also inconsistent with the prevalent ROL types in our data (see figure 4): Since RUMs predict that low-cost errors are the most likely, in our experiment, RUMs would typically predict low-cost NSFs, such as 12354 (whose expected value is only slightly below that of the SF strategy 12345) to be quite prevalent, but in reality these are rarely seen in the data. In fact, when choosing parameters for our experiment, one of our goals was to tell apart EBRD from RUMs; we chose parameters such that EBRD and logit (a specific RUM) make different predictions.

Recent models of cognitive limitations, strategic confusion, and misunderstanding—such as Li (2017), Pycia and Troyan (2021), Börgers and Li (2019) and Gonczarowski et al. (2022)—classify mechanisms, their variants, or their implementations (e.g., how the variant is described to subjects) by different notions of how simple they are. Therefore, they do not generally explain variation in behavior *within* a specific implementation of a specific variant of a specific mechanism. Moreover, while some such models suggest that dynamic implementation increases straightforward behavior, none of the models we are aware of tells apart dynamic-proposing and dynamic-receiving DA—implying that these models cannot easily explain the observed variation across treatments either. Finally, these models are also silent on the specific types of errors people make in specific choice situations (within a specific mechanism) and thus cannot easily explain the prevalence of specific NSF types in general or within specific matching problems.

4 Discussion

Our findings have several implications. First, we replicate previous findings that a substantial fraction of participants in simple allocation games choose seemingly dominated actions. Such findings are particularly striking in controlled lab settings, where said seemingly dominated actions are equivalent to choosing first-order stochastically dominated (FOSD) lotteries over sums of money. While in clear violation of the predictions of Classical-preferences models, our analysis suggests that the inclusion of EBRD preferences can explain a substantial fraction of such behavior. Moreover, the EBRD model can explain specific *types* of such FOSD-violating behavior, both predicted by the model and observed in the data. Importantly, these predictions are made while adding no additional degrees of freedom relative to the no-EBRD model; instead, we merely replace the default assumption of no loss-aversion—i.e., $\lambda = 1$ for everybody—with an empirical population distribution of λ estimated in a previous study.

At the same time, while our DR treatment is predicted to be EBRD-proof, a non-negligible fraction of participants still choose dominated actions—actions that remain dominated at *any* degree of loss aversion. We also find similar fractions of participants choosing such actions in "strong student" problems under our other DA variants, where the (parameter-constrained) EBRD model *also* predicts no such behavior. These findings underscore the need for additional explanations, for example, along the lines of Li (2017) and Gonczarowski et al. (2022).

Another important prediction of the EBRD model that our experiment confirms is that while the fraction of such non-straightforward (NSF) behavior varies significantly across both DA variants and matching problems, the fraction of *costly* NSF behavior does not vary much across either. From a theoretical point of view, the intuition behind this prediction is simple: for a fixed distribution of loss aversion λ, the EBRD model predicts more FOSD violations the lower the probabilities of acceptance at higher-value schools. However, the lower those probabilities, the lower the chance that such violations will end up costly.

From a policymaker's point of view, one potential reaction is that there is little value in reducing the prevalence of NSF if it is mainly bottom-line-outcome irrelevant. However, such a reaction may overlook two nuances. First, in our data, NSF behavior does not only *decrease* under DR; it also *changes*, and could have distributive consequences. The (reduced)

NSF behavior that we find under DR—and that the EBRD model cannot explain—may be consistent with, e.g., (potentially over-optimistic) subjects rejecting low-value offers and expecting to get better ones. In contrast, EBRD-consistent NSF in the non-DR variants is consistent with (potentially over-pessimistic) loss-averse subjects shying away from applying to high-value schools.

Second, informal conversations with policymakers suggest that at least some of them strongly view prevalent NSF behavior as a problem to be minimized, regardless of the actual cost. For example, "leveling the playing field" and other equity arguments favoring DA become more complicated when NSF behavior is prevalent. From that perspective, reducing the actual incidence of NSF behavior is a goal in itself.

Finally, for policymakers, the question of implementation feasibility looms large—a question that we do not address in this paper. By definition, dynamic implementations of a matching mechanism require it to run for more extended periods, during which participants are repeatedly asked to make decisions. For investigations like ours to have real-world impact, progress will have to be made on whether and how to implement such variants.

