Improving Transparency and Verifiability in School Admissions: Theory and Experiment

Rustamdjan Hakimov (WZB Berlin)

Madhav Raghavan (University of Lausanne)

Discussion Paper No. 376

January 27, 2023

Improving Transparency and Verifiability in School Admissions:

Theory and Experiment[*]

Rustamdjan Hakimov[†] and Madhav Raghavan[‡]

January 12, 2023

Abstract

Students participating in centralised admissions procedures do not typically have access to the information used to determine their matched school, such as other students' preferences or school priorities. This can lead to doubts about whether their matched schools were computed correctly (the 'Verifiability Problem') or, at a deeper level, whether the promised admissions procedure was even used (the 'Transparency Problem'). In a general centralised admissions model that spans many popular applications, we show how these problems can be addressed by providing appropriate feedback to students, even without disclosing sensitive private information like other students' preferences or school priorities. In particular, we show that the Verifiability Problem can be solved by (1) publicly communicating the minimum scores required to be matched to a school ('cutoffs'); or (2) using 'predictable' preference elicitation procedures that convey rich 'experiential' information. In our main result, we show that the Transparency Problem can be solved by using cutoffs and predictable procedures together. We find strong support for these solutions in a laboratory experiment, and show how they can be simply implemented for popular school admissions applications involving top trading cycles, and deferred and immediate acceptance.

1 Introduction

Market design has thrived in recent years, both in developing allocation mechanisms for a variety of applications, as well as in putting them into practice.[1] However, designers around the world have sometimes had to address concerns that the mechanisms in use are non-transparent, in that participants cannot verify that their assignments are correct. In two prominent examples, Google had to update its Ad Manager auction mechanism in 2019 in order to provide 'transparency to both publishers and advertisers,'[2] while the US Federal Communications Commission changed its auction for spectrum so that 'bidders and other interested parties [could] verify that the rules are followed' (Cramton and Schwartz, 2000).

Similar perceptions of lack of transparency have also affected one of the most successful market design applications—centralised school and college admissions. The New Orleans Recovery School District in the US ditched the top-trading-cycles (TTC) mechanism after just one year of use, likely because 'TTC was perceived as difficult... to understand... why a child was assigned a seat over a child who was not' (Abdulkadiroğlu et al., 2017). In France, an overhaul of the university admissions system followed protests against the lack of transparency of the existing deferred-acceptance (DA) based system:[3]

> A lot of people see [the algorithm] as a black box, they don't understand how the student selection process happens... [Y]ou have to create the conditions of fairness of the algorithm and of its full transparency... [Otherwise] people will eventually reject this innovation.[4]

In addition to the obvious problems of perception, non-transparency sometimes also imposes a direct cost. In 2018, the UK Department of Education received around 60,000 formal appeals (nearly 4% of all admissions cases).[5] A large fraction of these appeals are unsuccessful (nearly 77% in 2018), and many of these can be traced to a lack of transparency.[6] Improving transparency might reduce the time and resources spent in handling such appeals.

---

[1] A non-exhaustive list of applications includes Roth and Peranson (1999); Roth et al. (2004); Sönmez and Switzer (2013); Andersson and Svensson (2014); Budish et al. (2015); Hakimov et al. (2021a).

[2] "An update on first price auctions for Google Ad Manager", 10 May 2019, https://www.blog.google/products/admanager/update-first-price-auctions-google-ad-manager/ (last accessed 02 Oct 2020).

[3] The proposed law: LegiFrance, 8 Mar 2018. https://www.legifrance.gouv.fr/affichTexte.do?cidTexte =JORFTEXT000036683777 (last accessed 15 Nov, 2019). See also, The Guardian, 5 Apr 2018. https://www.theguardian.com/world/2018/apr/05/we-cant-back-down-french-students-dig-in-for-macron-battle .

[4] Interview with Emmanuel Macron (President of France), WIRED, 31 March 2018 (last accessed 15 November, 2019). http://www.wired.com/story/emmanuel-macron-talks-to-wired-about-frances-ai-strategy/.

[5] UK Department of Education. https://assets.publishing.service.gov.uk/government/uploads/system/uploads/attachment_data/file/826573/2019_Admissions_Appeals_Release.pdf (last accessed 15 Nov 2019).

[6] UK Office of the Schools Adjudicator, 2018. https://assets.publishing.service.gov.uk/government/uploads/system/uploads/attachment_data/file/680003/2017_OSA_Annual_Report_-_Final_23_January_2018.pdf (last accessed 15 Nov 2019).

We contribute to the school admissions literature by formalising the transparency problem for centralised school and college admissions, and proposing solutions for it. To motivate this exercise, consider the celebrated student-proposing deferred acceptance (SPDA) school admissions rule (Gale and Shapley, 1962; Abdulkadiroğlu and Sönmez, 2003). Assume initially that the designer knows each student's score at each school, and runs an elicitation procedure that asks each student to report a comprehensive ranking of their preferences over schools. Based on this information, the SPDA rule computes the matching using the simple and well-known SPDA algorithm:

> An application is initially sent from each student to her most-preferred acceptable school. Each school 'tentatively accepts' the applicants with the highest scores, up to its capacity, rejecting all excess applications. In each subsequent step, an application is sent from each student rejected in the previous step to her most-preferred remaining acceptable school. Each school considers its revised set of applications, including any previously tentatively accepted applications, and tentatively accepts the applications among these with the highest scores, up to its capacity, rejecting all others. When no more rejections are made, the last tentatively accepted applications become final.

To illustrate the lack of transparency, consider a simple example with three students $\{1, 2, 3\}$ and three schools $\{x, y, z\}$, each of unit capacity. A school admissions 'problem' (a collection of students' preferences and their scores at each school) is presented on the left of Figure 1 below. It can be quickly verified that the matching in boxes is the result of SPDA: Student 1 is matched to School z, Student 2 to School x, etc.

Figure 1: An Illustration

A school admissions problem

Preferences			Scores			
1	2	3	Student	x	y	z
x	$\boxed{x}$	z	1	40	95	$\boxed{85}$
$\boxed{z}$	y	$\boxed{y}$	2	$\boxed{80}$	45	95
y	z	x	3	90	$\boxed{80}$	50

Student 1's information

Preferences			Scores			
1	2	3	Student	x	y	z
x			1	40	95	$\boxed{85}$
$\boxed{z}$			2			
y			3			

Notes: On the left of each table, a school in a higher row represents a school preferred by the student to a school in a lower row, e.g., Student 1 prefers School x to School z. On the right side of each table, each row represents a student's score at each school, e.g., Student 1 has a score of 40 at School x, 95 at School y, etc. A student with a higher score has a higher priority at that school.

However, while the designer knows the entire problem, each student typically knows only her own score at each school and the preferences she reported to the elicitation procedure.[7] Student 1's 'intrinsic' information is thus quite limited, and is illustrated in the table on the right of Figure 1. Notice that,

---

[7]This is admittedly a 'worst case' scenario; she might know some reported preferences or scores of other students.

even if Student 1 believes that SPDA was used to compute the matching, her intrinsic information is too limited for her to check whether her match to School z is legitimate and that she should not, in fact, have been matched to School x, which she prefers. We call this the Verifiability Problem. At a deeper level, Student 1 might wonder if the SPDA rule was used at all.[8] We call this the Transparency Problem.[9] Thus, verifiability assumes participants believe the promised mechanism was used and want to verify the outcome. At the same time, transparency requires that participants could verify both the outcome and the use of the promised mechanism.

Intuitively, both the Verifiability Problem and the Transparency Problem are problems of limited information on the part of students, since they do not have access to the information used by the designer to compute the matching. As such, both problems should be solvable by providing students adequate feedback. Indeed, an easy way is by full disclosure: the designer reveals all the information on preferences and scores that she used to compute the matching. Each student can presumably recreate the SPDA algorithm and ascertain that her match was correct and determined by the promised rule. However, while easy to implement, this involves revealing sensitive private information, like reported preferences and priorities of all participants, and might face privacy challenges.

In this paper, we establish two simple-to-implement forms of feedback that address both problems without revealing private information directly. For the first, notice that the scores of students at their matched school are effectively the lowest score required to be matched to that school. In the example above, the score of Student 1 at School x (40) is not enough for her to be matched to it under SPDA (the cutoff score is 80), and similar is true for other students at any school they prefer to their match. In our first result, we show that publicly communicating these kinds of cutoffs solves the Verifiability Problem; each student can check that their matched school is the best possible for them under SPDA. The exact nature of cutoffs depends on the application, and has been derived explicitly for TTC (Dur and Morrill, 2018; Leshno and Lo, 2020) and DA (Azevedo and Leshno, 2016). Our result holds for these rules. We add to this literature by deriving cutoffs for the immediate acceptance (or 'Boston') mechanism. In general, while cutoffs have been informally justified as improving transparency, we are the first to formalise exactly how and why this is the case. As a side benefit, cutoffs are relatively privacy-preserving.

Our second contribution is to show that the Verifiability Problem can also be solved by changing the underlying preference elicitation procedure. In the above example, students report their preferences to the designer as a ranking over all schools in a single stage. In contrast, we introduce a multi-stage elicitation procedure in which, loosely speaking, students report only one school at a time, and are contacted again

---

[8] The 'college-proposing' DA (CPDA) rule, for example, is compatible with the same inputs but can yield a different match.

[9] One can think of verifiability as a problem of verifying the outcome, while transparency is a problem of verifying the mechanism.

only when necessary.[10] We model the feedback acquired by students through these multiple interactions as 'experiential information'.[11] In effect, these elicitation procedures provide rich feedback by being 'predictable': each student knows that either (a) she will be matched to the last school she reported or (b) she will be contacted again. We show that this experiential feedback is rich enough to solve the Verifiability Problem without requiring the designer to communicate any additional information (Proposition 1 and Theorem 2). These results rationalise the recent switch to multi-stage mechanisms in some college admissions jurisdictions (Gong and Liang, 2016; Bó and Hakimov, 2018; Grenet et al., 2022).

So far we have assumed that the student only requires proof of the validity of her match under the promised admission rule. We now turn to the Transparency Problem which, in addition to the Verifiability Problem, also involves demonstrating that the promised rule was indeed used.[12] Without taking a position on the honesty of the designer, it should be seen that solving the Transparency Problem is useful even for a perfectly honest designer, as it provides a greater reassurance to students about their match. However, neither solution to the Verifiability Problem—published cutoffs and predictable elicitation procedures—solve the Transparency Problem by themselves; we show that each retains the theoretical possibility that a different rule was used without contradicting any student's information. Our final theoretical contribution is to show that, for a class of school admissions environments that includes most popular applications, using a predictable elicitation procedure *and* publishing cutoffs afterwards can *together* guarantee that the promised rule was used, and used correctly.[13] This surprising result has obvious practical significance: predictable elicitation procedures are easy to implement and cutoffs are simple to compute at least for some mechanisms (as observed in the recent changes to college admissions in France). We use these results to show how transparency can be achieved for 'student-proposing' and 'college-proposing' DA, top-trading cycles, and immediate acceptance.

Our theoretical results rest on students' ability to use informational feedback to verify their matches. To test this ability, we run a laboratory experiment based on the widely used student-proposing DA (SPDA) rule in school admissions. Fixing SPDA, we use four treatments that vary in the designer's communication and the choice of elicitation procedure. The first environment uses direct (one-stage) DA with no cutoffs (DirNo). This is the classical implementation of DA, and is unverifiable through either experience or

---

[10]In the example, Student 1 reports School x when initially contacted, and she is contacted again only when School x is no longer possible for her (due to other students reported preferences and scores), at which point she reports School z.

[11]The idea of experiential information appears informally in studies on dynamic (ascending or descending) auctions, which are perceived to be more transparent than sealed-bid equivalents, as participants can observe bids and infer the actions of other bidders (Cramton, 1998; Ausubel, 2004). We are the first to formalise this idea for centralised admissions.

[12]We should point out that our definition of transparency is limited, in the sense that it excludes the possible manipulation of 'inputs', such as school capacities or students' scores at a school.

[13]Technically, we show that a rule that is 'outcome-equivalent' to the promised rule must have been used. Our result holds for 'non-wasteful' mechanisms and 'structured-offers' elicitation procedures.

communication. The second treatment involves a predictable elicitation procedure, with no feedback from communication (SeqNo), and is theoretically verifiable only through experience. The third treatment uses a direct DA procedure, but communicates 'step-cutoffs' corresponding to the cutoffs at each step of the underlying DA algorithm (DirCutoffs). This treatment is verifiable through communication. Finally, we combine the two by using a predictable elicitation procedure as well as communicating step-cutoffs (SeqCutoffs) which, theoretically, is transparent. In a novel design feature, in each treatment subjects are told that their matches are randomly determined half of the time. After learning their matches, subjects are given the option to incur a cost and appeal. Payoffs are constructed so that it is optimal to appeal if (and only if) the match is random. Thus, the correctness of decisions of whether or not to appeal provides a sense of whether subjects can correctly spot random matches (in particular, matches produced not by the explained mechanism), and provides an experimental measure of observed verifiability.

We find that, first, the proportion of correct decisions of whether or not to appeal in DirNo is 52%; subjects do no better than chance. This confirms the complexity of verifying DA matches and, thus, the public perception of its lack of transparency. Second, feedback from experience (SeqNo, 64%) and communication (DirCutoffs, 65%) significantly increases the proportion of correct decisions relative to DirNo. Third, the proportion of correct decisions is the highest under SeqCutoffs (83%), significantly higher than even SeqNo and DirCutoffs.[14] This strongly suggests that a second source of feedback plays an important role in empirical verifiability. Our experiment thus supports using both sources of feedback even when designer honesty is not in doubt. Moreover, we believe our novel experimental setting is of independent interest, especially to compare the relative verifiability of different environments.

1.1 RELATED LITERATURE

Transparency has long been proposed as a desirable normative criterion for allocation processes, for instance, in education (West, Pennell, and Noden, 1998), public administration (Meijer, 2013), and also decentralised allocation settings like financial markets (Asquith, Covert, and Pathak, 2019) and multilateral organisations (Nelson, 2001). We add to this literature by offering—to our knowledge—the first systematic analysis of transparency for centralised school and college admissions.

There are, however, several studies on related issues. Two recent papers study the same problem: transparency in the allocation markets without transfers. Grigoryan (2022) develops the notion of transparency based on the number of agents necessary to detect the deviations of the designer, showing that serial dictatorship and immediate acceptance are the most transparent mechanisms as just two agents could detect the deviations. A similar finding for the serial dictatorship mechanism is established by

---

[14]Notice that this is despite the fact that all treatments except DirNo are theoretically verifiable and should have 100% correct appeal decisions.

Pycia and Ünver (2020). Compared to Grigoryan (2022) our notions of transparency are independent and complement each other. Our notions of transparency are binary and should be used as a guide for the designer to improve the transparency of the environment, given the choice of an allocation rule. At the same time, Grigoryan (2022) allows relative comparison between allocation rules, as the notion is continuous (transparency index). Möller (2022) also introduces a notion of transparency in allocation markets based on the idea of the designer's commitment to an announced strategy-proof mechanism, where the designer and the agents communicate privately. Agents know the announced mechanism, all objects' priorities, and all agents' assignments. In this setup, only stable allocations might be transparent. At the same time, we consider public communication, and our notion is the feature of the information structure, which leads to the possibility of reaching transparency in efficient and not strategy-proof mechanisms. Another difference is that we do not assume that agents know the assignments of all other agents, which is essential, as it requires us to reveal the cutoffs for the verifiability of SPDA. Note that the assumption of Möller (2022) about complete information on priorities and assignments implies knowledge of the cutoffs in stable mechanisms. A close relation to our paper—although in the context of an auction —is Woodward (2020), who proposes a notion of an 'audited' auction as one in which enough information is disclosed post-auction for participants to check that the auction has been run as claimed. Auditability is similar to our notion of verifiability in our paper. Auditability is also studied by Pycia and Ünver (2020) in the context of Arrovian social welfare functions.

An important recent development is to consider the incentives of the designer, not just participants. Designer-incentive compatibility in auctions is studied as 'credibility' in Akbarpour and Li (2020). On one hand, credibility is stronger than transparency, as it preserves incentive-compatibility even under private communication. On the other hand, credibility is a feature of the mechanism alone, while transparency depends also on the communication of information. Naor, Pinkas, and Sumner (1999) also consider the incentives of an 'auction issuer' who facilitates the computation of auctions.

Another recent development is to consider whether mechanisms are easy for participants to understand and 'play optimally,' for instance 'obvious strategy-proofness' and simplicity of mechanisms (Li, 2017; Pycia and Troyan, 2019). Relatedly, Núñez (2019) proposes transparency of a mechanism as a participant's ability to (perfectly) understand the consequences of her actions, and proposes measuring it by comparing experimental behaviour to theoretical predictions. While transparency could be interpreted as a simplicity-like attribute of a mechanism, we abstract from strategic issues of truthful revelation of preferences.

The role of information communicated by the designer is studied, for instance, as optimal revelation in auctions (Milgrom and Weber, 1982; Gal-Or, Gal-Or, and Dukes, 2007), certifiable pre-play communication (Hagenbach, Koessler, and Perez-Richet, 2014), or as communication requirements for efficiency and

stability in auctions or matching (Nisan and Segal, 2006; Segal, 2007; Gonczarowski, Nisan, Ostrovsky, and Rosenbaum, 2014; Ashlagi, Braverman, Kanoria, and Shi, 2017). Important in this context is the role of cutoffs in improving transparency. In recent work, Azevedo and Leshno (2016) develop cutoffs as 'clearing prices' in the assignment of students to schools for the SPDA mechanism, with a similar 'competitive equilibrium' result established for TTC by Dur and Morrill (2018). Leshno and Lo (2020) identify the structure of cutoffs for TTC, which are more complicated than those for SPDA and are multi-dimensional, as they require specifying the minimum priority required at one school in order to be admitted to another. Immorlica, Leshno, Lo, and Lucier (2020) emphasise the importance of providing historical cutoffs in college admissions in helping students form their preferences, while Hakimov et al. (2021b) provide empirical support to this argument from a laboratory experiment. Cutoffs are a central feature of the communications from the designer in this paper.

Our paper also relates to work on multi-stage mechanisms. The role of ascending auctions in providing information to participants is studied for single units in Cramton (1998) and multiple units in Ausubel (2004, 2006). We support these arguments by offering a definition for experiential information. In school admissions, the sequential implementation of stable allocations is studied in Bó and Hakimov (2018); Haeringer and Iehlé (2019). Li (2017) considers the sequential implementation of the serial dictatorship mechanism, while Bó and Hakimov (2020b) do the same for TTC. Dur and Kesten (2019) consider sequential systems comprising two stages, where all unassigned participants from the first stage can participate in the second stage. An empirical evaluation of multi-stage mechanisms recently adopted in Inner Mongolia (China), Brazil, and Germany can be found in Gong and Liang (2016), Bó and Hakimov (2018), and Grenet, He, and Kübler (2022). The mechanisms in Inner-Mongolia and Brazil also emphasise the extensive use of intermediate cutoff communication in practice. We show that some of these sequential mechanisms can be used to improve transparency.

We also contribute to a growing experimental literature on matching mechanisms, recently surveyed in Hakimov and Kübler (2020) and Pan (2020).

2   The Model

2.1   Preliminaries

Throughout, concepts are presented in **boldface**, while definitions are underlined. There is a finite set of **students** N and a finite set of **schools** X, which includes an underlined unassigned or outside option $\emptyset \in X$. Each school $x \in X$ has a **capacity** $q_x \in \{1, \ldots, |N|\}$ that determines the maximum number of students it can accommodate. We assume $q_\emptyset = |N|$ (the unassigned option has unrestricted capacity). A **matching** is a vector $a \in X^N$ with $a_i \in X$ denoting the **match** of $i \in N$. A matching $a \in X^N$ is feasible if it

respects capacity constraints for schools, i.e., $|\{i \in N \mid a_i = x\}| \leq q_x$ for each $x \in X$. The set of **feasible matchings** is denoted $\mathcal{A} \subset X^N$.

Preferences for student $i \in N$ are given by a linear ordering $\succeq_i$ over X. By convention, $x \succeq_i y$ means x is 'at least as good as' y and $x \succ_i y$ means x is 'preferred to' y. A school $x \in X$ is <u>acceptable</u> to i under $\succeq_i$ if it is at least as good as being unassigned, i.e., $x \succeq_i \emptyset$. A **preference profile** $\succeq = (\succeq_i)_{i \in N}$ collects preferences for each student, and the set of preference profiles is denoted $\mathcal{R}$. Each student $i \in N$ has a **score** at each school $x \in X$, given by a real number $v_i^x \in \mathbb{R}_+$. For a school $x \in X$, the vector of scores for each student given by $v^x = (v_i^x)_{i \in N}$ represents the **priorities** for school $x \in X$.[15] A **priority profile** is a collection of priorities for each school and is denoted $v = (v_i^x)_{i \in N}^{x \in X}$, while the set of all priority profiles is denoted $\mathcal{V}$. The **domain of (admissions) problems** is $\mathcal{P} \equiv (N, X, (q_x)_{x \in X}, \mathcal{A}, \mathcal{R}, \mathcal{V})$. For simplicity, we denote a **problem** $P \in \mathcal{P}$ simply as $P = (\succeq, v)$. In line with a large part of the literature on school admissions, we focus in this paper on one-sided matching, where schools are treated as 'objects' to be consumed. Henceforth, we fix a priority profile v^*, and assume that the designer knows v^* (we also assume that each student initially knows only her own score v_i^{*x} at each school $x \in X$).

While the designer initially knows the priority profile v^*, she does not know students' preferences. In principle, any problem $P = (\succeq, v^*)$ could arise, depending on the preferences reported by students. We shortly describe how the designer acquires information on students' preferences, but notice already that any information she acquires refines the set of problems further. For example, if the designer learns the top-ranked school in each student's preferences, she can already eliminate all preference profiles (and thus problems) that are not compatible with this information. In general, we ignore the 'nature' of information and associate it directly with the subset of problems it induces; **information** is given by a subset of problems $\phi \subseteq \mathcal{P}$, where $P \in \phi$ is <u>compatible</u> with ϕ. Let $\Phi \equiv 2^{|\mathcal{P}|}$ be the set of all possible information.

2.2 SCHOOL ADMISSIONS MECHANISMS

As standard in the literature starting with Abdulkadiroğlu and Sönmez (2003), a **school admissions rule** is given by a function $f : \mathcal{P} \to \mathcal{A}$ that produces a feasible matching $f(P)$ for any problem $P \in \mathcal{P}$. However, we broaden the definition of a mechanism to incorporate this idea of information, in particular to allow us to describe interim stages of multi-stage elicitation procedures. Thus we denote a **mechanism** (for a school admissions rule f) by a correspondence $g^f : \Phi \to 2^{\mathcal{A}}$ that associates each information $\phi \in \Phi$ with a set of feasible matchings $g^f(\phi) \subseteq \mathcal{A}$, where $g^f(\phi) \equiv \{a \in \mathcal{A} \mid f(P) = a \text{ for some } P \in \phi\}$. It is easy

---

[15]This is a cardinal representation of priorities, where scores could be from an exam, or a metric that awards points for neighbourhood proximity, siblings in the school, and so on. Ordinal priorities, which are more common in the school admissions literature, can be recovered by ranking students in descending order of score. In any case, it is enough to treat the score for a student at a school as simply her position (or 'rank') in the ordinal priorities for a school.

to see that if $\phi = P$ (the information is 'full') then the feasible matching is uniquely determined.[16] In what follows, when the admissions rule is known, we will suppress reference to f.

Every school admissions mechanism is based on an 'eligibility criterion' that determines which of competing students have a higher claim on a school. This idea is implicit throughout the literature. For instance, the eligibility criterion under an SPDA mechanism and its school-proposing counterpart is simply the student scores; a student with a higher score at a school has a higher claim on the school (see, e.g., Azevedo and Leshno (2016)). Elsewhere, the eligibility criterion is complex, such as the matrix-based eligibility criterion for a TTC mechanism (Dur and Morrill, 2018; Leshno and Lo, 2020) or the rank-based one for the Boston mechanism (Doğan and Klaus, 2018). In this paper, we explicitly formulate an eligibility criterion for a school $x \in X$ under mechanism g as a statistic $\rho^{g,x}$ that for each information induces an ordering over students, i.e., for any $i, j \in N$ and any $\phi \in \Phi$, $\rho_i^{g,x}(\phi) > \rho_j^{g,x}(\phi)$ or vice versa. A school admissions mechanism always respects its own eligibility criteria: at each problem, each school is matched to the highest-ranked students according to its eligibility criterion among those who want it. Formally, for information $\phi \in \Phi$, each compatible problem $P = (\succeq, v^*) \in \phi$, any two students $i, j \in N$ and school $x \in X$, if $x = g_j(P) \succ_i g_i(P)$, then $\rho_j^{g,x}(\phi) > \rho_i^{g,x}(\phi)$.[17]

We restrict our attention to 'non-wasteful' mechanisms, which carries the usual meaning of requiring that no school is left unfilled as long as some student prefers to attend it. We describe the elicitation procedure in a mechanism shortly, but for now assume that the 'outcome' is (ϕ^g, a), where ϕ is the designer's final information and $a \in A$ is a matching. A mechanism g is non-wasteful if, for any outcome (ϕ, a) in g, any problem $P = (\succeq, v^*) \in \phi$, any $i \in N$ and any $x \in X$, we have that $x \succ_i a_i \implies |\{j \in N \mid a_j = x\}| = q_x$. One justification is that empty seats in schools are costly for the designer. Another is that empty seats in schools are observable, and can lead to complaints from students. It is straightforward that when $\phi^g = P$, we recover the classical definition of non-wastefulness.

Finally for this section, for a mechanism g, information $\phi \in \Phi$, and student $i \in N$, the possibility set for i collects all matches produced for her by g at ϕ, i.e., $Poss_i^g(\phi) = \{x \in X \mid g_i(P) = x$ for some $P \in \phi\}$. As information becomes more precise (i.e., the smaller the set of compatible problems), the smaller a student's possibility set becomes—an idea that will be key to our notions of verifiability and

---

[16]In fact, many school admissions mechanisms do not require full information in order to uniquely determine a matching. For example, if the top row of students' preferences is itself a feasible matching, then this information (along with students' scores) is sufficient for any mechanism for the SPDA rule to uniquely produce this matching.

[17]Respecting the eligibility criterion can be thought of as an eligibility-based 'no-justified-envy' condition. Notice that an SPDA mechanism produces no-justified-envy matchings with respect to scores at schools (Gale and Shapley, 1962). However, while a TTC mechanism might generate justified-envy matchings with respect to scores (Abdulkadiroğlu and Sönmez, 2003), it generates no-justified-envy matchings with respect to its *own* eligibility criterion. This observation is crucial for the definition of cutoffs for different mechanisms (see Section 7).

transparency.[18] The relation between information and possibility sets is illustrated in Figure 2.

Figure 2: Mechanisms, information and possibility

Notes: The mechanism g associates matchings $g(\phi) \subseteq \mathcal{A}$ with information $\phi \in \Phi$. For a student i, the possibility set $Poss_i^g(\phi) \subseteq X$ is the set of schools matched to i among matchings in $g(\phi)$.

3   THE VERIFIABILITY PROBLEM

At the outset, the designer announces to students that she will use a (non-wasteful) promised mechanism g^*, and explains its rules. Assume for now that each student $i \in N$ reports some preferences $\succeq_i$ to the designer,[19] while the (fixed) priority profile is v^*. Let the 'outcome' be (ϕ^{g^*}, a^*), where $\phi^{g^*} = (\succeq, v^*)$ is the information of the designer, and $a^* = g^*(\phi^{g^*})$ is the (unique) final matching (we will define the outcome more precisely when we later define the elicitation procedure underlying the mechanism). When student $i \in N$ learns that she is matched to a school $x \in X$, how can she establish that $x = a_i^*$, i.e., that her match is—or is not—the one determined by the promised mechanism? This is the Verifiability Problem.

A student's intrinsic information $\phi_i^{g^*} = (\succeq_i, v_i^*)$ consists of her reported preferences and scores at each school, and induces a possibility set $Poss_i^{g^*}(\phi_i^{g^*})$. For most reasonable mechanisms, the possibility set contains multiple schools (for instance, any acceptable school in the reported preferences could be her match, depending on the reports and scores of others). The only way to shrink the possibility set for a student is to provide information via feedback, and our objective in particular is to provide feedback rich enough to remove the possibility of her being matched to any school other than a_i^*, i.e., to provide her with information $\phi_i \in \Phi$ so that $Poss_i^{g^*}(\phi_i^{g^*} \cap \phi_i) = a_i^*$.

3.1   VERIFIABILITY THROUGH COMMUNICATION

The first source of feedback we consider is 'communications' from the designer, itself determined by a 'communication protocol'. We formalise this channel via a ***communication protocol***, given by a function $M : \Phi \to \Phi$, which prescribes a ***communication*** $M(\phi) \in \Phi$ for designer's information $\phi \in \Phi$. As examples

---

[18]It is important to note that the possibility set is different from a 'choice set' - it is not that a student can choose any of the schools in $Poss_i^g(\phi)$; instead, any of these schools could be her match, depending on which problem in ϕ arises.

[19]As we formalise later when we describe the elicitation procedure in detail, we do not take a position on whether students report their 'true' preferences or not; strategy-proofness is orthogonal to verifiability and transparency in our analysis.

of two extremes, for designer's information ϕ^g, the 'full disclosure' protocol communicates $M(\phi^g) = \phi^g$ while the 'null' protocol communicates $M(\phi^g) = \emptyset$. Importantly, communications are assumed to be public and common to all students, e.g., via announcements on public websites, newspapers, etc.[20]

We say that a mechanism is 'verifiable' by a communication protocol if, for any outcome of the mechanism and any student, the communication (along with the student's intrinsic information) is sufficient to uniquely pin down her match:

DEFINITION **1.** Mechanism g^* is <u>verifiable by communication protocol M^*</u> if, for outcome (ϕ^{g^*}, a^*) and each $i \in N$: $Poss_i^{g^*}(\phi_i^{g^*} \cap M^*(\phi^{g^*})) = a_i^*$.

3.1.1 CUTOFFS

It is easy to see that the full disclosure protocol achieves verifiability (the designer's information produces a unique matching for the promised mechanism). We now show that the simple-to-use cutoff protocol is in fact sufficient for verifiability. Cutoffs have of course been long-studied in the school admissions literature, but we redefine the idea in line with our information-based framework. Let g be a mechanism, and let $\rho^{g,x}$ be the corresponding eligibility criterion for each $x \in X$. A 'cutoff' for a school in a matching is the minimum value of the eligibility criterion among students matched to that school. Formally, a <u>cutoff</u> for $x \in X$ in $a \in \mathcal{A}$ under g at information $\phi \in \Phi$ is given by $c_x^g(a, \phi) \equiv \min\{\rho_i^{g,x}(\phi) \mid a_i = x\}$.[21] Let $c^g(a, \phi) = (c_x^g(a, \phi))_{x \in X}$ be the (vector of) cutoffs for a under g at ϕ. A communication protocol M is a <u>cutoff protocol</u> for g if it communicates the cutoffs at the outcome, i.e., $M(\phi^g) = c(a, \phi^g)$ for outcome (ϕ^g, a) in g. Then:

THEOREM **1.** *A non-wasteful school admissions mechanism is verifiable by a communication protocol if and only if the communication is 'informationally equivalent' for each student to cutoffs. Formally, let g^* be a non-wasteful mechanism, and let (ϕ^{g^*}, a^*) be the outcome in g^*. Then $Poss_i^{g^*}(\phi_i^{g^*} \cap M^*(\phi^{g^*})) = a_i^*$ for each $i \in N$ if and only if $Poss_i^{g^*}(\phi_i^{g^*} \cap M^*(\phi^{g^*})) = Poss_i^{g^*}(\phi_i^{g^*} \cap c^{g^*}(a^*, \phi^{g^*}))$ for each $i \in N$.*

The proof is in Appendix A and is fairly intuitive, since cutoffs at a matching are essentially the clearing prices according to the eligibility criterion. Of course, the full disclosure protocol also achieves verifiability through communication. So do communication protocols that communicate an eligibility criterion value

---

[20]This assumption is quite plausible in school admissions, where the designer is often constrained to make all announcements public and commonly accessible by students. In the wider mechanism design literature, this distinguishes us from Akbarpour and Li (2020), who assume communications between an auctions designer and bidders to be private and individualised. Private or semi-private communication between the designer and students is outside our present scope.

[21]This is consistent with the idea of cutoffs serving as 'clearing prices' for the mechanism, as described for the SPDA mechanism in Azevedo and Leshno (2016), and for the TTC mechanism in Dur and Morrill (2018) and Leshno and Lo (2020). The units of the price are determined by the eligibility criterion. We always set the cutoff for the unassigned option $\emptyset$ to zero.

that is $c_x(\phi^{g^*}, a^*) - \psi_x$ for each school x, when ψ_x is within limits.[22] However, it is straightforward to see that these communications are informationally equivalent to cutoffs, since they induce the same possibility set for each student. Theorem 1 establishes this result more generally.

EXAMPLE **1.** We show how cutoffs achieve verifiability through communication via a simple example. Let $N = \{1, 2, 3\}$ be three students, and let $X = \{x, y, z\}$ be three schools, each of unit capacity. Consider the school admissions problem in Figure 3 below. The unique feasible matching under a one-stage SPDA mechanism is marked in boxes, as is the score of the matched student at each school. Since only one student is matched to each school, these scores in boxes also represent the cutoffs for that school.

Figure 3: Cutoffs for SPDA

Preferences			Scores			
$\succeq_1$	$\succeq_2$	$\succeq_3$	Student	v^x	v^y	v^z
x	$\boxed{x}$	z	1	40	95	$\boxed{85}$
$\boxed{z}$	y	$\boxed{y}$	2	$\boxed{80}$	45	95
y	z	x	3	90	$\boxed{80}$	50

The cutoffs $(80, 80, 85)$ achieve verifiability through communication. To see this, consider Student 3, who is matched to School y, but prefers School z. The cutoffs demonstrate that she does not have a high enough score at School z (50 versus a cutoff of 85), but does have a high enough score for her match, School y (after all, her score is the cutoff). Even though she has a higher score at School x than the cutoff (90 versus a cutoff of 80), given the ranking of School y above School x in her actions, she cannot be matched to School x since the SPDA mechanism used respects its own eligibility criterion. Thus, her possibility set contains only her match. The same is true for other students. ∎

3.2 VERIFIABILITY THROUGH EXPERIENCE

In Theorem 1, we have shown that verifiability can be achieved by the common communication of cutoffs. In this section, we show how verifiability can also be achieved by the suitable design of the elicitation procedure by which students report their preferences to the designer. Multi-stage elicitation procedure have long been used in practice, and our contribution is to formalise the procedure in a way that allows us to describe both one-stage and multi-stage elicitation of preferences.

---

[22]In particular, ψ_x should not be large enough for $c_x(\phi^{g^*}, a^*) - \psi_x$ to fail to be a clearing price, i.e., to fall below the highest value of the eligibility criterion for x among students who ranked x higher than their matches in a^*.

3.2.1 THE ELICITATION PROCEDURE

Students participate in the elicitation procedure by 'taking actions' that involve ranking (a subset of) schools. Formally, an **action** for student $i \in N$ is given by a reflexive, complete, transitive, and anti-symmetric binary relation $\geq_i$ over a set $\hat{Y} \subseteq X$ (all schools in $\hat{Y}$ are ranked strictly).[23] Each elicitation procedure has a set of <u>permitted actions</u> determined by restrictions on $\hat{Y}$. For example, one-stage procedures require all schools to be ranked ($\hat{Y} = X$), while others might require students to rank their best k schools ($|\hat{Y}| = k$).[24] An **action profile** is denoted $\geq \equiv (\geq_i)_{i \in N}$. Students' actions in an elicitation procedure are interpreted as their preferences over schools, and thus are information in their own right.

We allow for elicitation of preferences in multiple stages. A variable indexed by k refers to the punctual variable for stage k, whereas a variable indexed by $\{k\}$ refers to the cumulative variable for all stages up to and including k. Let g be a mechanism, and let $\geq^{g,0}$ be an initial 'empty' action profile with $\hat{Y}_i^{g,0} = \emptyset$ for each $i \in N$. Thus $\phi^{g,0} \equiv (\geq^{g,0}, v^*)$ represents the initial information of the designer.

<u>Stage k, $k \geq 1$:</u>

1. The procedure <u>contacts</u> a set of students $Act^{g,k} \subseteq N$.

2. Each contacted student $i \in Act^{g,k}$ is <u>offered</u> a set of schools $Y_i^{g,k}$ which contains at least one school she has not already ranked, i.e., $Y_i^{g,k} \setminus \hat{Y}_i^{g,\{k\}} \neq \emptyset$.

3. Each contacted student $i \in Act^{g,k}$ <u>takes a permitted action</u> $\geq_i^{g,k}$ such that (1) actions are based on schools offered to her at that stage: $\hat{Y}_i^{g,k} \subseteq Y_i^{g,k}$ (2) actions are consistent with previous actions: $x >_i^{g,\{k-1\}} y \implies x >_i^{g,k} y$ for all $x, y \in \hat{Y}_i^{g,k}$; and (3) all schools in her actions so far are ranked higher than all offered schools that are not: $x >_i^{g,\{k\}} y$ for all $x \in \hat{Y}_i^{g,\{k\}}$ and $y \in \hat{Y}_i^{g,\{k\}} \setminus Y_i^{g,k}$.[25]

4. Non-contacted students do not take part. Therefore, set $Y_j^{g,k} = \emptyset$ and $\geq_j^{g,k} = \emptyset$ for all $j \notin Act^{g,k}$, and let $\geq^{g,k}$ be the action profile in this stage.

5. The designer's information is updated to $\phi^{g,\{k\}} = (\geq^{g,\{k\}}, v^*)$, which includes newly taken actions. If the designer's information produces a unique **realised matching** $a \in g(\phi^{g,\{k\}})$, the designer (privately) announces the **realised match** a_i to each $i \in N$, and the procedure terminates. Otherwise, the procedure goes to Stage $k+1$.

---

[23] This definition of an action is not universal; we rule out actions in which schools in $\hat{Y}$ are bifurcated into an acceptable and unacceptable subset (all schools in $\hat{Y}$ must be ranked strictly), or in which students respond to yes/no questions.

[24] For the technically-minded reader, actions comprise the message space for mechanism g. Importantly, we also restrict the strategy space by assuming that student i with preferences ('type') $\succeq_i$ has a type $\succeq_i^g$ for mechanism g, such that, for each action $\geq_i$ she takes in g and each $x, y \in \hat{Y}$ ranked in the action, $x >_i y \implies x \succ_i^g y$. Strategy-proofness requires $\succeq_i^g = \succeq_i$ to be a weakly dominant strategy; we do not impose this condition.

[25] The third condition assumes that actions reveal the ranking of preferred schools. We make this assumption for notational parsimony, noting that our results can accommodate elicitation procedures that involve students ranking their 'worst' schools.

Since students provide strictly more information each time they are contacted, the elicitation procedure terminates finitely. In general, procedures can differ in: (1) which students are contacted in a stage, (2) which schools they are offered, and (3) permitted actions. Each elicitation procedure effectively defines a new mechanism, so we often refer to the two interchangeably. We can now formalise the ***outcome*** in g as the pair (ϕ^g, a), where ϕ^g is the designer's terminal information in g, and $a \in g(\phi^g)$ is the (unique) realised matching. Figure 4 contains a handy reference of important notation from the elicitation procedure that we use frequently in what follows.

Figure 4: Frequently used notation from the elicitation procedure for mechanism g

Notation	Description
$Act^{g,k}$	Contacted students at stage k
$Y_i^{g,k}$	Schools offered to i at k
$\geq_i^{g,k}$	Action of i at k
$\geq_i^{g,\{k\}}$	Cumulative actions of i at k
K^g	Terminal stage of g
$\geq_i^g \equiv \geq_i^{g,\{K^g\}}$	Cumulative actions of i in g
$\phi_i^{g,k} = (\geq_i^{g,k}, v_i^*)$	Intrinsic information for i at k
$\phi_i^{g,\{k\}} = (\geq_i^{g,\{k\}}, v_i^*)$	Cumulative intrinsic information for i at k
$\phi_i^g = \phi_i^{g,\{K^g\}}$	Intrinsic information for i in g
$\phi^{g,k} = (\geq^{g,k}, v^*)$	Designer's information at k
$\phi^{g,\{k\}} = (\geq^{g,\{k\}}, v^*)$	Cumulative designer's information at k
$\phi^g \equiv \phi^{g,\{K^g\}}$	Designer's terminal information in g
$a \in g(\phi^g)$	Realised matching in g
(ϕ^g, a)	Outcome in g

EXAMPLE **2.** The SPDA rule has a number of variants in terms of elicitation procedures (and thus mechanisms). We call the mechanism with a one-stage procedure a <u>Direct DA mechanism (DirDA)</u>: all students are contacted and offered all schools that are initially possible for them, and permitted actions involve strictly ranking all offered schools (or at least all acceptable schools).

One can think of several multi-stage elicitation procedures for SPDA. For illustration, we highlight a simple mechanism—<u>Single-School Sequential SPDA (SeqDA)</u>—that effectively turns the steps of the SPDA algorithm into stages in the elicitation procedure. That is, in the first stage, all students are contacted and are offered all schools, and permitted actions involve selecting exactly one of these schools (or the outside option). This school is interpreted as the top-ranked school for the student, and the first step of the SPDA algorithm is performed. If any student is rejected in any stage, the procedure contacts her in the next stage, and offers her all schools from which she has not been rejected already.[26] As before,

---

[26]This offer is to align this mechanism with DirDA. A quicker (and outcome-equivalent under SPDA) version instead offers

each contacted student selects one of these schools, which is interpreted as her 'next-best' school. The next step of the SPDA algorithm is consequently performed, and the procedure continues until a unique feasible matching is determined (i.e., when the SPDA algorithm no longer performs rejections). ∎

3.2.2 EXPERIENTIAL INFORMATION

Our elicitation procedure operates under the closed-room assumption: students take part alone in a closed room, with no way to observe other students or even the stages of the procedure in which they are not contacted. This is again a 'worst-case' scenario; in some situations it might be possible to observe the actions of other students or the stages of the procedure, which would convey more information. However, even this limited participation might convey useful information. As an example, suppose the rules of g are that a contacted student is to be offered all schools which are possible for her at that stage. If a contacted student observes an offer $Y \subsetneq X$, she can immediately infer that schools in $X \setminus Y$ *not offered* to her are *not possible* for her under the rules of g. Formally, for a student $i \in N$, a mechanism g, and a stage k where i has been contacted, the ***experiential information*** for i under g at k is denoted $\epsilon_i^{g,k} \in \Phi$. The closed-room assumption implies that no experiential information is acquired in stages in which a student is not contacted. The cumulative experiential information for i under g is $\epsilon_i^g \equiv \epsilon_i^{g,\{K\}}$.[27] A definition for verifiability through experience follows:

DEFINITION **2.** A non-wasteful mechanism g^* is verifiable by experience if, for outcome (ϕ^{g^*}, a^*) and each student $i \in N$: $Poss_i^{g^*}(\phi_i^{g^*} \cap \epsilon_i^{g^*}) = a_i^*$.