Figure 5: NSF by ROL Type

(a) Theoretical Predictions (2 on Top): Classical & EBRD Preferences

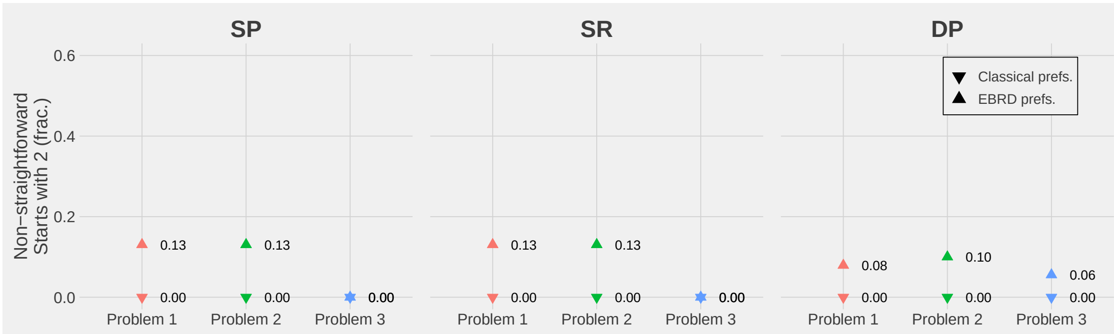

(b) Results (2 on Top)

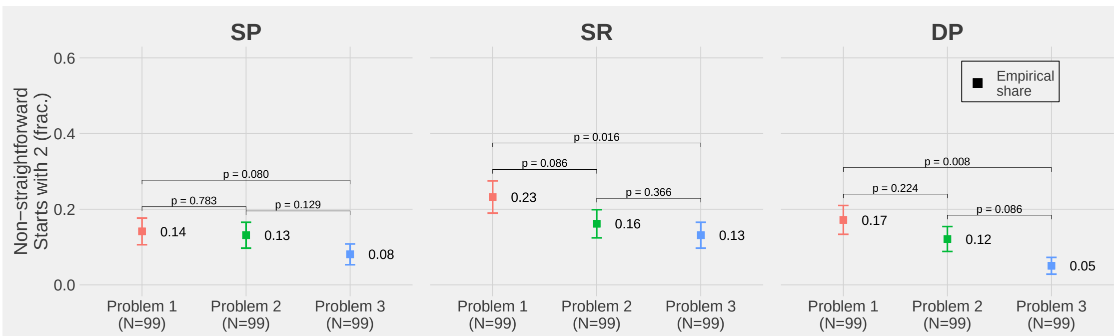

(c) Theoretical Predictions (3 on Top): Classical & EBRD Preferences

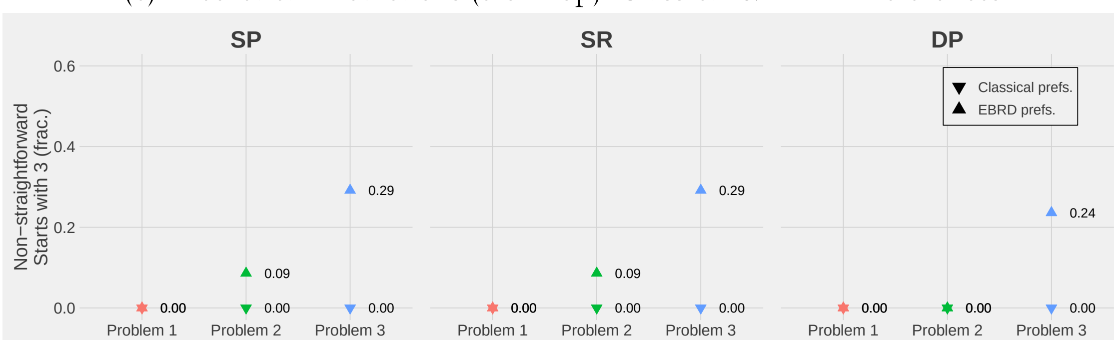

(d) Results (3 on Top)

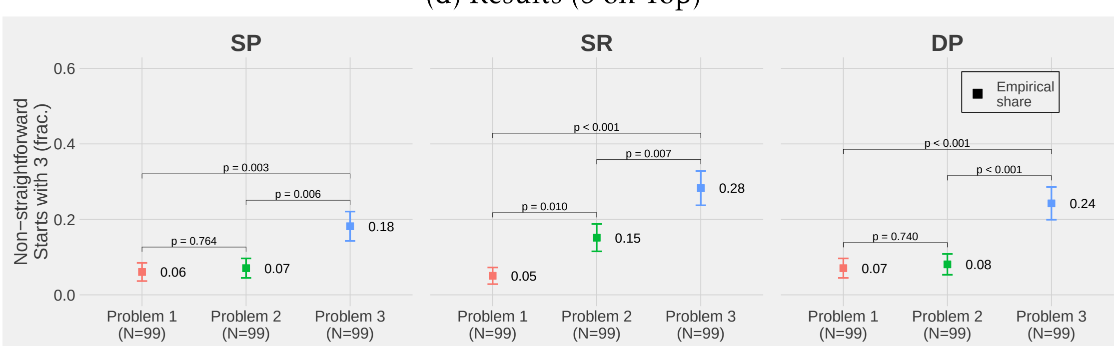

Notes: Fraction of NSF ROLs where the 2nd (3rd) highest-valued school is the one ranked first (*weak* problems only). Panels (a) and (c) present chosen NSF strategy predictions by problem type under Classical (downward triangles) and EBRD (upward triangles) preferences. Panels (b) and (d) present the empirical shares. Error bars: standard errors from a regression of NSF behavior on problem type, clustering at the individual level. *p*-values: Wald tests.