Experience is a weaker channel of feedback than the designer's communication, which can directly provide information on other's actions and scores. To achieve verifiability through this source of feedback, therefore, we require some additional structure on the elicitation procedure. We start with a stronger notion of possibility. We say a school is in the 'core possibility' set for a student (given some information and a mechanism) if she is matched by the mechanism to that school at any compatible problem in which she ranks that school as her most-preferred. Formally, given information ϕ, a school $x \in X$ is core possible for student $i \in N$ at ϕ if for any $P = (\succeq, v^*) \in \phi$ such that $x \succ_i y$ for all $y \in X$, we have that $g_i(P) = x$.[28] The core possibility set for i under g at ϕ is denoted $Core_i^g(\phi)$. Then:

1. A mechanism g is descending if, at any stage of the elicitation procedure, any contacted student

---

her only those schools which are still possible for her.

[27] Students acquire experiential information even in one-stage elicitation procedures, though of course this information is richer and possibly more salient in multi-stage procedures—an observation we shall exploit.

[28] In effect, student i can guarantee for herself school x by reporting preferences in which x is top-ranked. A related idea of guaranteeing—or 'clinching'—a school appears in Morrill (2015) for the TTC mechanism, and for obviously strategy-proof mechanisms in Pycia and Troyan (2019).

is offered all schools that are possible for her under the designer's information at that stage, i.e., $Poss_i^g(\phi^{g,\{k-1\}}) = Y_i^{g,k}$ for any stage k in g and any contacted student $i \in Act^{g,k}$.

2. A mechanism g is <u>ascending</u> if, at any stage of the elicitation procedure, any contacted student is offered all schools that are core possible for her under the designer's information up to and including that stage, i.e., $\bigcup_{l<k} Core_i^g(\phi^{g,l}) = Y_i^{g,k}$ for any stage k in g and any contacted student $i \in Act^{g,k}$.

3. A mechanism is <u>structured-offers</u> if it is either ascending or descending.

We note that all non-wasteful mechanisms with a one-stage elicitation procedure are also trivially structured-offers, since each student is offered all schools.

3.2.3 PREDICTABLE ELICITATION PROCEDURES

In general, among structured-offers mechanisms, verifiability through experience can be achieved if permitted actions in the elicitation procedure involve ranking essentially one school at a time. We call such elicitation procedures 'predictable,' in the sense that the experiential information is rich enough for each student to pinpoint precisely at every stage which school she will be matched to if the procedure were to terminate at that point. Formally, the elicitation procedure for mechanism g is <u>predictable</u> if:

P1: Possibility sets under experiential information are always single-valued: For any stage k and any $i \in Act^{g,k}$, $|Poss_i^g(\phi_i^{g,\{k\}} \cap \epsilon_i^{g,\{k\}})| = 1$; and

P2: Each student is matched to the unique possible school at the terminal stage: If K^g is the terminal stage of the elicitation procedure of g, and the outcome is (ϕ^g, a), then $Poss_i^g(\phi_i^{g,\{K^g\}} \cap \epsilon_i^{g,\{K^g\}}) = a_i$.

Predictable elicitation procedures allow students to rank multiple schools in a stage, provided only one of them is possible for the student under her induced experiential information. A useful and intuitive class of predictable procedures involve students ranking *exactly* one school at a time. Formally, a mechanism g is a <u>single-school mechanism (SSM)</u> if, for outcome (ϕ^g, a), each stage k in g and each contacted student $i \in Act^{g,k}$: (1) permitted actions involve only one school: $|\hat{Y}_i^{g,k}| = 1$ and (2) the last school in her actions is her match: if $i \notin Act^{g,l}$ for any $l > k$ then $\hat{Y}_i^{g,k} = a_i$. It is straightforward to show that SSMs are predictable.

PROPOSITION **1.** *A non-wasteful and structured-offers mechanism g^* is predictable if it is an SSM.*

The proof is in Appendix A. More generally, predictability of a structured-offers mechanism is necessary and sufficient to make the promised environment verifiable through experience.

THEOREM **2.** *Let g^* be a non-wasteful and structured-offers mechanism, and let (ϕ^{g^*}, a^*) be the outcome in g^*. Then g^* is verifiable through experience (i.e., $Poss_i^{g^*}(\phi_i^{g^*} \cap \epsilon_i^{g^*}) = a_i^*$ for each $i \in N$) if and only if the elicitation procedure for g^* is predictable.*

The proof is in Appendix A. It is straightforward from the definitions that predictable elicitation procedures for non-wasteful and structured-offers mechanisms induce verifiability through experience. In the other direction, we show separately for ascending and descending mechanisms that a mechanism that is verifiable through experience necessarily provides experiential information through its elicitation procedure that induces single-valued possibility sets at each stage of the mechanism, satisfying predictability. Proposition 1 and Theorem 2 have profound practical implications for achieving verifiability in school admissions, as SSMs are easy to design (see Example 3 below), and can usually be just as easily implemented, for instance via Internet platforms, like in Germany or France.

EXAMPLE **3.** Consider the SeqDA mechanism introduced in Example 2. Firstly, it is easy to see that SeqDA is an SSM, as each student selects only one school at a time, and her last-chosen school is her final match. To see how SeqDA is predictable, and thus induces verifiability through experience, recall the problem in Example 1. In the first stage of SeqDA, school x is the only school possible for Student 1 under her experiential information, since it is the only school in her actions. In the second stage, her experiential information is updated such that neither School x (which she is not offered) nor School y (which she has not yet ranked) are possible for her. School z is thus the only school possible for her under her experiential information, and this is also her final match. A similar reasoning can be applied to Students 2 and 3. ∎

4   THE TRANSPARENCY PROBLEM

A *(matching) environment* is denoted $\mathcal{E} = (\mathcal{P}, g, M)$, where $\mathcal{P}$ is a domain of matching problems, g is a mechanism, and M is a communication protocol. Since $\mathcal{P}$ is fixed, we will henceforth refer to an environment as simply $\mathcal{E} = (g, M)$. We can extend our verifiability definitions to environments quite naturally:

DEFINITION **3.** Let $\mathcal{E}^* = (g^*, M^*)$ be the promised environment, and let (ϕ^{g^*}, a^*) be the outcome. We say $\mathcal{E}^*$ is verifiable by experience if g^* is verifiable by experience, and is verifiable by communication if g^* is verifiable by the communication protocol M^*. We say that $\mathcal{E}^*$ is verifiable if it is either verifiable by experience or communication (or both).

Our proposed solutions to the verifiability problem rely on students assuming that the designer has indeed used the promised environment $\mathcal{E}^*$ to compute the matching, i.e., using the promised mechanism g^* to realise the outcome (ϕ^{g^*}, a^*) and communicating information $M^*(\phi^{g^*})$ according to the promised

protocol M^*, and simply requiring information enough to verify their matches. In this section, we drop this assumption, in effect allowing the possibility of students who also doubt whether the promised environment was used or not. In particular, what if a student believes that her match is $a_i \neq a_i^*$, where $a = (a_i)_{i \in N}$ is the result of the designer using an environment $\mathcal{E} \neq \mathcal{E}^*$? We call this the Transparency Problem.

A few notes before we continue our analysis. Firstly, we are agnostic as to whether the designer actually deviates in this way; instead, we seek to identify conditions on the promised environment $\mathcal{E}^*$ that prevent *any* deviation from a^*, no matter the reason. Tying the designer's hands in this way is a useful exercise even if the designer is completely honest, as it amounts to providing students comfort that the designer could not deviate from the outcome even if she tried. Secondly, we do not consider all imaginable deviations by the designer. In particular, possible deviations involving manipulating student scores or school capacities are beyond the scope of this analysis; we treat these 'inputs' to the matching problem as fixed and immutable. Thirdly, we impose non-wastefulness even on any possible deviation of the designer, following on from our assumption that empty seats in schools are costly for the designer and also observable by students. We acknowledge that these are strong restrictions: a truly dishonest designer might indeed choose these kinds of manipulations. However, not only can such deviations only be prevented by requiring full disclosure, but our admittedly limited definition of the Transparency Problem is nevertheless useful as it can assuage students that the correct output was produced relative to the (fixed) inputs and the promised environment.

Back to our analysis. For a matching $a \neq a^*$ to be produced, the designer must have used an environment $\mathcal{E} = (g, M)$ with outcome (ϕ^g, a). Moreover, for no student to have immediately detected this change, $\mathcal{E}$ would have to be 'plausible' in both the mechanism and communication dimension. A mechanism g is indistinguishable from the promised mechanism g^* if (1) g is non-wasteful, and (2) offer sets observed by any student in g 'look similar' to offer sets they might observe in g^*.[29] On the other hand, a communication protocol M is indistinguishable from the promised protocol M^* if (1) $M(\phi^g)$ is of the same type as $M^*(\phi^{g^*})$, and (2) g is verifiable by M if g^* is verifiable by M^*. An environment $\mathcal{E} = (g, M)$ is plausible if g and M are indistinguishable from g^* and M^*, respectively.[30]

Suppose a student $i \in N$ believes the designer has run the environment $\mathcal{E}$. Let $\phi_i^{\mathcal{E}}$ be the total information (intrinsic plus feedback) acquired by a student in $\mathcal{E}$. The possibility set induced for student

---

[29]Notice that if g^* is ascending, then offer sets weakly increase in size, while they weakly decrease in size if g^* is descending. Formally, the second condition translates to the requirement that, for any student $i \in N$, and any two stages $k < l$ of the elicitation procedure of g: (1) $Y^{g,l} \subseteq Y^{g,k}$ if g^* is descending; and (2) $Y^{g,l} \supseteq Y^{g,k}$ if g^* is ascending.

[30]From a strategic point of view, we assume that students use the same type in indistinguishable mechanisms. That is, if g is indistinguishable from g^*, then the pre-committed type $\succeq_i^g$ for i in g is the same as the pre-committed type $\succeq_i^{g^*}$ for i in g^*. As before, we ignore any relationship between the used type and any true type.

$i \in N$ is given by $Poss_i^{g^*}(\phi_i^{\mathcal{E}})$.[31] 'Transparency' of the promised environment then essentially means that *only* the correct match a_i^* is possible for $i \in N$, no matter which plausible environment the designer uses.

DEFINITION 4. The promised environment $\mathcal{E}^*$ is <u>transparent</u> if, for the promised outcome (ϕ^{g^*}, a^*), for each plausible environment $\mathcal{E}$ and each student $i \in N$: $Poss_i^{g^*}(\phi_i^{\mathcal{E}}) = a_i^*$.

4.1 MAIN RESULT

We first demonstrate how the solutions that we have identified for the Verifiability Problem—mechanisms with predictable elicitation procedures and cutoff protocols—fail to achieve transparency.

EXAMPLE 4. Consider a problem as given in Figure 5. The matching prescribed by the SPDA rule for this problem is marked in boxes, as are the cutoffs corresponding to this matching. A false matching is underlined, as are the cutoffs corresponding to the false matching.[32]

Figure 5: The SPDA matching and a false matching

Actions			Scores			
$\geq_1$	$\geq_2$	$\geq_3$	Student	v^x	v^y	v^z
x	$\boxed{x}$	$\boxed{z}$	1	40	$\boxed{\underline{75}}$	85
$\underline{y}$	$\underline{z}$	$\underline{x}$	2	$\boxed{80}$	45	$\underline{95}$
z	y	y	3	$\underline{90}$	60	$\boxed{\underline{50}}$

Suppose the designer promises an environment $\mathcal{E}^* = (g^*, M^*)$ where g^* is SeqDA, and M^* is the 'null' protocol. Consider instead a mechanism g with the same first stage as g^*: all students are contacted and offered all schools. In this first stage, both Student 1 and Student 2 select School x. According to g^*, Student 1 should be rejected, but suppose instead that Student 2 is rejected in g. Student 2 has no way to tell from her experience alone that this rejection is illegitimate. In the second stage of g, Student 2 is contacted and offered School z and School y, and she selects School z. In every subsequent stage, g rejects students as g^* would. Eventually, g terminates with the false matching underlined in Figure 5. Thus predictability of the promised mechanism is by itself insufficient to make the environment transparent.

Interestingly, cutoffs are also not sufficient to make an environment transparent by themselves. Suppose the designer promised to use an environment $\mathcal{E}^* = (g^*, M^*)$ where M^* is a cutoff protocol, and g^* is the DirDA mechanism. If the designer were to select the false matching underlined in Figure 5, and com-

---

[31]Notice that the information $\phi_i^{\mathcal{E}}$ comes from the environment $\mathcal{E}$ used by the designer, but possibility sets are generated under the promised mechanism g^*.

[32]The knowledgeable reader would identify the false matching as the outcome of using the college-proposing DA rule for this problem.

municate the cutoffs $(90, 75, 95)$ corresponding to this matching, each student's match would nevertheless be justified by this communication, and the deviation cannot be detected. ∎

In one sense, transparency of the promised environment can be achieved quite easily if we insist on full disclosure by the designer, as she cannot misrepresent scores or actions without being detected. We prove this formally in Theorem 3. But we also show that transparency can often be achieved even without full disclosure, and sometimes with surprisingly little feedback for students. In particular, for some environments, all it requires is verifiability through *both* communication *and* experience.

THEOREM **3.** *The promised environment $\mathcal{E}^* = (g^*, \mathcal{M}^*)$ is transparent if:*

1. *M* is full disclosure; or*

2. *g^* is a non-wasteful and structured-offers mechanism, and $\mathcal{E}^*$ is verifiable through both experience and communication.*

Showing that full disclosure achieves transparency is easy, as it cannot be falsified without detection. The main intuition behind the rest of the proof is that the used mechanism g needs to be multi-stage (as a consequence of predictability of g^*), which limits the designer in significant ways. In deciding to deviate from the promised mechanism at any stage, she has to work with all possibilities of continuation actions of students, since she does not know them in advance. This is in contrast to a direct mechanism, where she acquires all actions in one stage, and has more degrees of freedom to manipulate the matching. In particular, we show that there is always the risk that the deviation might either lead to the final matching being wasteful, or a communication of information that does not justify the matching at the false outcome.

It should be pointed out that these are sufficient conditions for transparency; it might be possible to make environments transparent, at least in some applications, by providing students with *even less* information than required for verifiability through communication and experience. As such, the nature of the *minimum* information required by students for transparency is an open question. Nevertheless, we believe our sufficiency result is useful, because of the following corollary with practical implications:

COROLLARY **1.** *A promised environment with a non-wasteful, structured-offers, single-school mechanism, and a cutoff protocol is transparent.*

5   ADDITIONAL RESULTS

We have shown in Theorem 3 that a predictable mechanism can form an important part of a transparent environment. Such a mechanism is necessarily multi-stage. But in many applications, it might not be feasible for the designer to use a multi-stage mechanism: for instance, if there are time constraints on

producing the matching, or if it is costly to set up multiple interactions for students. In this section, we propose a stronger communication protocol that can be used even with a one-stage mechanism. This protocol achieves a stronger notion of verifiability (though we show that it falls short of transparency).

To motivate this notion, recall Example 4, in which we showed how a false matching under DirDA could be justified by providing the corresponding cutoffs, even without manipulating them in any way. In particular, the cutoff protocol is not rich enough in informational terms to avoid justifying a false matching. We formalise this idea. Formally, let $\mathcal{E} = (g^*, M^*)$ be the promised environment, and let (a^*, ϕ^{g^*}) be the outcome. Let $M^*(a^*, \phi^{g^*})$ be the communication prescribed by the protocol at this outcome. We say that a matching $a \in \mathcal{A}$ is justifiable by M^* if a_i is justified by $M^*(a, \phi^{g^*})$ for each $i \in N$, where $M^*(a, \phi^{g^*})$ is the communication prescribed by M^* at outcome (a, ϕ^{g^*}). Strong verifiability requires that only the correct matching a^* be justifiable by the communication protocol.

DEFINITION **5.** Let $\mathcal{E}^* = (g^*, M^*)$ be the promised environment, and let (a^*, ϕ^{g^*}) be the outcome. $\mathcal{E}^*$ is strongly verifiable through communication if a^* is the unique matching that is justifiable by M^* in $\mathcal{E}^*$.

An environment strongly verifiable through communication is also verifiable through communication, though the converse is not true. Moreover, as Example 4 suggests, a cutoff protocol is not enough for strong verifiability. To arrive at a communication protocol that does achieve strong verifiability, we need a few additional concepts. For a mechanism g, define a single-school equivalent (SSE) mechanism g^S such that (1) g^S is an SSM and (2) $g(P) = g^S(P)$ for each $P \in \mathcal{P}$. Let the designer's information at any stage k in g^S be $\phi^{g^S,\{k\}}$. Since g^S is an SSM, we can define a temporary matching at k given by $a^{g^S,k} \in X^N$, comprising the last-chosen school for each student (the empty school if none exists). Let $c^g(a^{g^S,k}, \phi^{g^S,\{k\}})$ be the cutoffs at k for this temporary matching. Then, the step-cutoffs for g are given by the collection $(c^{g^S}(a^{g^S,k}, \phi^{g^S,\{k\}}))_{k \leq K^{g^S}}$ of cutoffs at every stage of the SSE g^S. A communication protocol M is a step-cutoff protocol for g if it communicates the collection of step-cutoffs generated by the SSE for g, i.e., $M(\phi^g) = (c^{g^S}(a^{g^S,k}, \phi^{g^S,\{k\}}))_{k \leq K^{g^S}}$ where g^S is the SSE for g. Then:

PROPOSITION **2.** *A promised environment $\mathcal{E}^* = (g^*, \mathcal{M}^*)$, where g^* is non-wasteful and the SSE g^{*S} is structured-offers, is strongly verifiable through communication if $\mathcal{M}^*$ is a step-cutoff protocol for g^*.*

The proof is in Appendix A. Essentially, step-cutoffs allow students to infer more detailed information about the actions and scores of other students, helping them to rule out the possibility of another justified match. However, step-cutoffs are not sufficient for transparency, as we demonstrate below:

EXAMPLE **5.** Consider the problem in the table on the left of Figure 6. The matching generated by the DirDA mechanism is given in boxes, and the centre table of Figure 6 gives the step-cutoffs generated by SeqDA, which is the SSE for DirDA. However, we show that the designer can manipulate this outcome

without detection, by manipulating the step-cutoffs. In particular, let g be an SSM that differs from SeqDA in that Students 2 and 3 are rejected in the first step. The designer manipulates the cutoffs for School y and School z at this step by raising them high enough to justify the rejections of both Students 2 and 3 (see the table on the right of Figure 6). In subsequent steps, g operates as if it were SeqDA. Thus in stage 2, Student 2 ranks z and Student 3 ranks x in her actions, respectively. This causes Student 1 to be rejected from x, and in stage 3 she ranks y, which leads to the underlined matching. This matching is justified by the false step-cutoffs given in the table on the right, and moreover no student can individually detect the deviation. Therefore, this environment is not transparent. ∎

Figure 6: Manipulating step-cutoffs

A problem

Actions			Scores			
$\geq_1$	$\geq_2$	$\geq_3$	Student	v^x	v^y	v^z
x	y	z	1	40	70	90
$\underline{y}$	$\underline{z}$	$\underline{x}$	2	80	45	65
z	x	y	3	55	80	50

True step-cutoffs

Step	School		
	x	y	z
1:	40	45	50

False step-cutoffs

Step	School		
	x	y	z
1:	40	60	52
2:	70	60	55
3:	70	65	55

6 Verifiability: A Laboratory Experiment

In this section, we present a laboratory experiment based on the SPDA rule in school admissions. Recall that our theory assumes students can rationally, fully, and accurately process feedback and refine their possibility sets. As our theoretical solutions rest on this ability, it is essential to know if it is possible (and, consequently, if our solutions work). Our experiment is designed as a proof of concept for the usefulness of feedback for verifiability, and not only to test our theoretical results.

In our experimental design, we opt for honesty of the designer; the promised mechanism is always run. While a random draw might determine the match, we ensure that verifiability is enough to spot random matches. We test separately for the effect of experiential information and communicated information on 'observed verifiability.' Experiential information is generated by the use of a predictable mechanism, SeqDA, which theoretically ensures verifiability. The complexity of decision-making arises from the need to understand the rules of the promised mechanism and infer the possibility set from the sequence of observed offers and taken actions. The information communicated is via step-cutoffs, which theoretically ensures strong verifiability. In this case, the complexity of decision-making comes from the need to interpret cutoffs at every step of the algorithmic procedure and correspondingly infer the possibility set. We also use a

theoretically transparent environment to provide both sources of feedback.[33] Our experimental setting is a novel empirical tool to compare the relative verifiability of an environment. Thus, it can be used to test the observed verifiability of newly proposed environments before putting them into practice. In this sense, it is of independent interest.

6.1 Treatments

Consider the following four environments involving the SPDA, which correspond to treatments in the experiment. These environments contain two different mechanisms and two different communication protocols. In the mechanism dimension, we use either a 'direct' mechanism, where participants submit rank-order lists of schools, or a 'single-school' mechanism, where participants apply to schools one by one (Echenique, Wilson, and Yariv, 2016; Bó and Hakimov, 2020a; Klijn, Pais, and Vorsatz, 2019). The other dimension is on the communication protocol. The first 'null' protocol provides no feedback. The other is a step-cutoff protocol where we provide the cutoff grades of each school at each step of the underlying DA algorithm.[34] The four environments described below correspond to our four treatments in the experiment.

Treatment 1: The direct SPDA mechanism with no feedback (DirNo)

Every student submits a rank-order list of schools. Schools' capacities and strict priorities over students are given exogenously. The mechanism collects students' submitted rank-order lists of schools simultaneously and in one stage, which are used by the algorithm below:

- Step 1: An application is sent for each student to her most-preferred school. Each school rejects the lowest-score students that are in excess of its capacity, and temporarily holds the others.

- Step $k > 1$: An application is sent for every student rejected in step $k - 1$ to the next school in her submitted rank-order list. Each school pools together new applicants and those who are held from step $k - 1$ and rejects the lowest-score students in excess of its capacity. Those who are not rejected are temporarily held.

The algorithm terminates when there are no rejections. Each school is then matched to the students

---

[33]Note that transparency implies verifiability. We perceive the verifiability of the environment as the primary behavioral question, as transparency relates to the designer's ability to manipulate the allocation, not to participants' ability to distinguish correct and wrong assignments.

[34]Theoretically, we could have used terminal-cutoffs. However, our design involves unit capacity for schools, and so terminal-cutoffs would make detecting violations either trivial (if cutoffs were reported truthfully) or impossible (if cutoffs were adjusted to the random matching). This is further discussed when we provide the exact description of step-cutoffs in the case of random matchings.

it holds, and students who are not held at any school are left unmatched. This is the unique realised matching. This environment is not verifiable through experience or communication.

Treatment 2: The direct SPDA mechanism with step-cutoffs (DirCutoffs)

Every student submits a rank-order list of schools, and the algorithm as in DirNo above is run. After the matching is determined, students also receive the table of cutoff grades for each school corresponding to each step of the algorithm (step-cutoffs). This environment is strongly verifiable through communication (Proposition 2). Thus, this treatment tests the ability of participants to use communicated step-cutoffs to refine their possibility sets under SPDA.

Treatment 3: The sequential SPDA mechanism with no feedback (SeqNo)

Each student applies to schools one at a time. After each step, a student learns through the rules of the mechanism whether she was temporarily accepted or rejected by the school to which she applied at that step. More specifically:

- Step 1: Each student applies to one school. Each school rejects the least-ranked students according to its priority from among those who applied to it, in excess of its capacity, and temporarily holds the others. If no application is rejected by any school, the procedure will stop at this step, matching the schools to the students they hold.

- Step $t > 1$: Each student who is not held at some school applies to any school that has not rejected her previously. If all schools have already rejected her, she is no longer asked to make choices. Each school rejects the least-ranked students according to its priority among those held and those who applied to it, in excess of its capacity, and temporarily holds the others. If no application is rejected, the procedure stops at this step, matching the schools to the students they hold and leaving students who are not held at any school unmatched.[35]

This environment is predictable (Proposition 1) and thus verifiable through experience. This treatment tests the ability of participants to use experiential information to refine their possibility sets under SPDA.

Treatment 4: The sequential SPDA mechanism with step-cutoffs (SeqCutoffs)

Every student participates in SeqNo, and cutoff grades of each school are revealed at each step. After the procedure terminates, students also receive the table of step-cutoff grades for each school. This environment is verifiable through both experience and communication, and thus is also transparent.[36]

---

[35]Subjects in the experiment had to make a decision at each step, with no option of not making a choice.

[36]Recall that terminal-cutoffs are enough for Transparency when using Sequential DA. However, we wished to preserve the 2x2 feature of the design, and thus chose step-cutoffs for this treatment.

6.2 Experimental Design

In the experiment, there were six schools with one seat each and six competing students. Five students were computer-simulated ('sims'), while the sixth student was a human subject. We use sims because we are interested in the individual decision-making of participants with respect to transparency, and it allows us to form relatively large markets and increase the number of independent observations. Together, sims and subjects are called 'students' in what follows. Sims were programmed to submit truthful rank-order lists in DirNo and DirCutoffs, which are weakly dominant strategies under Direct DA (Roth, 1982). In SeqNo and SeqCutoffs, at each step, they were programmed to apply to the best school among those that did not previously reject them. This strategy is an ordinal perfect Bayesian equilibrium under sequential DA (Bó and Hakimov, 2018). Subjects were informed that sims follow the strategy that maximises their payoff.[37] In all treatments, subjects received 31 Swiss Francs (CHF) if they were matched to their most-preferred school, 26 CHF for the second most-preferred school, 21 CHF for the third most-preferred school, and so on in decrements of 5 CHF, receiving 6 CHF for the least-preferred school.

Subjects were informed that in 50% of cases the match would be determined not by the explained SPDA procedure but instead randomly. This probability was common across all treatments. After the match was determined, subjects were asked to make a decision on whether or not to appeal.[38] Payoffs were constructed to make appealing optimal if the match was determined randomly. If a subject appealed and the appeal was incorrect (the match was determined by the explained procedure), she incurred a cost of 6 CHF, which was deducted from her match payoff.[39] If a subject appealed and the appeal was correct (the match was determined randomly), her total payoff for the round was 40 CHF, replacing whatever match payoff she might have otherwise received. If a subject decided not to appeal, she received her match payoff. Each session contained 10 rounds of the school admissions game. All subjects had to submit an appeal decision in each round, and each round represented a new market. The preferences for each student in each

---

[37]SeqNo and SeqCutoffs truthful behavior maximizes payoffs of computer players under the assumption that human players do not condition the choice from the menus on timing of the procedure or specific values of cutoffs. This is a reasonable assumption, as a profitable deviation from truthful behavior would require computer players to know who and exactly how the human players condition their choices on the timing of the mechanism.

[38]One could argue that a random match might be the easiest to spot for subjects, as the designer would try to tailor deviations in a more sophisticated manner. This would bias our results towards 'higher verifiability.' However, the fact that given a random draw students might receive a better university than the true match might preclude participants from appealing, thus biasing our results towards 'lower verifiability.' Importantly, the effects are the same between treatments, and should not affect the validity of treatment comparisons.

[39]We opted for the additive cost of an appeal to approximate potentially different incentives to appeal for those who receive, justifiably or not, a better or a worse placement. Notably, the equilibrium and random assignments are the same between treatments, which balances the expected incentives to appeal.

market were generated using the designed-markets principle (Chen and Sönmez, 2006). For each market, we generated the qualities of schools uniformly and randomly between [0,40], corresponding to a common utility for each student from being matched to each school. Additionally, for each student and school, we generated a random component of utility from the interval [0,20]. The resulting total utilities were transformed into ordinal preferences. The procedure above ensured some correlation between preferences. The grades were independently drawn in each round from the uniform distribution with support [1,100] for math and languages.[40]

The decision to appeal is a costly decision with probabilistic payoff, and thus might be affected by the risk preferences of the participants. In order to control for a potentially different distribution of risk attitudes between treatments we measure subjects' risk-aversion. After 10 rounds of the school admissions game, subjects participated in a risk-aversion measurement task similar to 'multiple price lists' (Holt and Laury, 2002).[41] At the end of the experiment, one round of either the school admissions game or the results of the risk-aversion task was randomly drawn to determine payoffs.

In the experiment, subjects could see tables with all ordinal preferences. Complete information on preferences allowed subjects to infer the popularity of each school, resembling real-life conditions. They were also told the distribution of exam grades, but could see the realisation for only their own grades. We believe that this too approximates real-life informational conditions. After each round, subjects received feedback only about their own match. Additionally, in the DirCutoffs and SeqCutoffs treatments, they observed the table of step-cutoff grades of all schools. After submitting an appeal decision, they learned whether their decision was correct, as well as their final payoff for the round.

The random match for each round was pre-generated and fixed for all subjects. It never coincided with the student-optimal stable matching. We used the same random allocations for a particular round in all treatments. We ran two sessions for each treatment. For each session, we also predetermined the rounds in which the match was determined randomly, and these rounds were the same in all treatments. In order to make sure that allocations were determined randomly precisely 50% of the time in our data, we ran the two sessions for each treatment as follows: if the match was determined randomly in a round in one session (a 'random round'), then the match was determined through the explained procedure for this round (a 'determined round') in the other session for this treatment. Thus, for each market, each round was a random round in one session and a determined round in the other. In particular, in the first session of each treatment, the random rounds were 3, 7, 8, and 10, but were rounds 1, 2, 4, 5, 6, and 9 in

---

[40]The details of all markets are presented in Appendix C.

[41]Note that payoffs were designed to make appealing optimal if the match was random, independently of risk preferences. However, if subjects are not sure whether a violation took place, risk preferences are important in determining their choice of whether or not to appeal. Thus, we measure subjects' risk-aversion to control for their influence on appeal decisions.

the second session. The order of treatments was randomised between sessions 1 to 4 and 5 to 8.

In random rounds, we modified the cutoffs in the final step for random rounds. The final cutoff of the random match showed the student's grade (to correspond with the match), and the cutoff was randomly generated at the school corresponding to the correct match.[42] Participants were informed about this modification in the instructions.

In the risk-aversion task, each subject had to choose between two options: either a certain amount of CHF or a lottery between two amounts. One row was randomly chosen, and the payoffs for the task were according to the subjects' choice: they received either a certain amount or the outcome of the chosen lottery. The choices for the risk-aversion task are presented in Table 1.

Table 1: Choices in Round 11

	Option A	Option B
Row 1:	100% of CHF36	50% of CHF40 and 50% of CHF30
Row 2:	100% of CHF34	50% of CHF40 and 50% of CHF28
Row 3:	100% of CHF31	50% of CHF40 and 50% of CHF25
Row 4:	100% of CHF26	50% of CHF40 and 50% of CHF20
Row 5:	100% of CHF21	50% of CHF40 and 50% of CHF15
Row 6:	100% of CHF16	50% of CHF40 and 50% of CHF10
Row 7:	100% of CHF11	50% of CHF40 and 50% of CHF5
Row 8:	100% of CHF6	50% of CHF40 and 50% of CHF0

We designed choices in Rows 3 to 8 to correspond to the implicit tradeoffs faced by subjects in their decision to appeal or not, under the belief that the match was random with a 50% probability. For example, Row 3 presents the tradeoff for the case when the subject is allocated to their most-preferred school. If a subject chooses not to appeal, she receives the assured payoff of 31 CHF. However, if she chooses to appeal, her appeal is correct with 50% probability (with a 40 CHF payoff) and incorrect with 50% probability (payoff 25 CHF). Row 4 corresponds to the choice for being matched to the second most-preferred school, down to Row 8, which corresponds to a match to the least-preferred school.

We ran four treatments between subjects: DirNo, DirCutoffs, SeqNo, and SeqCutoffs. Thus, each subject participated in only one session, under only one environment. The experiment was run at the experimental economics lab at LABEX at the University of Lausanne. We recruited subjects from our pool with the help of ORSEE (Greiner, 2003). The experiments were programmed in z-Tree (Fischbacher, 2007). Each session consisted of either 20, 22, or 24 participants.[43] Since subjects played against sims

---

[42]That is why we opted for step-cutoffs and not terminal-cutoffs, as it allowed participants to spot violations through the communication, despite the modification of final cutoffs.

[43]We invited 26 subjects to every session in order to have 24 subjects per session, but some sessions had a high rate of

with no interaction between subjects, we treat each subject as an independent observation. In total, we conducted eight sessions with 186 subjects. With two sessions per treatment, we have 46 subjects in DirNo, 48 subjects in DirCutoffs and SeqNo, and 44 subjects in SeqCutoffs.[44] On average, the experiment lasted 70 minutes, and the average earnings per subject was 38 CHF, including a show-up fee of 10 CHF.

At the beginning of the experiment, subjects were given printed instructions (see Appendix C). Subjects were informed that the experiment was about the study of decision-making. The instructions were identical for all subjects in a treatment, explaining the experimental setting in detail. First, the mechanism and the game of appeals were explained. They were then also provided a detailed example of the procedure, including tables of cutoffs where applicable. Questions were answered in private. The instructions for Round 11 (the risk-aversion task) were provided at the beginning of the experiment. They were also repeated on the screen after Round 10.

6.3 EXPERIMENTAL RESULTS

The significance level of all our results is 5%, unless otherwise stated.

6.3.1 CORRECT APPEALS DECISIONS

Our main focus is the comparison across treatments of the proportion of correct decisions of whether to appeal or not. Note that a correct appeal decision involves two components: *appeal* in random rounds and *not appeal* in determined rounds. We argue that the proportion of correct appeal decisions[45] measures the degree of observed (empirical) verifiability of the environment, as only in a verifiable environment it is possible for subjects to make a correct decision on whether or not to appeal. Our main observation is given below and is also illustrated in Figure 7:

The proportion of correct appeal decisions is the highest under SeqCutoffs (83%), with the difference being significant relative to all other treatments. It is the lowest under DirNo (52%), with the difference being significant relative to all other treatments. There is no significant difference between the proportion of correct appeal decisions under DirCutoffs (65%) and SeqNo (64%).

The proportion of 52% correct appeal decisions in DirNo corresponds effectively to a random chance of a correct decision. This evidence supports our hypothesis regarding the lack of transparency of the direct DA mechanism with no feedback. Once step-cutoff grades are provided to subjects, the proportion

---

no-shows.

[44]Our number of participants, with 10 decisions per subject, allows us to identify treatment differences of above 10.5 percentage points in the proportion of correct appeal decisions with the statistical power of at least 80% for a 5% significance level.

[45]That is, the number of rounds with a correct appeal decision, divided by the total number of rounds.

Notes: Vertical gray bars represent the 95% confidence intervals. DirNo is a treatment with direct DA and no cutoffs. SeqNo uses sequential DA with no cutoffs. DirCutoffs and SeqCutoffs use direct and sequential DA, respectively, and provide step cutoffs.

Figure 7: Proportion of correct appeal decisions

increases significantly to 65%. This difference suggests that subjects are at least partially able to use step-cutoffs to judge their match's legitimacy. A similar tendency is observed for the SeqNo treatment. The proportion of correct appeal decisions is 64% on average, which is significantly higher than under DirNo, but not significantly different from DirCutoffs. This is the effect of switching the mechanism, and thus experiential information. Again, in line with our theory, subjects can use their experience in a sequential mechanism to judge their matches' legitimacy.

Our theoretical predictions find support in terms of treatment differences, but not the levels. Note that the verifiability of SeqNo and DirCutoffs should lead to 100% correct appeal decisions.[46] However, we observe a high rate of mistakes in DirCutoff that might be driven by the complexity of processing cutoff information, given that cutoffs even in random rounds are made such that the subject's grade always corresponds to the announced match. Similar mistake rates are also observed in SeqNo, suggesting that many subjects did not understand the relation between their matches and their choices in the steps of the mechanism. This interpretation is also suggested in Table 2, which presents the proportion of correct appeal decisions split by the first five and last five rounds for each treatment. There is a significant increase in the proportion between the first five and the last five rounds of both the DirCutoff and SeqNo

---

[46]In fact, there were 29 of 1,860 outcomes in the experiment when the random outcome was the same as the match outcome. This might make the identification of manipulation impossible, though not necessarily so. By design, we made sure this would not happen if students played the mechanism optimally, but some participants manipulated reports such that the outcome coincided with the random draw. Given how small the number is, we ignore it in subsequent analysis, using the whole sample, since excluding these observations do not qualitatively change the results.

treatments. This suggests that subjects might need some time to understand how to correctly use this information in their appeal decisions.

Table 2: Proportions of correct appeal decisions

	DirNo	DirCutoffs	SeqNo	SeqCutoffs	1=2	1=3	1=4	2=3	2=4	3=4
	(1)	(2)	(3)	(4)	p-val.	p-val.	p-val.	p-val.	p-val.	p-val.
First half	52%	59%	57%	82%	0.12	0.30	0.00	0.63	0.00	0.00
Last half	53%	71%	72%	84%	0.00	0.00	0.00	0.79	0.00	0.01
Overall	53%	65%	64%	83%	0.00	0.00	0.00	0.89	0.00	0.00
First=last	0.85	0.01	0.00	0.57						

Notes: All the p-values in columns 6 to 11 are p-values for the coefficient of the corresponding treatment dummy in the probit regression of the dummy for the correct appeal decision on the treatment dummy, with the sample restricted to the treatments involved in the test. The p-values in row 5 are p-values for the dummy, which equals 1 for rounds 6 to 10, and 0 for rounds 1 to 5 in a probit regression of the correct appeal decision, with the sample, restricted to one treatment involved in the test. The standard errors of all regressions are clustered at the subjects' level.

The proportion of correct appeal decisions is the highest (on average, 83%) in the transparent SeqCutoffs treatment. The difference is significant relative to the other three treatments. Once subjects receive information from both experience and communication from the designer, the effect is much stronger than under each source alone. For some subjects this might be because it is easier to spot violations through cutoff grades, while for others the experience of the mechanism is more intuitive and guides their decisions. On the other hand, it might be that all subjects receive signals about the correctness of the match from both sources, but signals are followed only when both are present. This result supports using a transparent environment even when the honesty of the designer is not a concern, as it improves observed verifiability of the environment by providing two sources of information to the participants. Note also that there is no significant difference across rounds in SeqCutoffs, with the rate remaining stable at around 83% in the first five and last five rounds.

6.3.2 DETERMINANTS OF CORRECT APPEAL DECISIONS

Next, we look at the determinants of correct appeal decisions. Our main observations are that:

1. *The proportion of correct appeal decisions is higher when the correct decision is to not appeal rather than to appeal. This is true in all treatments.*

2. *The lower the match payoff, the more likely a subject is to appeal.*

3. *Subjects are more likely to make correct appeal decisions when they play truthfully.*

Figure 8 presents the proportion of correct appeal decisions split by random and determined rounds. The left panel presents proportions of correct appeal decisions in determined rounds, i.e. when the correct

decision is to not appeal. Under DirNo, DirCutoffs, and SeqNo, the proportions are 71%, 75%, and 77% of the time respectively, with no significant differences between treatments. Under SeqCutoffs, the proportion is 90%, which is significantly higher than in all other treatments. The right panel presents proportions of correct appeal decisions in random rounds, i.e. when the correct decision is to appeal. In each treatment, the decisions are correct less frequently than in determined rounds. Note that under DirNo, subjects' decisions are less correct than if they were generated by chance. This suggests that subjects are biased towards not appealing. The treatment differences in the proportion of correct appeal decisions between DirNo and DirCutoffs, and between DirNo and SeqNo are driven by the higher appeal rates when appealing is optimal. Thus, spotting procedure violations is simpler for subjects in these treatments relative to DirNo, which is in line with our theoretical predictions, as each violation is identifiable.

Notes: Vertical gray bars represent the 95% confidence intervals. The left figure presents only rounds when the correct match was communicated to subjects, and thus the optimal decision was not to appeal. The right figure presents only rounds when the false match was communicated to subjects, and thus the optimal decision was to appeal. DirNo is a treatment with direct DA and no cutoffs. SeqNo uses sequential DA with no cutoffs. DirCutoffs and SeqCutoffs use direct and sequential DA, respectively, and provide step cutoffs.

Figure 8: Proportion of correct appeal decisions by optimal decision

What explains the aversion to appeals? It is useful to have a benchmark. If there is no further information to suggest whether the match is random or not, appeal decisions should be driven purely by subjects' risk preferences based on a 50% probability of random match. The tradeoff that subjects face in appeal decisions depends on the rank of the school that they received. Thus, for each decision, we can construct an appeal decision based purely on a 50% probability of a successful appeal. For instance, if a subject is matched to his or her most-preferred school, then the decision of whether to appeal based on a 50% probability of success is equivalent to the choice between the sure payoff of 31 CHF (not appealing)

and the lottery with a 50% chance of 40 CHF (correct appeal) and a 50% chance of 25 CHF (incorrect appeal). As discussed, this tradeoff corresponds to the lottery choice in Row 3 in the risk-aversion task of the experiment (Table 1). In particular, if the subject prefers the lottery, this suggests that she should appeal when she believes the match was generated randomly with a 50% probability. Similar choices for lower ranked matches correspond to later rows in Table 1.

Risk preferences suggest that subjects would appeal 73% of the time, with no significant differences between treatments. However, we observe appeals rates of 32%, 40%, 38%, and 42% under DirNo, Dir-Cutoffs, SeqNo, and SeqCutoffs, respectively. Thus, subjects appeal less than the choice in the lottery suggests. Such appeal-aversion cannot be explained by risk preferences alone, and is arguably even stronger if one accounts for this possibility. The appeal-aversion is consistent with two explanations:

- Under all environments, subjects update beliefs in favour of the match being generated by the mechanism, and not randomly. This might be driven by confirmation bias when subjects try to rationalise the match through the prism of the mechanism.

- Appeal-aversion is explained by the loss aversion of subjects with the match payoff as a reference point. Under this interpretation, loss aversion is not present in the choice between lotteries in the risk-aversion task, as it does not generate a reference point (as every row provides a different payoff alternative).[47] If appeal decisions are based on such a reference point, subjects will be less likely to appeal than in the case of the lottery with the equivalent tradeoff in payoffs.

So what are the determinants of correct appeal decisions? Table 3 presents the marginal effects of the probit regression of the dummy for the correct appeal decision on treatment dummies and controls. Models (1) and (2) present regression results for the full sample of decisions. Models (3) and (4) present results only for determined rounds, where the optimal decision is to not appeal. Models (5) and (6) present the results for random rounds, where the optimal decision is to appeal.

First, the significance of treatment dummies for DirCutoffs, SeqNo, and SeqCutoffs does not change when adding controls, which confirms the robustness of our main result. Second, several controls have a significant correlation with correct appeal decisions. On average, the higher the preference rank of the match, the less likely participants are to appeal correctly (the coefficient for the rank of the match is negative and significant in Model (2)). Moreover, there is a negative effect in determined rounds and a positive effect in random rounds—the worse the match, the more likely the appeal.

---

[47]Sprenger (2015) argues that a strong reference point can be formed in lottery choice when the same alternative is present in all rows on the multiple price list. In Round 11, we change both the lottery and the certainty equivalent, and thus the strong reference point is less likely to be present.