Figure 6: Alternative Measures

(a) Theoretical Predictions (Costly NSF): Classical & EBRD Preferences

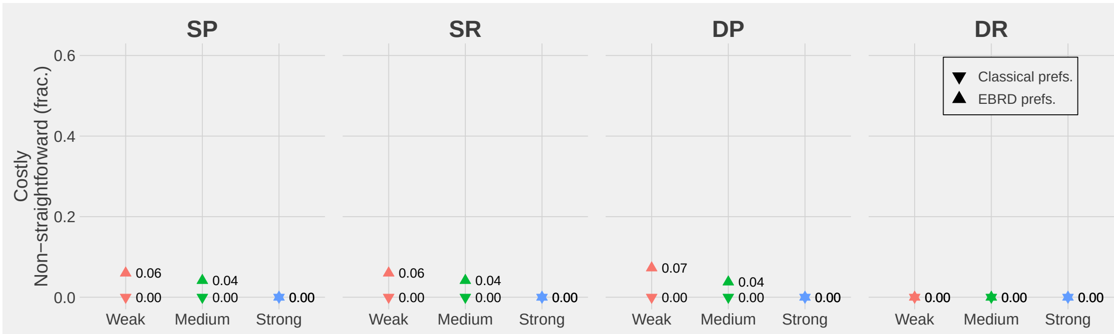

(b) Results (Costly NSF)

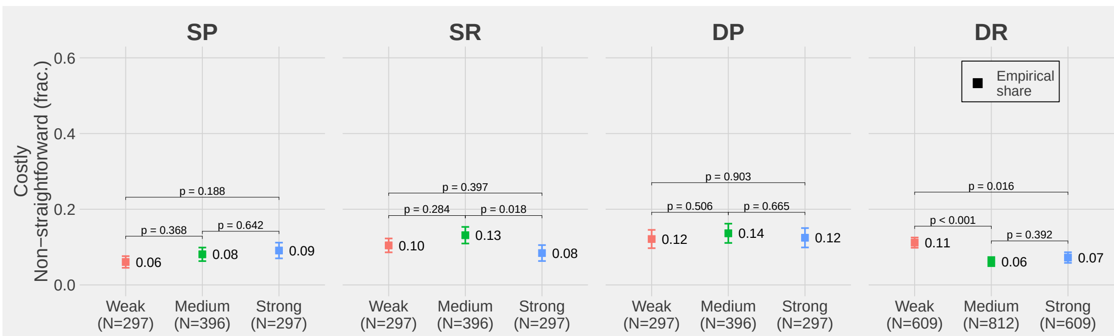

(c) Theoretical Predictions (Dynamically Observable NSF): Classical & EBRD Preferences

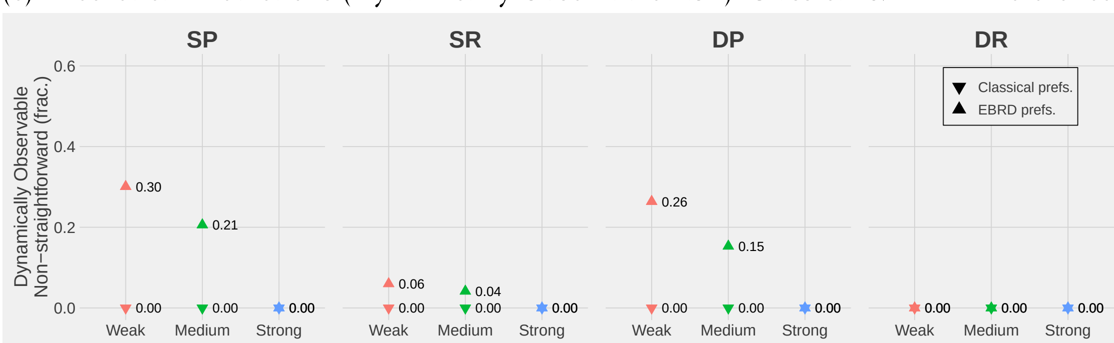

(d) Results (Dynamically Observable NSF)