Table 3: Determinants of proportions of correct appeal decisions

Correct appeal decision	(1) Full sample	(2) Full sample	(3) Determined r.	(4) Determined r.	(5) Rand. r.	(6) Rand. r.
DirCutoffs	.110***	.109***	.034	.033	.198***	.208***
	(.028)	(.028)	(.038)	(.037)	(.051)	(.051)
SeqNo	.106***	.095***	.048	.031	.175***	.158***
	(.030)	(.030)	(.033)	(.033)	(.055)	(.057)
SeqCutoffs	.278***	.271***	.178***	.170***	.386***	.380***
	(.023)	(.023)	(.028)	(.028)	(.041)	(.042)
Dummy for random		-.250***				
		(.029)				
Rank of match		-.018**		-.075***		.083***
		(.008)		(.010)		(.015)
Period		.015***		-.011**		.013**
		(.004)		(.005)		(.007)
Dummy for truthful		.088***		-.049		.156***
		(.024)		(.036)		(.040)
Dummy for appeal inlottery		.025		-.026		.089*
		(.026)		(.032)		(.049)
Observations	1860	1860	932	932	928	928
No. of individuals	186	186	186	186	186	186
log(likelihood)	-1143	-1056	-474	-436	-601	-566

Notes: These are marginal effects of probit regressions on correct appeals. Models (1) and (2) include the full sample. Models (3) and (4) restrict the sample to determined rounds, when the match is correct. Models (5) and (6) restrict the sample to random rounds. The dummy for random is 1 in random rounds and 0 otherwise. The rank of match is the rank of the resulting match in the true preferences of the subject. The dummy for truthful is 1 if all participants played the truthful strategy (see Appendix B for an explanation). The dummy for appeal by lottery is 1 if the participants chose a lottery in Round 11 in the choice corresponding to the payoffs of their matches. * $p < 0.10$, **$p < 0.05$, *** $p < 0.01$. Standard errors are clustered on the individual level and are presented in parentheses.

Third, on average, subjects who play 'truthfully'[48] are more likely to make correct appeal decisions. The effect is driven by a significantly higher likelihood of appeal in random rounds. One possible reason could be that subjects who play truthfully are better at understanding the mechanism and thus at spotting random matches.

We present the results regarding the optimality of mechanism strategies, stability, and efficiency in Appendix B. The main takeaway is that verifiability is not necessarily aligned with reporting preferences truthfully: the proportion of stable allocations and the average efficiency is significantly higher under sequential DA than direct DA, independent of the cutoff provision. This result replicates the findings that

---

[48] A student plays truthfully when she submits the truthful list of all six schools under DirNo and DirCutoffs and when under SeqNo and SeqCutoffs she follows the straightforward strategy of applying to the best school among those that are still available.

suggest sequential DA performs better relative to direct DA (Bó and Hakimov, 2020a; Klijn, Pais, and Vorsatz, 2019).

7 Other School Admissions Applications

7.1 College-Proposing Deferred Acceptance

The corresponding college-proposing deferred acceptance (CPDA) rule inverts the role of students and schools in the CPDA algorithm, which determines feasible matchings as follows:

Let $P = (\succeq, v^*)$ be a problem. In the first step, an offer is sent from each school x to the q_x students with the highest score at that school. Any student i with multiple offers is tentatively matched with the top-ranked of these according to her preferences $\succeq_i$, and other offers are rejected. In subsequent steps, schools with rejected offers make offers to the students with the next-highest scores, if there are any, and subject to total offers less than or equal to q_x. As before, students with multiple offers are tentatively matched with the top-ranked of these according to their preferences, and other offers are rejected. As before, the algorithm terminates when there are no more rejections. The vector of tentative matches at this step is the feasible matching produced for this problem by CPDA.

The direct CPDA mechanism involves just one stage, in which all students are contacted and offered all schools, and are required to rank at least all acceptable schools. The designer's information at the end of this stage is sufficient to compute the feasible matching according to the above algorithm. As with SeqDA, there are many multi-stage mechanisms for CPDA. One such single-school CPDA mechanism converts steps of the CPDA algorithm into stages. In any stage, students are offered all schools that make offers to them according to the CPDA algorithm. Permitted actions involve selecting one of these schools, which is interpreted as the top-ranked school among offered schools. Other offered schools are rejected. The mechanism terminates when no more rejections are made.

Importantly, any school offered to a student at a stage of the single-school CPDA mechanism belongs to her core possibility set under the designer's information at that stage. To see this, notice that a student receiving an offer from a school can guarantee herself that school by simply top-ranking it in her actions, since she is one of the top-priority students at that school among those who have not rejected it already. No matter what actions other students take that remain compatible with the designer's information, the CPDA mechanism matches this student to this school. This is the very definition of a core possible school, and thus this mechanism is an ascending (structured-offers) mechanism. It is also non-wasteful and predictable.

Cutoffs for CPDA are the same as those for SPDA, namely, the minimum score for each school among

students matched to that school. The following result then follows immediately as a corollary to our main results.

PROPOSITION **3.** *Let $\mathcal{E}^* = (g^*, M^*)$ be a promised CPDA environment.*

1. If g^* is the single-school CPDA mechanism, then $\mathcal{E}^*$ is verifiable through experience.

2. If M^* is the cutoff protocol, then $\mathcal{E}^*$ is verifiable through communication.

3. If both, then $\mathcal{E}^*$ is transparent.

4. If M^* is the step-cutoff protocol, then $\mathcal{E}^*$ is strongly verifiable through communication.

Note that France switched from direct CPDA to a dynamic mechanism based on single-school CPDA for college admissions, motivated by the former mechanism's lack of transparency. Our results support the reform, in the sense that it indeed improved the transparency of the system.

7.2 TOP TRADING CYCLES

The Top Trading Cycles (TTC) procedure produces a feasible matching for each school admissions problem using the following algorithm:

Let $P = (\succeq, v^*)$ be a problem. A school is 'available' if it has unfilled capacity, while a student is 'remaining' if she is as yet unmatched. In the first step, each available school x makes an offer by pointing to the remaining student with the highest score at that school according to v^{*x}. Each student with an offer points at the most-preferred school among these, according to her preferences $\succeq_i$. There always exists at least one 'cycle,' i.e., a list of students $i_1, \ldots, i_k, i_{k+1} = i_1$ such that each i_j receives an offer from i_{j+1}'s most-preferred school. Each student in the cycle is matched to her most-preferred school, matched students are removed, and schools matched in this stage have their capacities reduced by one. All unfilled schools and remaining students go on to the next step, which proceeds the same way. The algorithm terminates when no more matches are made. The resulting matches comprise the feasible matching produced by TTC for this problem.

The <u>direct TTC</u> mechanism has only one stage, in which all students are contacted and offered all schools, and are required to rank at least all acceptable schools. The designer's information is sufficient at the end of this stage to produce a unique feasible matching according to the above algorithm. There can be many multi-stage variants of TTC. One such <u>single-school TTC</u> mechanism works as follows: in each stage, each remaining student is contacted and is offered all remaining schools. Each contacted student ranks exactly one of these schools (or the unassigned option), and this school is interpreted as being her most-preferred among remaining schools. The TTC algorithm is performed for this stage (cycles

are identified, matches made, students removed, and capacities adjusted). If there are still unmatched students, the mechanism goes to the next stage. It is straightforward to see that single-school TTC is an SSM, is structured-offers, and is also predictable, since each student is either matched to her last-chosen school, or is asked to choose again. These mechanisms are also non-wasteful.

Leshno and Lo (2020) establish that cutoffs for TTC take the form $\{p_b^c\}$, one for each pair of schools b, c, such that a student can use her priority at school b to gain admission to school c if her priority at school b is above the cutoff p_b^c. Cutoffs are thus more complicated under TTC than under DA. The cutoffs themselves are calculated in Leshno and Lo (2020) as the clearing prices in competitive equilibrium, where the units of the price are student scores at schools (a competitive equilibrium result for TTC is also established in Dur and Morrill (2018)). In our framework, the eligibility criterion for school x is thus multidimensional, and is based on students' scores at *other* schools. The following results follow as immediate corollaries to our main results:

PROPOSITION **4.** *Let $\mathcal{E}^* = (g^*, M^*)$ be a promised TTC environment.*

1. If g^* is the single-school TTC mechanism, then $\mathcal{E}^*$ is verifiable through experience.

2. If M^* is the cutoff protocol, then $\mathcal{E}^*$ is verifiable through communication.

3. If both, then $\mathcal{E}^*$ is transparent.

4. If M^* is the step-cutoff protocol, then $\mathcal{E}^*$ is strongly verifiable through communication.

7.3   BOSTON (IMMEDIATE ACCEPTANCE)

The Boston (or Immediate Acceptance (IA)) procedure produces a feasible matching for each problem according to the following algorithm:

> Let $P = (\succeq, v^*)$ be a problem. In the first step, an application is sent from each student $i \in N$ to her most-preferred acceptable school in X according to her preferences $\succeq_i$. Each school $x \in X$ accepts up to q_x applicants with the highest scores, rejecting all excess applications. Accepted applications are finalised and thus accepted students leave the problem with their matches. In each subsequent step, an application is sent from each rejected student i in the previous step to her most-preferred acceptable school according to $\succeq_i$ from among those that have not rejected her already, if any such school remains. Each school x accepts the applicants with the highest score, up to its remaining capacity, rejecting all others. The algorithm terminates when no more rejections are made. Accepted applications form the feasible matching for the problem P.

The direct IA mechanism has only one stage, in which all students are contacted, are offered all schools, and are required to at least rank all acceptable schools. The designer's information at the end of this stage

is sufficient to compute a unique feasible matching according to the above algorithm. One multi-stage IA mechanism is the single-school IA mechanism, which converts the steps of the IA algorithm into stages in the obvious way.[49] It is easy to see that the single-school IA mechanism is an SSM, structured-offers, and predictable, because each student is either matched to her last-chosen school, or is asked to choose again.

As noted by Doğan and Klaus (2018), 'a key feature of [IA] is that the relative rank of a school in the students' preferences is crucial in determining who receives a seat: seats are allocated to students who rank the school first, followed by those who rank it second *only when* seats are still available, and so on.' Scores of students at a school are used to break ties among those who assign the same rank to the school. This suggests that a student i has a higher value of the eligibility criterion for a school x under IA than a student j if: (1) she ranks x higher in her actions than j does; or (2) if i and j rank x at the same position in their actions, then i has a higher score at x than j. Thus, cutoffs for a school x under IA take the form (k, p), where k is the step in the IA algorithm in which x reaches capacity, and p is the minimum score among students matched to x in stage k. A school is in the possibility set for a student under these cutoffs only if she ranks x higher than rank k or, if she ranks x k^{th}, if she has a higher score than the value p. The following results then follow directly from our main results:

PROPOSITION **5.** *Let $\mathcal{E}^* = (g^*, M^*)$ be a promised IA environment.*

1. If g^* is the single-school IA mechanism, then $\mathcal{E}^*$ is verifiable through experience.

2. If M^* is the cutoff protocol, then $\mathcal{E}^*$ is verifiable through communication.

3. If both, then $\mathcal{E}^*$ is transparent.

4. If M^* is the step-cutoff protocol, then $\mathcal{E}^*$ is strongly verifiable through communication.

8   CONCLUSION

In this paper, we identify verifiability and transparency as desirable attributes for centralised admissions. We emphasise the importance of broadening the focus in mechanism design from simply choosing a mechanism to also designing the associated communication structure. We propose simple ways to achieve verifiability and transparency. These solutions provide simple blueprints for admissions policymakers concerned about transparency. Our experimental results support the idea that the two sources of feedback

---

[49]Importantly, in each stage, the designer offers all schools to which a student did not apply in previous stages. This is necessary to keep the mechanism close to the direct IA mechanism. If we instead restricted offers in a stage only to schools with remaining capacity, which might be more natural in a sequential setting, the mechanism becomes analogous to the 'modified' Boston mechanism (Dur, 2019).

improve observed verifiability, especially when provided together. Moreover, the experimental setup can be of independent interest as a tool to compare the relative transparency of alternative environments.

An open theoretical question is whether transparency is possible with less feedback. In particular, what would be the *minimum* feedback required for transparency? A related open question is to formulate a way to compare applications in terms of the informational requirements for transparency. We believe the intuition behind our results is also generalisable to other applications in market design. In auctions, for instance, cutoffs could be represented by winning prices, while predictable mechanisms might be similar to ascending or descending auction formats. We leave the formal derivation to future work.

Even within school admissions, our framework does not cover the efficiency-adjusted deferred acceptance mechanism (EADAM) (Kesten, 2010), based on obtaining consent from students to waive potential priority violations. Notably, the final matching produced by EADAM might contain instances where some students are matched to schools for which their scores do not meet the original SPDA cutoff. Our framework does not suggest a transparency solution for EADAM. For instance, a sequential mechanism might lead to some students being matched to schools that had rejected them previously, and thus might not be predictable (or even structured-offers). We leave transparency in EADAM to future research.

REFERENCES

ABDULKADIROĞLU, A., Y.-K. CHE, P. A. PATHAK, A. E. ROTH, AND O. TERCIEUX (2017): "Minimising Justified Envy in School Choice: The Design of New Orleans's OneApp," Technical report, National Bureau of Economic Research.

ABDULKADIROĞLU, A. AND T. SÖNMEZ (2003): "School choice: A mechanism design approach," *American Economic Review*, 93, 729–747.

AKBARPOUR, M. AND S. LI (2020): "Credible Auctions: A Trilemma," *Econometrica*, 88, 595–618.

ANDERSSON, T. AND L.-G. SVENSSON (2014): "Non-Manipulable House Allocation With Rent Control," *Econometrica*, 82, 507–539.

ASHLAGI, I., M. BRAVERMAN, Y. KANORIA, AND P. SHI (2017): "Communication Requirements and Informative Signaling in Matching Markets," in *EC*.

ASQUITH, P., T. COVERT, AND P. PATHAK (2019): "The Effects of Mandatory Transparency in Financial Market Design: Evidence from the Corporate Bond Market," NBER Working Paper 19417.

AUSUBEL, L. M. (2004): "An Efficient Ascending-Bid Auction for Multiple Objects," *American Economic Review*, 94, 1452–1475.

——— (2006): "An efficient dynamic auction for heterogeneous commodities," *American Economic Review*, 96, 602–629.

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

Bó, I. and R. Hakimov (2018): "The iterative deferred acceptance mechanism," Available at SSRN 2881880.

——— (2020a): "Iterated versus Standard Deferred Acceptance: Experimental Evidence," *The Economic Journal*, 130, 356–392.

——— (2020b): "Pick-an-Object Mechanisms," Working paper.

Budish, E., P. Cramton, and J. Shim (2015): "The high-frequency trading arms race: Frequent batch auctions as a market design response," *Q J Econ*, 130, 1547–1621.

Chen, Y. and T. Sönmez (2006): "School choice: an experimental study," *Journal of Economic Theory*, 127, 202–231.

Cramton, P. (1998): "Ascending auctions," *European Economic Review*, 42, 745–756.

Cramton, P. and J. A. Schwartz (2000): "Collusive bidding: Lessons from the FCC spectrum auctions," *Journal of Regulatory Economics*, 17, 229–252.

Doğan, B. and B. Klaus (2018): "Object allocation via Immediate-Acceptance: Characterizations and an affirmative action application," *Journal of Mathematical Economics*, 79, 140–156.

Dur, U. and O. Kesten (2019): "Sequential versus simultaneous assignment systems and two applications," *Economic Theory*, 68, 251–283.

Dur, U. and T. Morrill (2018): "Competitive equilibria in school assignment," *Games and Economic Behavior*, 108, 269–274.

Dur, U. M. (2019): "The modified Boston mechanism," *Mathematical Social Sciences*, 101, 31–40.

Echenique, F., A. J. Wilson, and L. Yariv (2016): "Clearinghouses for two-sided matching: an experimental study," *Quantitative Economics*, 7, 449–492.

Fischbacher, U. (2007): "z-Tree: Zurich toolbox for ready-made economic experiments," *Experimental Economics*, 10, 171–178.

Gal-Or, E., M. Gal-Or, and A. Dukes (2007): "Optimal information revelation in procurement schemes," *RAND Journal of Economics*, 38, 400–418.

Gale, D. and L. S. Shapley (1962): "College admissions and the stability of marriage," *American Mathematical Monthly*, 69, 9–15.

Gonczarowski, Y. A., N. Nisan, R. Ostrovsky, and W. Rosenbaum (2014): "A Stable Marriage Requires Communication," *ArXiv*, abs/1405.7709.

Gong, B. and Y. Liang (2016): "A Dynamic College Admission Mechanism in Inner Mongolia: Theory and Experiment," Working Paper.

Greiner, B. (2003): "An online recruitment system for experimental economics," *Research and Scientific Computing*, 63, 79–93.

GRENET, J., Y. HE, AND D. KÜBLER (2022): "Preference Discovery in University Admissions: The Case for Dynamic Multi-offer Mechanisms," .

GRIGORYAN, A. (2022): "Transparency in Allocation Problems," .

HAERINGER, G. AND V. IEHLÉ (2019): "Gradual college admission," Available at SSRN 3488038.

HAGENBACH, J., F. KOESSLER, AND E. PEREZ-RICHET (2014): "Certifiable pre-play communication: Full disclosure," *Econometrica*, 82, 1093–1131.

HAKIMOV, R., C.-P. HELLER, D. KÜBLER, AND M. KURINO (2021a): "How to avoid black markets for appointments with online booking systems," *American Economic Review.*

HAKIMOV, R. AND D. KÜBLER (2020): "Experiments on centralized school choice and college admissions: a survey." *Experimental Economics.*

HAKIMOV, R., D. KÜBLER, AND S. PAN (2021b): "Costly Information Acquisition in Centralized Matching Markets," Working paper.

HOLT, C. A. AND S. K. LAURY (2002): "Risk aversion and incentive effects," *American Economic Review*, 92, 1644–1655.

IMMORLICA, N., J. LESHNO, I. LO, AND B. LUCIER (2020): "Information acquisition in matching markets: The role of price discovery," Working paper.

KESTEN, O. (2010): "School choice with consent," *Q J Econ*, 125, 1297–1348.

KLIJN, F., J. PAIS, AND M. VORSATZ (2019): "Static versus dynamic deferred acceptance in school choice: Theory and experiment," *Games and Economic Behavior*, 113, 147–163.

KOUTOUT, K., M. V. DER LINDEN, A. DUSTAN, AND M. WOODERS (2019): "Mechanism performance under strategy advice and sub-optimal play: A school choice experiment," Working paper.

LESHNO, J. D. AND I. LO (2020): "The Cutoff Structure of Top Trading Cycles in School Choice," *Review of Economic Studies*, rdaa071.

LI, S. (2017): "Obviously Strategy-proof Mechanisms," *American Economic Review*, 107, 3257–87.

MEIJER, A. (2013): "Understanding the complex dynamics of transparency," *Public Administration Review*, 73, 429–439.

MILGROM, P. R. AND R. J. WEBER (1982): "A Theory of Auctions and Competitive Bidding," *Econometrica*, 50, 1089–1122.

MÖLLER, M. (2022): "Transparent Matching Mechanisms," *Available at SSRN 4073464.*

MORRILL, T. (2015): "Two simple variations of top trading cycles," *Economic Theory*, 60, 123–140.

NAOR, M., B. PINKAS, AND R. SUMNER (1999): "Privacy preserving auctions and mechanism design," *EC*, 99, 129–139.

NELSON, P. J. (2001): "Transparency Mechanisms at the Multilateral Development Banks," *World Development*, 29, 1835–1847.

NISAN, N. AND I. SEGAL (2006): "The communication requirements of efficient allocations and supporting prices," *Journal of Economic Theory*, 129, 192–224.

NÚÑEZ, M. (2019): "Towards Transparent Mechanisms," in *The Future of Economic Design*, ed. by J.-F. Laslier, H. Moulin, M. Sanver, and W. Zwicker, Springer, 341–346.

PAN, S. (2020): "Experiments in market design," in *Handbook of Experimental Game Theory*, Edward Elgar Publishing.

PYCIA, M. AND P. TROYAN (2019): "A Theory of Simplicity in Games and Mechanism Design," Working paper.

PYCIA, M. AND M. U. ÜNVER (2020): "Arrovian efficiency and auditability in the allocation of discrete resources," Working paper.

ROTH, A. E. (1982): "The economics of matching: stability and incentives," *Mathematics of Operations Research*, 7, 617–628.

ROTH, A. E. AND E. PERANSON (1999): "The Redesign of the Matching Market for American Physicians: Some Engineering Aspects of Economic Design," *American Economic Review*, 89, 748–780.

ROTH, A. E., T. SÖNMEZ, AND M. U. ÜNVER (2004): "Kidney Exchange," *Q J Econ*, 119, 457–488.

SEGAL, I. (2007): "The communication requirements of social choice rules and supporting budget sets," *Journal of Economic Theory*, 136, 341–378.

SÖNMEZ, T. AND T. SWITZER (2013): "Matching with (branch-of-choice) contracts at the United States Military Academy," *Econometrica*, 81, 451–488.

SPRENGER, C. (2015): "An endowment effect for risk: Experimental tests of stochastic reference points," *Journal of Political Economy*, 123, 1456–1499.

WEST, A., H. PENNELL, AND P. NODEN (1998): "School admissions: Increasing equity, accountability and transparency," *British Journal of Educational Studies*, 46, 188–200.

WOODWARD, K. (2020): "Self-Auditable Auctions," Working paper.

A APPENDIX A: OMITTED PROOFS

A.1 PROOF OF THEOREM 1

Let $\mathcal{E}^* = (g^*, M^*)$ be the promised environment, let g^* be non-wasteful, and let (ϕ^{g^*}, a^*) denote the outcome of g^*. Let $i \in N$ be a student. In one direction, let $M^*(\phi^{g^*})$ be informationally equivalent to cutoffs, i.e., $Poss_i^{g^*}(\phi_i^{g^*} \cap M^*(\phi^{g^*})) = Poss_i^{g^*}(\phi_i^{g^*} \cap c^{g^*}(a^*, \phi^{g^*}))$. We show that $Poss_i^{g^*}(\phi_i^{g^*} \cap c^{g^*}(a^*, \phi^{g^*})) = a_i^*$. First, by construction of cutoffs, we have that i has a higher value of the eligibility criterion for a_i^* than its cutoff value, i.e., $\rho_i^{g^*, a_i^*} \geq c_{a_i^*}^{g^*}(a^*, \phi^{g^*})$. It follows that a_i^* is possible for i under cutoffs, i.e., $a_i^* \in Poss_i^{g^*}(c^{g^*}(a^*, \phi^{g^*}))$. Moreover, since g^* is non-wasteful, $a_i^* \succeq_i \emptyset$ at each problem $P = (\succeq, v^*)\phi^{g^*}$. To guarantee this, a_i^* must be actively ranked by i in her actions at some stage of g^*, and so a_i^* is possible

for i under her intrinsic information $\phi_i^{g^*}$. Thus $a_i^* \in Poss_i^{g^*}(\phi_i^{g^*} \cap c^{g^*}(a^*, \phi^{g^*}))$. We show that no other match is possible for i under her information, i.e., $x \neq a_i^* \implies x \notin Poss_i^{g^*}(\phi_i^{g^*} \cap c^{g^*}(a^*, \phi^{g^*}))$. There are three cases to consider:

Case 1: x is not offered to i at any stage in g^*. Then x cannot be ranked by i in g^*, and there is a problem $P = (\succeq, v^*) \in \phi^{g^*}$ in which $x \succ_i a_i^*$. Since g^* is non-wasteful, there is a set $B = \{j \in N \mid a_j^* = x\}$ of students matched to x in a^* with $|B| = q_x$. Since $x \succ_i a_i^*$, each student in B has a higher value of the eligibility criterion for x at ϕ^{g^*} than i does, i.e., $\rho_j^{g^*,x}(\phi^{g^*}) > \rho_i^{g^*,x}(\phi^{g^*})$ for each $j \in B$. By definition of cutoffs, therefore, it follows that $c_x^{g^*}(a^*, \phi^{g^*}) \equiv \min_{j \in B}\{\rho_j^{g^*,x}\} > \rho_i^{g^*,x}(\phi^{g^*})$, and so x is not possible for i under cutoffs for a^*, i.e., $x \notin Poss_i^{g^*}(c^{g^*}(a^*, \phi^{g^*}))$. Thus $x \notin Poss_i^{g^*}(\phi_i^{g^*} \cap c^{g^*}(a^*, \phi^{g^*}))$ in this case.

Case 2: x is offered to i and ranked higher in i's actions than a_i^*. Then, for each problem $P = (\succeq, v^*) \in \phi^{g^*}$, we have that $x \succ_i a_i^*$, and the logic of the previous case follows to imply that $x \notin Poss_i^{g^*}(c^{g^*}(a^*, \phi^{g^*}))$. Thus $x \notin Poss_i^{g^*}(\phi_i^{g^*} \cap c^{g^*}(a^*, \phi^{g^*}))$ in this case as well.

Case 3: x is offered to i and ranked lower in i's actions than a_i^*. Suppose for contradiction that x is possible for i, i.e., $x \in Poss_i^{g^*}(\phi_i^{g^*} \cap c^{g^*}(a^*, \phi^{g^*}))$. By definition of the possibility set, there is some compatible problem $P = (\succeq, v^*) \in (\phi_i^{g^*} \cap c^{g^*}(a^*, \phi^{g^*}))$ with $g_i^*(P) = x$. Notice that $a_i^* \succ_i x$. Since g^* is non-wasteful, there is a set $B = \{j \in N \mid g_j^*(P) = a_i^*\}$ of students matched to a_i^* in $g^*(P)$, with $|B| = q_{a_i^*}$. Since $x \succ_i a_i^*$, each student in B has a higher value of the eligibility criterion than i, i.e., $\rho_j^{g^*,a_i^*}(\phi^{g^*}) > \rho_i^{g^*,a_i^*}(\phi^{g^*})$ for each $j \in B$. By construction of cutoffs, it follows that $c_{a_i^*}^{g^*}(g^*(P), \phi^{g^*}) > \rho_i^{g^*,a_i^*}(\phi^{g^*})$. Thus $g^*(P)$ does not generate the observed cutoffs $c^{g^*}(a^*, \phi^{g^*})$, which means P is not compatible with $c^{g^*}(a^*, \phi^{g^*})$. This contradicts our assumption that $P \in (\phi_i^{g^*} \cap c^{g^*}(a^*, \phi^{g^*}))$. Thus, our supposition was false, and $x \notin Poss_i^{g^*}(\phi_i^{g^*} \cap c^{g^*}(a^*, \phi^{g^*}))$.

In effect, we have shown that $Poss_i^{g^*}(\phi_i^{g^*} \cap c^{g^*}(a^*, \phi^{g^*})) = a_i^*$, from which it follows that $Poss_i^{g^*}(\phi_i^{g^*} \cap M^*(\phi^{g^*})) = a_i^*$. Thus, $\mathcal{E}^*$ is verifiable through communication. In the other direction, let $\mathcal{E}^*$ be verifiable through communication. Then a_i^* is justified in $\mathcal{E}^*$ for each $i \in N$ by $M^*(\phi^{g^*})$. In particular, $Poss_i^{g^*}(\phi_i^{g^*} \cap M^*(\phi^{g^*})) = a_i^*$ by definition. Since we have demonstrated that $Poss_i^{g^*}(\phi_i^{g^*} \cap c(a^*, \phi^{g^*})) = a_i^*$ as well, it follows that $M^*(\phi^{g^*})$ is informationally equivalent to cutoffs. ∎

A.2 PROOF OF PROPOSITION 1

Let $\mathcal{E}^* = (g^*, M^*)$ be the promised environment, let g^* be a non-wasteful and structured-offers mechanism, and let g^* be an SSM. Let (ϕ^{g^*}, a^*) be the outcome in g^*. We consider descending and ascending mechanisms in turn:

Descending mechanism: Let $i \in N$ be a student. Let k be the first stage in g^* in which she is contacted. By definition of descending mechanisms, she is offered all schools that in her possibility set under the designer's information at that stage, i.e., $Y_i^{g^*,k} = Poss_i^{g^*}(\phi^{g^*,\{k-1\}})$. Let $\geq_i^k$ denote her action in this stage. Since g^* is an SSM, i ranks exactly one of these schools in her actions, say school x. Since she ranks only one school, x is the only school in her possibility set under experiential information at this stage, i.e., $x = Poss_i^{g^*}(\phi_i^{g^*,\{k\}} \cap \epsilon_i^{\{k\}})$, which satisfies P1. Since g^* is non-wasteful, she can only be matched to a school she has actively ranked in her actions. Thus, if i is not contacted again, then $a_i^* = x$, which satisfies P2. Instead, let k' be the stage in which she is contacted again. By definition of descending mechanisms, she is offered a subset of schools she was previously offered, i.e., $Y_i^{g^*,k'} \subset Y_i^{g^*,k}$. Let $\geq_i^{k'}$ denote her action in this stage. If she continues to be offered x, i.e., $x \in Y_i^{g^*,k'}$, then she continues to rank x in this stage, and the previous reasoning applies. On the other hand, if she is not offered x, i.e., $x \notin Y_i^{g^*,k'}$, her experiential information is updated such that this school is not possible for her, and so $x \notin Poss_i^{g^*}(\epsilon_i^{g^*,k'})$. Since g^* is an SSM, she again ranks a single school y, and since g^* is non-wasteful, y is the only school in her cumulative actions which is possible for her, thus $y = Poss_i^{g^*}(\phi_i^{g^*,\{k'\}} \cap \epsilon_i^{g^*,\{k'\}})$, which satisfies P1. Again, if she is not contacted again, $a_i^* = y$, which satisfies P2, otherwise a repeated argument follows. Since this is true for each student, we have that P1 and P2 are always satisfied, and g^* is predictable.

Ascending mechanism: Let $i \in N$ be a student. Let k be the first stage in g^* in which she is contacted. By definition of ascending mechanisms, she is offered all schools that in her core possibility set under the designer's information at that stage, i.e., $Y_i^{g^*,k} = Core_i^{g^*}(\phi^{g^*,\{k-1\}})$. Let $\geq_i^k$ denote her action in this stage. Since g^* is an SSM, i ranks one of these schools in her actions, say x. Since she ranks only one school, x is the only school in her possibility set under experiential information at this stage, i.e., $x = Poss_i^{g^*}(\phi_i^{g^*,\{k\}} \cap \epsilon_i^{\{k\}})$, which satisfies P1. Since g^* is non-wasteful, she can only be matched to a school she has actively ranked in her actions. Thus, if i is not contacted again, then $a_i^* = x$, which satisfies P2. On the other hand, let k' be the stage in which she is contacted again. Let $\geq_i^{k'}$ denote her action in this stage. By definition of ascending mechanisms, she continues to be offered x, i.e., $x \in Y_i^{g^*,k'}$. Since g^* is an SSM, she ranks one school in her action at this stage. If she ranks x again, then the previous reasoning applies. If she ranks some $y \neq x$, then her experiential information is updated such that x is not possible for her, i.e., $x \notin Poss_i^{g^*}(\epsilon_i^{g^*,k'})$. Since g^* is non-wasteful, y is the only school in her cumulative actions which is possible for her, $y = Poss_i^{g^*}(\phi_i^{g^*,\{k'\}} \cap \epsilon_i^{g^*,\{k'\}})$, which satisfies P1. Again, if she is not contacted again, $a_i^* = y$, which satisfies P2, otherwise a repeated argument follows. Since this is true for each student, we have that P1 and P2 are always satisfied, and g^* is predictable. ∎

A.3 PROOF OF THEOREM 2

Let $\mathcal{E}^* = (g^*, M^*)$ be the promised environment, let g^* be a non-wasteful and structured-offers mechanism, and let (ϕ^{g^*}, a^*) be the outcome. In one direction, let g^* be predictable. Then g^* satisfies P1 and P2. Together, these imply that there is a unique school for each student in her possibility set under experiential information at the terminal stage of the mechanism, and moreover that this school is her match in a^*, which makes $\mathcal{E}^*$ verifiable through experience. In the other direction, let $\mathcal{E}^*$ be verifiable through experience. We show that g^* is predictable by treating descending and ascending mechanisms separately:

Descending mechanism: Let $i \in N$ be a student, and let k be a stage in which she is contacted. Let $\geq_i^k$ denote her action in this stage. By definition of descending mechanisms, each offered school $x \in Y_i^{g^*,k}$ is possible for her under the designer's information. If she ranks only one school, then P1 is satisfied for this stage. Instead, suppose she ranks more than one school. Since these schools are ranked strictly, there exists x among these such that $x \succ_i^k y$ for all other y ranked in this stage. If x remains possible for i under the designer's terminal information, then it is her match, i.e., $x \in Poss_i^{g^*}(\phi^{g^*}) \implies a_i^* = x$. Thus, in particular, x is possible for i under her experiential information at k, i.e., $x \in Poss_i^{g^*}(\phi_i^{g^*,k} \cap \epsilon_i^{g^*,k})$. On the other hand, suppose i is not matched to x, and $a_i^* = y$ for some $y \neq x$. Since $\mathcal{E}^*$ is verifiable through experience, this means that x is removed from the possibility set for i under experiential information at some future stage. In particular, there is a stage $k' > k$ in which she is not offered x, i.e., $x \notin Y_i^{g^*,k'}$. At stage k, therefore, student i knows that if she is to be matched to a school other than x, she will be contacted again, and since g^* is non-wasteful, she will be offered this school. Knowing this, $y \notin Poss_i^{g^*}(\phi_i^{g^*,k} \cap \epsilon_i^{g^*,k})$ for each $y \neq x$. Thus $Poss_i^{g^*}(\phi_i^{g^*,k} \cap \epsilon_i^{g^*,k}) = x$, which satisfies P1. In particular, this is true for the last stage in which i is contacted. Since i does not take any further actions or acquire experiential information if she is not contacted again, the unique school in her possibility set under experiential information at the terminal stage of the mechanism is her match, which satisfies P2. Thus g^* is predictable.

Ascending mechanism: Let $i \in N$ be a student, and let k be a stage in which she is contacted. Let $\geq_i^k$ denote her action in this stage. By definition of ascending mechanisms, each offered school $x \in Y_i^{g^*,k}$ is core possible for her under the designer's information. If she ranks only one school, then P1 is satisfied for this stage. Instead, if she ranks more than one school, there exists x among these such that $x \succ_i^k y$ for all other y ranked in this stage. As before, if x remains possible for i under the designer's terminal information, then it is her match, i.e., $x \in Poss_i^{g^*}(\phi^{g^*}) \implies a_i^* = x$. In particular, x is possible for i under her experiential information at stage k, i.e., $x \in Poss_i^{g^*}(\phi_i^{g^*,k} \cap \epsilon_i^{g^*,k})$. On the other hand, suppose i is not matched to x, and $a_i^* = y$ for some $y \neq x$. By definition of core possibility, this means y is ranked above x in i' actions. This means that y is offered to i at some stage, which means y is core possible for i at some stage. In particular, y is not offered to i in stage k, and since $\mathcal{E}^*$ is verifiable through experience,

there is a stage $k' > k$ in which i is offered y, i.e., $y \in Y_i^{g^*,k'}$, and she ranks y above x in her actions. At stage k, therefore, student i knows that if she is to be matched to a school $y \neq x$, she will be contacted again, and since g^* is non-wasteful, she will be offered y. Knowing this, $y \notin Poss_i^{g^*}(\phi_i^{g^*,k} \cap \epsilon_i^{g^*,k})$ for each $y \neq x$. Thus $Poss_i^{g^*}(\phi_i^{g^*,k} \cap \epsilon_i^{g^*,k}) = x$, which satisfies P1. In particular, this is true for the last stage in which i is contacted. Since i does not take any further actions or acquire experiential information if she is not contacted again, the unique school in her possibility set under experiential information at the terminal stage of the mechanism is her match, which satisfies P2. Thus g^* is predictable. ∎

A.4 Proof of Theorem 3

Let $\mathcal{E}^* = (g^*, \mathcal{M}^*)$ be the promised environment, and let g^* be a non-wasteful and structured-offers mechanism. Let (ϕ^{g^*}, a^*) be the outcome in $\mathcal{E}^*$. Let $\mathcal{E} = (g, M)$ be a plausible environment, and let (ϕ^g, a) be the outcome in $\mathcal{E}$. To show that $\mathcal{E}^*$ is transparent, we must show that $a_i \neq a_i^*$ is not justified by $\phi_i^{\mathcal{E}}$ for some student $i \in N$, i.e., produces a detectable violation.

Firstly, let $\mathcal{M}^*$ be full disclosure. Then indistinguishability requires that the used communication also be full disclosure, i.e., $M(\phi^g) = \phi^g$, otherwise some student can spot that her preferences or scores were not accurately communicated. Since $a^* \in g^*(\phi^{g^*})$ is unique, so must be $a \in g^*(\phi^g)$, otherwise the deviation is immediately detectable. Moreover, under our behavioural assumption that students use the same type in g as in g^*, and since scores are fixed, it follows that $\phi^g = \phi^{g^*}$, and so $M(\phi^g) = \phi^{g^*}$. Since a_i^* is the unique match justified by $\phi_i^{\mathcal{E}}$ for each $i \in N$, $\mathcal{E}^*$ is transparent. Instead, let $\mathcal{E}^*$ be verifiable through both experience and communication. Thus g^* is a predictable mechanism (Theorem 2), and M^* is possibility-equivalent to cutoffs (Theorem 1). Recall that indistinguishability requires that g is non-wasteful, and has the same permitted actions and dynamic of offers as g^*, and moreover that a_i is justified by $M(\phi^g)$ for each $i \in N$.

Since g is indistinguishable from g^*, experiential information from g is treated by each student as if it were generated by g^*. In particular, since g^* is predictable, possibility sets under experiential information are single-valued at each stage k in g, i.e., for each $i \in N$ and each stage k in g $Poss_i^{g^*}(\phi_i^{g,k} \cap \epsilon_i^{g,k}) = x$ for some $x \in X$. Set $b_i^k = Poss_i^{g^*}(\phi_i^{g,k} \cap \epsilon_i^{g,k})$ for $i \in N$, and call it her 'temporary match' at this stage. The 'temporary matching' at stage k in g collects these temporary matches, and is given by $b^k \equiv (b_i^k)_{i \in N}$. Notice that predictability of g^* ensures that $a = b^K$ for the terminal stage K of g, otherwise experiential information is contradicted for some student, and the deviation is detected. Thus we essentially have to show that $b^K = a^*$ at the terminal stage of g. We prove the result separately for descending and ascending mechanisms.

Case 1: g^* is descending:

LEMMA **1.** *In any stage in g, any contacted student is offered all schools in her possibility set under the designer's information and g^* at that stage. Formally, for any stage k in g and any $i \in Act^{g,k}$, we have that $Poss_i^{g^*}(\phi^{g,\{k-1\}}) \subseteq Y_i^{g,k}$.*

Proof: Suppose not. Then there is a stage k in g, a student $i \in Act^{g,k}$ and a school $x \in Poss_i^{g^*}(\phi^{g,\{k-1\}})$, such that $x \notin Y_i^{g,k}$. Notice that $x \notin Y_i^{g,l}$ for all $l > k$ (i.e., x cannot be offered to i in any future stage of g, since g^* is descending and g is indistinguishable from g^*).

Suppose x is not filled in the temporary matching b^{k-1}. We show that this could lead to waste. Since x is not offered to i in stage k, it is not in her possibility set under experiential information, i.e., $x \notin Poss_i^{g^*}(\epsilon_i^{g,k})$. However, we can find preferences $\succeq_i$ for i in which x is ranked above all schools she has not already ranked in her actions in g by stage k, and all schools she has not yet ranked are unacceptable. Moreover, let J be the set of students who have not yet ranked x in their actions up to stage k. We can find preferences $\succeq_j$ for $j \in J$ such that x is unacceptable. Defining the problem $P = (\succeq, v^*)$, we see that P is compatible with the designer's information $\phi^{g,\{k-1\}}$. By non-wastefulness, no student in J can be matched to x in a. Since x was not filled in b^{k-1}, x is not filled in a at this problem P. But then, non-wastefulness implies $a_i = x$ at P. However, this would violate the experiential information for i, since i can only be matched to schools which she actively ranks, and she cannot rank x. This violates verifiability through experience, and is a contradiction.

So x is filled in b^{k-1}. In particular, there is a set of agents J with $|J| = q_x$ such that $b_j^{k-1} = x$ for each $j \in J$. We show that this leads to a communication $M(\phi^g)$ that does not justify a_i for i. As before, there exists a problem $P = (\succeq, v^*)$ compatible with the designer's information $\phi^{g,k}$ in which x is ranked by i above all schools she has not yet ranked in her actions g, each of which is also unacceptable to her, and moreover x is unacceptable for each j who has not yet ranked x. By non-wastefulness, $a_j = x$ for all $j \in J$. Since the communication $M(\phi^g)$ justifies a_j for each j, it follows that $x \in Poss_j^{g^*}(M(\phi^g))$ for each $j \in J$. However, since x is possible for i under the designer's information, there is at least one $j \in J$ such that $\rho_i^{g^*,x} \geq \rho_j^{g^*,x}$. Thus $x \in Poss_i^{g^*}(M(\phi^g))$. But $a_i \neq x$, which violates verifiability through communication. Thus, our assumption that x is not offered to i at k in g is false, which proves the lemma. ∎

We now show that $a \neq a^*$. We prove this by an induction argument on stages of g.

CLAIM **1.** $Poss_i^{g^*}(\phi^{g,\{k\}}) = Poss_i^{g^*}(\phi^{g^*,\{k\}})$ *for each $i \in N$ and each stage k in g.*

Proof: **Basis step:** Consider Stage 1 of g. Any student $i \in Act^{g,1}$ contacted in this stage is offered all schools by Lemma 1 and the fact that g^* is a non-wasteful mechanism and so all schools are possible for her under the designer's initial information. Let $\geq^{g,1}$ be the action profile in ^{g} at this stage and let $\phi^{g,1}$ be the designer's information. In particular, $\phi^{g,1}$ is possibility equivalent to $\phi^{g^*,1}$, the information the designer

would have obtained in g^* by contacting students in $Act^{g,1}$. That is, $Poss_i^{g^*}(\phi^{g,1}) = Poss_i^{g^*}(\phi^{g^*,1})$ for each $i \in N$. Moreover, since g^* is predictable and g is indistinguishable from g^*, the temporary matching b^1 in this stage of g contains the top-ranked school among all schools in the actions of each $i \in Act^{g,1}$.

Induction hypothesis: Let $k \geq 1$ be a stage in g. Suppose $Poss_i^{g^*}(\phi^{g,\{k-1\}}) = Poss_i^{g^*}(\phi^{g^*,\{k-1\}})$ for each $i \in N$.