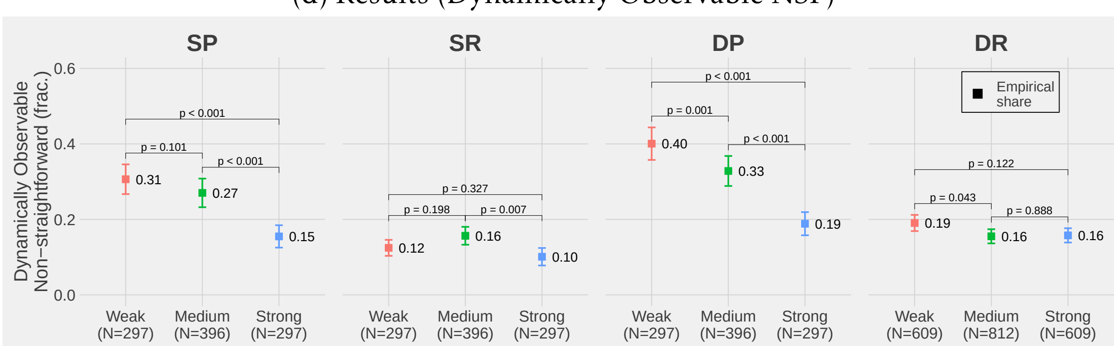

Notes: Fraction of costly and dynamically observable NSF. Panels (a) and (c) present predictions for the share of costly NSF and dynamically observable NSF by problem type under Classical (downward triangles) and EBRD (upward triangles) preferences. Panels (b) and (d) present the empirical shares. Error bars: standard errors from a regression of NSF behavior on problem type, clustering at the individual level. p-values: Wald tests.

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.

Tilman Börgers and Jiangtao Li. Strategically simple mechanisms. *Econometrica*, 87(6): 2003–2035, 2019. doi: https://doi.org/10.3982/ECTA15897. URL https://onlinelibrary.wiley.com/doi/abs/10.3982/ECTA15897.

Bnaya Dreyfuss, Ori Heffetz, and Matthew Rabin. Expectations-based loss aversion may help explain seemingly dominated choices in strategy-proof mechanisms. *American Economic Journal: Microeconomics*, forthcoming.

Federico Echenique, Alistair J Wilson, and Leeat Yariv. Clearinghouses for two-sided matching: An experimental study. *Quantitative Economics*, 7(2):449–482, 2016. https://doi.org/10.3982/QE496.

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.

Yannai Gonczarowski, Ori Heffetz, and Clayton Thomas. Self-explanatory strategyproof mechanisms. Presentation at Harvard Theory of Computation Seminar on March 23, 2022.

Rustamdjan Hakimov and Dorothea Kübler. Experiments on centralized school choice and college admissions: a survey. *Experimental Economics*, 24(2):434–488, July 2020. doi: 10.1007/s10683-020-09667-7. URL https://doi.org/10.1007/s10683-020-09667-7.

Avinatan Hassidim, Assaf Romm, and Ran I. Shorrer. The limits of incentives in economic matching procedures. *Management Science*, 67(2):951–963, February 2021. doi: 10.1287/mnsc.2020.3591. URL https://doi.org/10.1287/mnsc.2020.3591.

Alexander Karpov. A necessary and sufficient condition for uniqueness consistency in the stable marriage matching problem. *Economics Letters*, 178:63–65, 2019. ISSN 0165-1765.

doi: https://doi.org/10.1016/j.econlet.2019.02.022. URL https://www.sciencedirect.com/science/article/pii/S0165176519300679.

Flip Klijn, Joana Pais, and Marc Vorsatz. Static versus dynamic deferred acceptance in school choice: Theory and experiment. *Games and Economic Behavior*, 113:147–163, 2019. ISSN 0899-8256. doi: https://doi.org/10.1016/j.geb.2018.08.009. URL https://www.sciencedirect.com/science/article/pii/S0899825618301453.

Botond Kőszegi 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 Kőszegi 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 Kőszegi 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.

Marek Pycia and Peter Troyan. A Theory of Simplicity in Games and Mechanism Design, 2021.

Alex Rees-Jones and Samuel Skowronek. An experimental investigation of preference misrepresentation in the residency match. *Proceedings of the National Academy of Sciences*, 115(45):11471–11476, 2018.

Antonio Rosato. Loss aversion in sequential auctions: Endogenous interdependence, informational externalities and the" afternoon effect". *working paper*, 2014.

Ran I Shorrer and Sándor Sóvágó. Dominated choices in a strategically simple college admissions environment: The effect of admission selectivity. 2017. http://dx.doi.org/10.2139/ssrn.3068672.