Induction Step: Consider Stage k in g. Any student $i \in Act^{g,k}$ is offered all schools that are possible for her under g^* and the designer's information $\phi^{g,\{k-1\}}$ (Lemma 1). Since $Poss_i^{g^*}(\phi^{g,\{k-1\}}) = Poss_i^{g^*}(\phi^{g^*,\{k-1\}})$ for each $i \in N$ by the induction hypothesis, each contacted student is offered all schools which she would have been offered in g^* if she had been contacted. Since g is indistinguishable from g^*, each contacted student takes an action as she would in g^*, and since g^* is predictable, the single school in her possibility set under experiential information is the top-ranked of these schools. This gives us the temporary match b^k at this stage. Moreover, $Poss_i^{g^*}(\phi^{g,\{k\}}) = Poss_i^{g^*}(\phi^{g^*,\{k\}})$ for each $i \in N$. ∎

Predictability ensures that $b^K = a$ at the terminal stage K in g. Since g terminates when the possibility set for each student is single-valued, Claim 1 implies that K is such that $Poss_i^{g^*}(\phi^{g,\{k\}}) = a_i^*$ for each $i \in N$. Thus $a = a^*$, and since $\mathcal{E}$ was an arbitrary plausible environment, we have that $\mathcal{E}^*$ is transparent.

Case 2: g^* is ascending:

We start with a useful lemma that shows that at any stage of g, a school cannot be offered to a student for whom it is not in her core possibility at that stage under the designer's information and g^*.

LEMMA **2.** *For any stage k in δ^g, and any contacted student $i \in Act^{g,k}$, $x \notin Core_i^{g^*}(\phi^{g,\{k-1\}}) \implies x \notin Y_i^{g,k}$.*

Proof: Suppose for contradiction that there is a stage k in g, a contacted student $i \in Act^{g,k}$, and a school $x \notin Core_i^{g^*}(\phi^{g,\{k-1\}})$, such that $x \in Y_i^{g,k}$.

CLAIM **2.** The school x is unfilled in b^{k-1}.

Proof: Suppose for contradiction that x is filled in b^{k-1}. Let J be the students matched to x in b^{k-1}. We can find a problem $P = (\succeq, v^*) \in \phi^{g,\{k-1\}}$ compatible with the designer's information in which each student $j \in J \cup \{i\}$ ranks x above any school she has not yet ranked, each of which she considers unacceptable. Since g^* is predictable, x is the unique school in the possibility set induced by experiential information for each such student $j \in J \cup \{i\}$, i.e., $Poss_j^{g^*}(e_j^{g,\{k\}} \cap \phi_j^{g,k}) = x$. Since g is non-wasteful, g^* is verifiable by experience, and g is indistinguishable from g^*, we have that $a_j = x$ for each $j \in J \cup \{i\}$. This violates capacity constraints for x, as $q_x + 1$ students are matched to x in a. ∎

Claim 2 establishes that x is unfilled in b^{k-1}. Since $x \notin Core_i^{g^*}(\phi^{g,\{k-1\}})$, by definition there is some set

of students B such that $x \in Core_j^{g^*}(\phi^{g,\{k-1\}})$ for each $j \in B$. Moreover, $\rho_j^{g^*,x}(\phi^{g,\{k-1\}}) \geq \rho_i^{g^*,x}(\phi^{g,\{k-1\}})$ for each $j \in B$. We show that this leads to either x being matched to more students than its capacity, which violates capacity constraints, or that some student in B prefers x to her match and is not matched to x, which leads to a violation in verifiability through communication. Notice that there is a problem $P = (\succeq, v^*) \in \phi^{g,\{k-1\}}$ compatible with the designer's information in which each $j \in B \cup \{i\}$ ranks x higher than any school y that she has not yet ranked, and each such school y is unacceptable. By predictability, it follows that $a_i = x$ at this problem. Moreover, if x is offered to each $j \in B$ at some stage of g, $x = Poss_j^{g^*}(\phi_j^{g,\{k\}} \cap \epsilon_j^{g,\{k\}})$, and then predictability implies $a_j = x$ for each $j \in B$, which violates capacity constraints as $q_x + 1$ students are matched to x in a. Thus there is some student $j \in B$ who is not offered x at any stage in g. Thus j cannot rank x, and by predictability, $a_j \neq x$. Since a_i is justified by $M(\phi^g)$ for each $i \in N$, and since $\rho_j^{g^*,x}(\phi^g) \geq \rho_i^{g^*,x}(\phi^g)$ by definition of core possibility, $x \in Poss_j^{g^*}(M(\phi^g \cap \phi_j^g))$, which violates verifiability through communication, as $a_j \neq x$. This is a contradiction. Thus, our supposition was false, and $x \notin Core_i^{g^*}(\phi^{g,\{k-1\}}) \implies x \notin Y_i^{g,k}$. ∎

By Lemma 2, each school x is offered only to students for whom it is in their core possibility at that stage. In the first stage of g, each contacted student is offered some subset of her initial core possibility. Any school in her core possibility not offered to her at this stage cannot be ranked by her, and remains in her core possibility by definition. Thus each such school must be offered to her in a future stage before it is offered to any student for whom it is not in the core possibility (Lemma 2). In particular, there is a stage k in g in which each student has been offered all schools in her initial core possibility according to g^*. Since g^* is predictable and g is indistinguishable from g^*, the unique school in each student's possibility set at k under experiential information is the top-ranked of these schools. The temporary match b_i^k for i at k is precisely this school. By our assumption that students use the same type in g as in g^*, this school is the same as it would be in g^* under the same information. Moreover, by Lemma 2, any student contacted in a future stage is offered any school that enters her core possibility according to g^* and the designer's information. By predictability, $b^K = a$ at the terminal stage K of g. However, it also follows that $b_i^K = a_i^*$. Thus $a = a^*$, and since $\mathcal{E}$ was an arbitrary plausible environment, we have the required result. ∎

A.5 PROOF OF PROPOSITION 2

For a mechanism g, define a single-school equivalent (SSE) mechanism g^S such that (1) g^S is an SSM; and (2) $g(P) = g^S(P)$ for each $P \in \mathcal{P}$. Let the designer's information at any stage k in g^S be $\phi^{g^S,\{k\}}$. Since g^S is an SSM, we can define a temporary matching at k given by $a^{g^S,k} \in X^N$, comprising the last-chosen school for each student (the empty school if none exists). Let $c^g(a^{g^S,k}, \phi^{g^S,\{k\}})$ be the cutoffs at k for this tempo-

rary matching. Then, the underline{step-cutoffs} for g are given by the collection $(c^{g^S}(a^{g^S,k}, \phi^{g^S,\{k\}}))_{k \leq K^{g^S}}$ of step-cutoffs at every stage of the SSE g^S. A communication protocol M is a underline{step-cutoff protocol for g} if it communicates the collection of step-cutoffs generated by the SSE for g, i.e., $M(\phi^g) = (c^{g^S}(a^{g^S,k}, \phi^{g^S,\{k\}}))_{k \leq K^{g^S}}$ where g^S is the SSE for g. Then:

Let $\mathcal{E}^* = (g^*, \mathcal{M}^*)$ be the promised environment, let g^* be a non-wasteful mechanism. Let (ϕ^{g^*}, a^*) be the outcome in $\mathcal{E}^*$, let g^{*S} be the SSE for g^*, and let $\mathcal{M}^*$ be the step-cutoff protocol for g^*. Thus $M^*(\phi^{g^*}) = (c^{g^{*S}}(a^{g^{*S},k}, \phi^{g^{*S},\{k\}}))_{k \leq K^{g^{*S}}}$. Since $\mathcal{E}^*$ is verifiable through communication by Theorem 1, a_i^* is justified for each $i \in N$ by $M^*(\phi^{g^*})$. We show that a^* is the unique justifiable matching in $\mathcal{E}^*$. Since g^{*S} is an SSM, it is enough to consider only schools actively ranked by students in their actions in g^{*S}. In particular, even though g^* might be a direct mechanism in which all schools are ranked, we only need to consider schools that students rank higher in their actions in g^{*S} than their matches in a^*.

Case 1: g^{*S} is descending: Take student $i \in N$ and a school $x >_i^{g^*} a_i^*$ ranked higher in her actions. Since $x \notin Poss_i^{g^*}(\phi^{g^*})$, there is a first stage k in g^{*S} such that $x \notin Poss_i^{g^*}(\phi^{*S,\{k\}})$. By construction of step-cutoffs, it follows that cutoffs for x at the previous stage are higher than i's value of the eligibility criterion for x. Moreover, since g^{*S} is descending, the cutoffs for x are weakly increasing in subsequent stages. Thus, there is no outcome (a, ϕ^{g^*}) with $a_i = x$ that is justifiable according to these step-cutoffs. Since i and x were arbitrary, this is always true. Thus $\mathcal{E}^*$ is strongly verifiable through communication.

Case 2: g^{*S} is ascending: Take a student $i \in N$ and a school $x >_i^{g^*} a_i^*$. Since g^{*S} is ascending, this means that x is never offered to i in g^{*S}, which in turn means that x is not in the core possibility for i at any stage of g^{*S}. Thus, at every stage in g^{*S}, the cutoffs for x at that stage are higher than the value of the eligibility criterion for x at that stage, and thus x never enters the possibility set for i under step-cutoffs. In particular, there is no outcome (a, ϕ^{g^*}) with $a_i = x$ that is justifiable according to these step-cutoffs. Since i and x are arbitrary, this is always true. Thus, $\mathcal{E}^*$ is strongly verifiable through communication. ∎

B   Appendix B: Additional Experimental Results

In this subsection, we turn to notions of strategy, stability, and efficiency observed in the experiment. Our main observations are that:

1. *Transparency is not necessarily aligned with reporting preferences truthfully.*

2. *There is a significant increase in truthfulness between the first five and last five rounds of the experiment in all treatments. In all rounds, the proportion of truthful strategies is significantly higher under sequential DA than direct DA, independent of the cutoff provision.*

3. *The proportion of stable allocations and the average efficiency is significantly higher under sequential*

DA than direct DA, independent of the cutoff provision.

To simplify the language, we introduce the notion of a 'truthful strategy.' A student follows a **truthful strategy** when she submits under DirNo and DirCutoffs the truthful list of all six schools,[50] and when under SeqNo and SeqCutoffs she follows the **straightforward strategy** of applying to the best school among those that are still available (restricted either by previous rejections or by overly high intermediate cutoff grades).

Table 4: Proportions of truthful strategies

	DirNo	DirCutoffs	SeqNo	SeqCutoffs	1=2	1=3	1=4	2=3	2=4	3=4
	(1)	(2)	(3)	(4)	p-val.	p-val.	p-val.	p-val.	p-val.	p-val.
First half	12%	18%	24%	35%	0.34	0.02	0.00	0.17	0.00	0.02
Last half	33%	35%	57%	51%	0.83	0.00	0.00	0.00	0.01	0.27
Overall	23%	26%	40%	43%	0.56	0.00	0.00	0.00	0.00	0.53
First=last	0.00	0.00	0.00	0.00						

Notes: All p-values in columns 6 to 11 are for the coefficient of the corresponding treatment dummy in the probit regression of the truthfulness dummy on the treatment dummy, with the sample restricted to the treatments involved in the test. The p-values in row 5 are p-values for the dummy, which equals one for rounds 6 to 10, and equals zero for rounds 1 to 5 in probit regression of the truthfulness dummy, with the sample restricted to one treatment that is involved in the test. The standard errors of all regressions are clustered at the subject level.

Table 4 presents the proportion of truthful strategies by treatments in total and for the first five and last five rounds of the experiments separately. First, the proportion of truthful strategies is quite low, with an average of only 24% in direct DA and 42% in sequential DA. These rates are lower than typically observed in the literature (see Hakimov and Kübler (2020) for a survey) but comparable with the rates of the first 10 rounds in Bó and Hakimov (2020a) with 41%, and Koutout, der Linden, Dustan, and Wooders (2019) with 30.5%.[51] There is, however, a significant increase in the proportion of truthful strategies in all treatments. One possible reason for relatively low rates is that our subject pool has only 40% of students in economics and business and technical majors. Another reason could be that subjects might perceive that using a 'skipping' strategy is even more beneficial in our experiment than in a typical DA experiment. In particular, they might think that it is easier to identify random allocations when using skipping strategies. However, this is not consistent with the data, as we observe more correct appeals on average for random rounds under truthful strategies. Note the higher percentage of truthful strategies under SeqNo and SeqCutoffs relative to DirNo and DirCutoffs replicates findings of Bó and Hakimov

---

[50]Note that the only undominated strategy, given the information available, is to submit the full truthful list. In our setting, there is no 'safe' option, as the grades of the other students are unknown.

[51]Note that Bó and Hakimov (2020a) use the same informational setup, while Koutout, der Linden, Dustan, and Wooders (2019) get similar rates under incomplete information.

(2020a) and Klijn, Pais, and Vorsatz (2019).

These findings suggest that transparency is not necessarily aligned with optimal strategies. Moreover, while the presence of cutoffs significantly increases the proportion of correct appeal decisions, it does not affect the proportion of truthful strategies of participants. In particular, strategic simplicity and transparency appear to be two independent desiderata for centralised allocation.

We turn to the stability and efficiency of realised allocations. Row 1 of Table 5 presents the proportions of stable allocations by treatments. Note that stability can be distorted only by a subjects' suboptimal reports, as all sims are programmed to play truthfully. Row 2 of Table 5 presents the average efficiency of allocations by treatments. Efficiency is calculated as the ratio of the actual payoff to the payoff for the student-optimal-stable match.[52] As in the case of truthful strategies, the rate of stability and average efficiency are significantly higher under sequential DA than direct DA, independent of cutoff provision. The high rate of stability points to the fact that many deviations from truthful strategies were not payoff-relevant. Thus, subjects correctly anticipate their chances of being accepted in some schools that they skip.

Table 5: Proportions of stable allocations and average efficiency by treatments

	DirNo	DirCutoffs	SeqNo	SeqCutoffs	1=2	1=3	1=4	2=3	2=4	3=4
	(1)	(2)	(3)	(4)	p-val.	p-val.	p-val.	p-val.	p-val.	p-val.
Stability	70%	75%	84%	83%	0.23	0.00	0.00	0.01	0.04	0.97
Efficiency	83%	86%	91%	91%	0.30	0.00	0.00	0.01	0.02	0.89

Notes: Efficiency is calculated as the ratio of actual payoff to the payoff of the student-optimal-stable match. All the p-values in columns 6 to 11 in row 1 (row 2) are p-values for the coefficient of the dummy for the corresponding treatment in the probit (OLS) regression of the dummy for stable allocations (efficiency) on the treatment dummy, with the sample restricted to the treatments involved in the test. The standard errors of all regressions are clustered at the subject level.

C ONLINE APPENDIX: EXPERIMENT INSTRUCTIONS

C.1 COMMON INSTRUCTIONS - 1

Instructions

Welcome! This is an experiment in the economics of decision-making. If you follow the instructions carefully, you may earn a considerable amount of money. These instructions are identical for every participant. Please turn off your electronic devices. Do not communicate with each other or ask questions aloud during the experiment. If you have questions at any point, raise your hand and we will come to you

---

[52]Note that efficiency can never exceed 100%, as sims are programmed to report truthfully, and thus efficiency improvements over the stable matching like in Kesten (2010) are not possible.

and answer them.

Overview

In this experiment, we simulate an environment where students are allocated to universities.

- All of you will be making decisions as students.

- Each participant will participate in an admissions process for universities, competing with five simulated computer players.

- There are six simulated universities, namely M1, M2, L1, L2, H1, and H2.

- Every university has one seat available.

- The allocation procedure will allocate you and five simulated computer players each to one of the six places at the universities.

- The experiment consists of multiple independent rounds. Each round represents a new admissions process.

Exam Grades

- All universities admit students based on their grades in admissions exams. Each student has a grade for the Math exam and for the Language exam.

- The grades of each student (you and the computer players) for each exam are drawn independently and randomly from the set 1, 2, 3, …, 100. Each number is equally likely to be drawn. The computer will avoid ties when drawing grades. That is, each student in a group will have a different grade.

- You will learn your grades, but not the grades of computer players in your group.

- Universities M1 and M2 rank students solely based on the Math exam grade. Universities L1 and L2 rank students solely based on the Language exam grade. Universities H1 and H2 rank students based on the average between Math and Language exam grades.

- Later, when we describe the allocation procedure in detail, you will see that a higher grade may give you an advantage when competing with other students who wish to enter the same university.

Your Preferences over Universities

- You can obtain a higher university payoff if you are assigned a seat at a university you prefer more.

- As shown in the table below, your university payoff equals CHF31, CHF26, CHF21, CHF16, CHF11, or CHF6 if you hold a seat at the university ranked 1st, 2nd, 3rd, 4th, 5th, or 6th according to your preferences, respectively. If you are not assigned a seat at any university, your university payoff equals zero.

	Your university payoff
If you hold a seat at the university of your 1st Preference	CHF31
If you hold a seat at the university of your 2nd Preference	CHF26
If you hold a seat at the university of your 3rd Preference	CHF21
If you hold a seat at the university of your 4th Preference	CHF16
If you hold a seat at the university of your 5th Preference	CHF11
If you hold a seat at the university of your 6th Preference	CHF6
If you do NOT hold a seat at any university	CHF0

- You will learn your preferences over universities in the beginning of each round. You will also learn the preferences of computer students.

C.2   TREATMENT-SPECIFIC INSTRUCTIONS

C.2.1   TREATMENT DIRNO

Your Decisions before Allocation Procedure

1. You will be asked to indicate your preferences over the six universities by listing them as your 1st, 2nd, 3rd, 4th, 5th, and 6th choices.

Allocation Procedure

1. In each round, an allocation procedure will be used to allocate students to universities. The outcome of an allocation procedure depends on:

2. the ranked lists of universities submitted by you and by the other five students which are played by computers; computer players submit their lists in line with the strategy that maximises their expected payoff.

3. the exam grades of you and the other students.

Specifically, the allocation procedure follows the steps below (all the steps take place without any further interactions with the students):

The allocation procedure is implemented in the following way:

1. The mechanism sends applications from all students to the university of their top choice (the one which is stated first in the submitted list sent to the allocation mechanism).

2. Throughout the allocation process, a university can hold no more applications than its number of seats. If a university receives more applications than its capacity, then it rejects the students with the lowest relevant scores (math grade for M1, M2; language grade for L1, L2 and average grade for H1, H2) up to its capacity, and retains the remaining application(s).

3. Whenever an applicant is rejected at a university, her application is sent to the next highest university on her submitted list.

4. Whenever a university receives new applications, these applications are considered **together with the retained applications for that university**. Among the retained and new applications, the ones with the lowest relevant grades in excess of the number of the seats are rejected, while the remaining applications are retained.

5. The allocation is finalised when no more applications can be rejected. Each participant is assigned a seat at the university that holds his/her application at the end of the process.

Your Decisions after Allocation Procedure

The allocation procedure will determine your assignment (the university to which you are admitted). However, **there is a random chance that your assignment in a round is not determined according to the procedure but is determined randomly.** The probability of this happening is 50%, and is determined independently for each round.

After learning your assignment, you have a chance to submit an appeal if you think that your assignment was not determined by the procedure described above. The appeal is costly and **costs 6 CHF**. If you submit an appeal, and your assignment was indeed determined randomly and not according to the procedure described above, your appeal will be deemed correct, and your earnings for the round will be CHF40. If, however, your assignment was determined by the procedure, thus, your appeal will be deemed incorrect and you will keep your earnings from the assignment minus 6 CHF for the cost of the appeal.

Thus, you have to decide whether to accept your assignment or appeal.

Now we illustrate how the procedure works with an example.

An Example of the Allocation Procedure

Example:

In order to understand the mechanism better, let us go through an example together. If you have any questions about any step of the allocation procedure please feel free to ask at any point.

There are six students (ID numbers from 1 to 6) on the market, and three universities (University M1, University L1, and University H1) with two seats in each university.

Students have the following grades in their exams:

	Student1	Student2	Student3	Student4	Student5	Student6
Math	80	90	60	90	70	40
Language	50	20	80	30	76	82
Average	65	55	70	60	73	61

University M1 ranks students based on the Math grade only, University L1 grades students based on the Language grade only and University H1 ranks students based on the average of the two grades.

Students submit the following school rankings in their decision sheets:

Student ID	1	2	3	4	5	6
Top choice	L1	H1	M1	H1	H1	M1
Middle choice	H1	M1	H1	L1	M1	H1
Last choice	M1	L1	L1	M1	L1	L1

This allocation method consists of the following steps:

Step 1.

Students 3 and 6 apply for a seat at M1. University M1 has two seats available for allocation and two applicants, thus students 3 and 6 are retained at University M1.

Student 1 applies to University L1. University L1 has two seats and only one applicant, thus Student 1 is retained at University L1.

Students 2, 4, and 5 apply for University H1, but it has only two seats available, thus one of the applicants must be rejected. University H1 ranks students based on average grade for Math and Language: Student 2 has an average grade of 55, Student 4 has 60, and Student 5 has 73. Among the applicants Student 2 has the lowest average grade, thus Student 2 is rejected, and students 4 and 5 are retained at University H1.

	Retained students in the beginning of the round	Applications in the step	Rejected students
University M1	-	3,6	-
University L1	-	1	-
University H1	-	2,4,5	2

Step 2.

Student 2 is the only student who was rejected in the previous step. She applies to her second choice—University M1. Now University M1 considers Student 2 together with the retained students who applied to University M1 in the previous step—students 3 and 6. So the university has three applications for two seats, thus one of the applicants must be rejected. University M1 ranks students based on Math grades: Student 2 has a Math grade of 90, Student 3 has 60, and Student 6 has 40. Student 6 has the lowest Math grade among the applicants, thus Student 6 is rejected from University M1, while students 2 and 3 are retained.

	Retained students in the beginning of the round	Applications in the step	Rejected students
University M1	3,6	2	6
University L1	1	-	-
University H1	4,5	-	-

Step 3.

Student 6 applies to University H1. So the university has three applications for two seats, thus one of the applicants must be rejected. University H1 ranks students based on the average grade: Student 4 has an average grade of 60, Student 5 has 73, and Student 6 has 61. Student 4 has the lowest average grade among applicants and is thus rejected from University H1.

	Retained students in the beginning of the round	Applications in the step	Rejected students
University M1	2,3	-	-
University L1	1	-	-
University H1	4,5	6	4

Step 4.

Student 4 applies for University L1. Thus, there are 2 applications for two seats at University L1. No one is rejected. All current retained applications are finalised.

	Retained students in the beginning of the round	Applications in the step	Rejected students
University M1	2,3	-	-
University L1	1	4	-
University H1	5,6	-	-

Thus, the final allocation is as follows: University M1—students 2, 3; University L1- students 1, 4; University H1—students 5, 6. This is summarised in the following table:

Finalised assignments	University M1	University L1	University H1
seat 1	2	1	5
seat 2	3	4	6

C.2.2 TREATMENT DIRCUTOFFS

Your Decisions before Allocation Procedure

1. You will be asked to indicate your preferences over the six universities by listing them as your 1st, 2nd, 3rd, 4th, 5th, and 6th choices.

Allocation Procedure

1. In each round, an allocation procedure will be used to allocate students to universities. The outcome of an allocation procedure depends on:

2. The ranked lists of universities submitted by you and by the other five students, which are played by computers; Computer players submit their lists in line with the strategy that maximises their expected payoff.

3. the exam grades of you and the other students.

Specifically, the allocation procedure follows the steps below (all the steps take place without any further interaction with the students):

The allocation procedure is implemented in the following way:

1. The mechanism sends applications from all students to the university of their top choice (the one which is stated first in the submitted list sent to the allocation mechanism).

2. Throughout the allocation process, a university can hold no more applications than its number of seats. If a university receives more applications than its capacity, then it rejects the students with the lowest relevant scores (the math grade for M1, M2; the language grade for L1, L2, and the average grade for H1, H2) up to its capacity, and retains the remaining application(s).

3. Whenever an applicant is rejected at a university, her application is sent to the next highest university on her submitted list.

4. Whenever a university receives new applications, these applications are considered **together with the retained applications for that university**. Among the retained and new applications, the ones with the lowest relevant grades in excess of the number of seats are rejected, while the remaining applications are retained.

5. The allocation is finalised when no more applications can be rejected. Each participant is assigned a seat at the university that holds his/her application at the end of the process.

After the procedure is run, you will learn your placement and all intermediate cutoff grades - the minimum corresponding grades of retained students of all universities at each step of the procedure.

Your Decisions after Allocation Procedure

The allocation procedure will determine your assignment (the university to which you are admitted). However, **there is a random chance that your assignment in a round is not determined according to the procedure but is determined randomly.** The probability of this happening is 50% and is determined independently for each round.

After learning your assignment, you have a chance to submit an appeal if you think that your assignment was not determined by the procedure described above. The appeal is costly and **costs 6 CHF**. If you submit an appeal, and your assignment was indeed determined randomly and not according to the procedure described above, your appeal will be deemed correct, and your earnings for the round will be 40 CHF. If, however, your assignment was determined by the procedure, thus, your appeal will be deemed incorrect and you will keep your earnings from the assignment minus 6 CHF for the cost of the appeal. **Thus, you have to decide whether to accept your assignment or appeal.**

Note that in the case of your assignment being determined randomly, the table of cutoff grades is also adjusted, such that your grade affects the cutoff in the university of your random assignment. Moreover, a random cutoff is generated at the university of your true allocation.

Now we illustrate how the procedure works with an example.

An Example of the Allocation Procedure

Example:

In order to understand the mechanism better, let us go through an example together. If you have any questions about any step of the allocation procedure please feel free to ask at any point.

There are six students (ID numbers from 1 to 6) on the market, and three universities (University M1, University L1, and University H1) with two seats in each university.

Students have the following grades in their exams:

	Student1	Student2	Student3	Student4	Student5	Student6
Math	80	90	60	90	70	40
Language	50	20	80	30	76	82
Average	65	55	70	60	73	61

University M1 ranks students based on the Math grade only, University L1 grades students based on the Language grade only and University H1 ranks students based on the average of the two grades.

Students submitted the following school rankings in their decision sheets:

Student ID	1	2	3	4	5	6
Top choice	L1	H1	M1	H1	H1	M1
Middle choice	H1	M1	H1	L1	M1	H1
Last choice	M1	L1	L1	M1	L1	L1

This allocation method consists of the following steps:

Step 1.

Students 3 and 6 apply for a seat at M1. University M1 has two seats available for allocation and two applicants, thus students 3 and 6 are retained at University M1.

Student 1 applies to University L1. University L1 has two seats and only one applicant, thus Student 1 is retained at University L1.

Students 2, 4, and 5 apply for University H1, but it only has two seats available, thus one of the applicants must be rejected. University H1 ranks students based on the average grade for Math and Language: Student 2 has an average grade of 55, Student 4 has 60, and Student 5 has 73. Among the applicants, Student 2 has the lowest average grade, thus Student 2 is rejected, and students 4 and 5 are retained at University H1.

	Retained students in the beginning of the round	Applications in the step	Rejected students
University M1	-	3,6	-
University L1	-	1	-
University H1	-	2,4,5	2

The minimum corresponding grades of retained students of all universities at step 1 are:

University	M1	L1	H1
Step 1	40	0*	60

* Note, that if a university has a free seat the minimum accepted cutoff grade is zero.

Step 2.

Student 2 is the only student who was rejected in the previous step. She applies to her second choice—University M1. Now University M1 considers Student 2 together with the retained students who applied to University M1 in the previous step—students 3 and 6. So the university has three applications for two seats, thus one of the applicants must be rejected. University M1 ranks students based on the Math

grades: Student 2 has a Math grade of 90, Student 3 has 60 and Student 6 has 40. Student 6 has the lowest Math grade among the applicants, thus Student 6 is rejected from University M1, while students 2 and 3 are retained.

	Retained students in the beginning of the round	Applications in the step	Rejected students
University M1	3,6	2	6
University L1	1	-	-
University H1	4,5	-	-

The minimum corresponding grades of retained students of all universities at step 2 are:

University	M1	L1	H1
Step 2	60	0	60

Step 3.

Student 6 applies to University H1. So the University has three applications for two seats, thus one of the applicants must be rejected. University H1 ranks students based on the average grade: Student 4 has an average grade of 60, Student 5 has 73, and Student 6 has 61. Student 4 has the lowest average grade among applicants and is thus rejected from University H1.

	Retained students in the beginning of the round	Applications in the step	Rejected students
University M1	2,3	-	-
University L1	1	-	-
University H1	4,5	6	4

The minimum corresponding grades of retained students of all universities at step 3 is:

University	M1	L1	H1
Step 3	60	0	61

Step 4.

Student 4 applies for University L1. Thus, there are two applications for two seats at University L1. No one is rejected. All current retained applications are finalised.

	Retained students in the beginning of the round	Applications in the step	Rejected students
University M1	2,3	-	-
University L1	1	4	-
University H1	5,6	-	-

The minimum corresponding grades of retained students of all universities at step 4 are:

University	M1	L1	H1
Step 4	60	30	61

Thus, the final allocation is as follows: University M1—students 2, 3; University L1—students 1, 4; University H1—students 5, 6. This is summarised in the following table:

Finalised assignments	University M1	University L1	University H1
seat 1	2	1	5
seat 2	3	4	6

After the procedure, each student is informed of her placement and the following table of the minimum corresponding grades of retained students of all universities at each step.

University	M1	L1	H1
Step 1	40	0*	60
Step 2	60	0	60
Step 3	60	0	61
Step 4	60	30	61

C.2.3 Treatment SeqNo

Allocation Procedure

1. In each round, an allocation procedure will be used to allocate students to universities. The outcome of an allocation procedure depends on:

2. The choices that you and five computer players will make during the procedure. The computer players will make choices that maximise their expected payoff.

3. The admission <u>exam grades</u> of you and the other students.

The allocation procedure is implemented in the following way:

1. In the first step, each student applies to one of the universities.

   (a) Throughout the allocation process, a university can hold no more applications than its number of seats. If a university receives more applications than its capacity, then it rejects the students with the lowest relevant scores (Math grade for M1, M2; Language grade for L1, L2 and average grade for H1, H2) up to its capacity, and retains the remaining application(s).

2. At the end of each step, each student is informed about whether her application was rejected or retained.

3. In the next step, a rejected applicant can send her application to any university, except the one(s) from where she has already been rejected. If an applicant is retained at any university in the previous step, she is not active at this step and does not act.

4. Whenever a university receives new applications, these applications are considered **together with the retained applications for that university**. Among the retained and new applications, the ones with the lowest relevant grades in excess of the number of seats are rejected, while the remaining applications are retained.

5. Steps 3 and 4 are repeated until the allocation is finalised. The allocation is finalised when no more applications are rejected. Each participant is assigned a seat at the university that holds his/her application at the end of the process, and is unassigned if her application is not held at any university.

Your Decisions after Allocation Procedure

The allocation procedure will determine your assignment (the university to which you have been admitted). However, **there is a random chance that your assignment in a round is not determined according to the procedure but is determined randomly.** The probability of this happening is 50% and is determined independently for each round.

After learning your assignment, you have a chance to submit an appeal if you think that your assignment was not determined by the procedure described above. Submitting an appeal **costs 6 CHF**. If you submit an appeal, and your assignment was indeed determined randomly and not according to the procedure described above, your appeal will be deemed correct, and your earnings for the round will be CHF 40. If, however, your assignment was determined by the procedure, your appeal will be deemed incorrect and you will keep your earnings from the assignment minus 6 CHF for the cost of the appeal.

Thus, you have to decide whether to accept your assignment or appeal.

Now we illustrate how the procedure works with an example.

An Example of the Allocation Procedure

In order to understand the mechanism better, let us go through an example together. If you have any questions about any step of the allocation procedure please feel free to ask at any point.

There are six students (ID numbers from 1 to 6) on the market, and three universities (University M1, University L1, and University H1) with two seats in each university.

Students have the following grades in their exams:

	Student1	Student2	Student3	Student4	Student5	Student6
Math	80	90	60	90	70	40
Language	50	20	80	30	76	82
Average	65	55	70	60	73	61

University M1 ranks students based on the Math grade only, University L1 grades students based on the Language grade only and University H1 ranks students based on the average of the two grades.

This allocation method consists of the following steps:

Step 1.

Students took the following decisions about their application: Students 3 and 6 apply to M1, Student 1 applies to University L1 and students 2, 4 and 5 apply to University H1.

Students 3 and 6 apply for a seat at M1. University M1 has two seats available for allocation and two applicants, thus students 3 and 6 are retained at University M1.

Student 1 applies to University L1. University L1 has two slots and only one applicant, thus Student 1 is retained in University L1.

Students 2, 4, and 5 apply for University H1, but it has only two slots available for allocation, thus one of the applicants must be rejected. University H1 ranks students based on the average grade for Math and Language: Student 2 has average grade of 55, Student 4 has 60, and Student 5 has 73. Among the applicants, Student 2 has the lowest average grade, thus Student 2 is rejected, and students 4 and 5 are retained at University H1.

	Retained students in the beginning of the round	Applications in the step	Rejected students
University M1	-	3,6	-
University L1	-	1	-
University H1	-	2,4,5	2

Step 2.

Student 2 is the only student who was rejected in the previous step, thus, she is the only one who is active at this step.

She decides to apply to University M1.

Now University M1 considers Student 2 together with the retained students who applied to University M1 in the previous step—students 3 and 6. So, the university has three applications for two slots, thus one of the applicants must be rejected. University M1 ranks students based on the Math grade: Student 2 has Math grade of 90, Student 3 has 60, and Student 6 has 40. Student 6 has the lowest grade among the applicants, thus Student 6 is rejected from University M1, while students 3 and 2 are retained.

	Retained students in the beginning of the round	Applications in the step	Rejected students
University M1	3,6	2	6
University L1	1	-	-
University H1	4,5	-	-

Step 3.

Student 6 is the only student who was rejected in the previous step, thus she is the only one who is active at this step.

Student 6 decides to apply to University H1.

University H1 considers Student 6 together with the retained students—students 4 and 5. So the university has three applications for two slots, thus one of the applicants must be rejected. University H1 ranks students based on the average grade: Student 4 has an average grade of 60, Student 5 has 73, and Student 6 has 61. Student 4 has the lowest average grade among applicants and is thus rejected from University H1.

	Retained students in the beginning of the round	Applications in the step	Rejected students
University M1	2,3	-	-
University L1	1	-	-
University H1	4,5	6	4

Step 4.

Student 4 is the only student who was rejected in the previous step, thus she is the only one who is active at this step.

Student 4 decides to apply for University L1.

University H1 considers Student 4 together with the retained students—Student 1. Thus, there are two applications for two seats at University L1. No one is rejected. All current retained allocations are finalised.

	Retained students in the beginning of the round	Applications in the step	Rejected students
University M1	2,3	-	-
University L1	1	4	-
University H1	5,6	-	-

Thus, the final allocation looks as follows: University M1—students 2, 3; University L1—students 1, 4; University H1—students 5, 6. This is summarised in the following table:

Finalised assignments	University M1	University L1	University H1
seat 1	2	1	5
seat 2	3	4	6

C.2.4 Treatment SeqCutoffs

Allocation Procedure

1. In each round, an allocation procedure will be used to allocate students to universities. The outcome of an allocation procedure depends on:

2. The choices that you and five computer players will make during the procedure. The computer players will make choices that maximise their expected payoff.

3. The admission <u>exam marks</u> of you and the other students.

The allocation procedure is implemented in the following way:

1. In the first step, each student applies to one of the universities.

    (a) Throughout the allocation process, a university can hold no more applications than its number of seats. If a university receives more applications than its capacity, then it rejects the students with the lowest relevant scores (Math grade for M1, M2; Language grade for L1, L2 and average grade for H1, H2) up to its capacity, and retains the remaining application(s).

2. At the end of the step, each student is informed about whether her application was rejected or retained. **Moreover, the minimum corresponding grades of the retained students of all universities (the 'cutoffs') are publicly announced.**

3. In the next step, a rejected applicant can send the application to any university, except the one(s) from which she has already been rejected. If an applicant is retained at any university in the previous step, she is not active at this step and does not act.

4. Whenever a university receives new applications, these applications are considered **together with the retained applications for that university**. Among the retained and new applications, the ones with the lowest relevant grades in excess of the number of the seats are rejected, while the remaining applications are retained. All students see the result of the step. Each university publishes the minimum corresponding grade of retained students.

5. Steps 3 and 4 are repeated until the allocation is finalised. The allocation is finalised when no more applications are rejected. Each participant is assigned a seat at the university that holds his/her application at the end of the process, and is unassigned if her application is not held at any university.

Your Decisions after Allocation Procedure

The allocation procedure will determine your assignment (the university to which you are admitted). However, **there is a random chance that your assignment in a round is not determined according to the procedure but is determined randomly.** The probability of this happening is 50% and is determined independently for each round.

After learning your assignment, you have a chance to submit an appeal if you think that your assignment was not determined by the procedure described above. Submitting an appeal **costs 6 CHF**. If you submit an appeal, and your assignment was indeed determined randomly and not according to the procedure described above, your appeal will be deemed correct and your earnings for the round will be CHF 40. If, however, your assignment was determined by the procedure, thus your appeal will be deemed incorrect and you will keep your earnings from the assignment minus 6 CHF for the cost of the appeal.

Thus, you have to decide whether to accept your assignment or appeal.

Note that in the case of your assignment being determined randomly, the final cutoff grades are also adjusted, such that your grade affects the cutoff in the university of your random assignment. Moreover, a random cutoff is generated at the university of your true allocation.

Now we illustrate how the procedure works with an example.

An Example of the Allocation Procedure

In order to understand the mechanism better, let us go through an example together. If you have any questions about any step of the allocation procedure please feel free to ask at any point.

There are six students (ID numbers from 1 to 6) on the market, and three universities (University M1, University L1, and University H1) with two seats in each university.

Students have the following grades in their exams:

	Student1	Student2	Student3	Student4	Student5	Student6
Math	80	90	60	90	70	40
Language	50	20	80	30	76	82
Average	65	55	70	60	73	61

University M1 ranks students based on the Math grade only, University L1 grades students based on the Language grade only and University H1 ranks students based on the average of the two grades.

This allocation method consists of the following steps:

Step 1.

Students took the following decisions about their application: Students 3 and 6 apply to M1, Student 1 applies to L1 and students 2, 4 and 5 apply to H1.

Students 3 and 6 apply for a seat at M1. University M1 has two seats available for allocation and two applicants, thus students 3 and 6 are retained at University M1.

Student 1 applies to University L1. University L1 has two seats and only one applicant, thus student 1 is retained in University L1.

Students 2, 4, and 5 apply for University H1, but it has only two seats available for allocation, thus one of the applicants must be rejected. University H1 ranks students based on average grade for Math and Language: Student 2 has an average grade of 55, Student 4 has 60, and Student 5 has 73. Among the applicants, Student 2 has the lowest average grade, thus Student 2 is rejected, and students 4 and 5 are retained at University H1.

	Retained students in the beginning of the round	Applications in the step	Rejected students	Published minimum accepted grade
University M1	-	3,6	-	40
University L1	-	1	-	0*
University H1	-	2,4,5	2	60

Note, that if a university has a free seat the minimum accepted cutoff grade is zero.

Step 2.

Student 2 is the only student who was rejected in the previous step, thus she is the only one who is active at this step.

She decides to apply to University M1.

Now University M1 considers Student 2 together with the retained students who applied to University M1 in the previous step—students 3 and 6. So the university has three applications for two slots, thus one of the applicants must be rejected. University M1 ranks students based on the Math grade: Student 2 has Math grade of 90, Student 3 has 60 and Student 6 has 40. Student 6 has the lowest grade among the applicants, thus Student 6 is rejected from University M1, while students 3 and 2 are retained.

	Retained students in the beginning of the round	Applications in the step	Rejected students	Published minimum accepted grade
University M1	3,6	2	6	60
University L1	1	-	-	0
University H1	4,5	-	-	60

Step 3.

Student 6 is the only student who was rejected in the previous step, thus she is the only one who is active at this step.

Student 6 decides to apply to University H1.

University H1 considers Student 6 together with the retained students—students 4 and 5. So the university has three applications for two slots, thus one of the applicants must be rejected. University H1 ranks students based on the average grade: Student 4 has average grade of 60, Student 5 has 73, and Student 6 has 61. Student 4 has the lowest average grade among the applicants, thus he is rejected from University H1.

	Retained students in the beginning of the round	Applications in the step	Rejected students	Published minimum accepted grade
University M1	2,3	-	-	60
University L1	1	-	-	0
University H1	4,5	6	4	61

Step 4.

Student 4 is the only student who was rejected in the previous step, thus she is the only one who is active at this step.

Student 4 decides to apply for University L1.

University H1 considers Student 4 together with the retained students—Student 1. Thus, there are two applications for two seats at University L1. No one is rejected. All current retained allocations are finalised.

	Retained students in the beginning of the round	Applications in the step	Rejected students	Published minimum accepted grade
University M1	2,3	-	-	60
University L1	1	4	-	30
University H1	5,6	-	-	61

Thus the final allocation looks as follows: University M1—students 2, 3; University L1—students 1, 4; University H1—students 5, 6. This is summarised in the following table:

Finalised assignments	University M1	University L1	University H1
slot 1	2	1	5
slot 2	3	4	6

C.3   COMMON INSTRUCTIONS (CONTD.)

Your Final Earnings

- The experiment consists of 11 rounds. Ten rounds represent the university admission game described above. Round 11 will be described in the next section.

- At the beginning of every round of the university admission game, the computer will randomly:

  – draw <u>new</u> grades, and

  – generate <u>new</u> preferences for every participant.

- Each round represents a new admission process.

- Each round has the same components: lists of universities, grades, the allocation procedure, and decision of whether to appeal.

- At the end of each of the first 10 rounds, you will observe the following information.

  – **Your university payoff, which** equals 31 CHF, 26 CHF, 21 CHF, 16 CHF, 11 CHF, or 6 CHF if you are assigned a seat at the university ranked 1st, 2nd, 3rd, 4th, 5th, and 6th according to your preferences; and equals 0 CHF if you are not assigned a seat at any university;

  – **You will know whether your appeal decision was correct or not.**

  – **Your total payoff in the round, which** equals your university payoff if you do not appeal, equals 40 CHF if you appeal and it was deemed correct, and it equals your university payoff minus 6 CHF if your appeal was deemed incorrect.

- At the end of the experiment, ONE out of 11 rounds will be randomly selected. Your payoff in that round will determine your actual earnings. 10 CHF will be added to your earnings as a show-up fee.

Final additional task in round 11

In this task you will be shown a table with eight rows each on your screen in sequential order. In each of the rows, you are given the choice between option A and B. You need to decide which of the two options you prefer for every row. Option A will always represent a certain amount of money in CHF, while option B will always represent a lottery. At the end, only one of the rows from the table will determine your earnings, but you do not know in advance which row it will be. Every row is drawn with the same probability. Thus, after you have made your decision in each of the row of the table, the computer will randomly determine which row determines your payoff. Afterwards, the computer will draw your earnings given your decision for one of the rows, which is either A or B. This will be your payoff for Round 11.

For example, consider the following choice (each row of the table will have a similar choice):

Option A: 100% of 15 CHF; Option B: 50% of 25 CHF and 50% of 10 CHF

You will be asked to choose an option (A or B). If you specify option A, and this row is selected, your payoff for Round 11 is 15 CHF. If you specify Option B and this row is selected, the computer will pay out the lottery and your payoff for Round 11 will be either 25 CHF or 10 CHF.

After Round 11 you will be informed of the payoff for the round.