Transparency in Centralised Allocation: Theory and Experiment*

Rustamdjan Hakimov† and Madhav Raghavan‡

18 November 2020

Abstract

Many algorithmic allocation mechanisms suffer from a verifiability problem: participants cannot check if their assignments are correct. This problem is compounded if there are suspicions that the designer has deviated from the true allocation. We formalise these concerns and propose solutions in an information-based framework. A participant's assignment is 'verifiable' by her if any other assignment contradicts her information. A stronger requirement is 'transparency', where the designer cannot deviate from the true allocation without being detected. We show how the communication of 'terminal-cutoffs' and the use of 'predictable' multi-stage mechanisms each provide information to participants that verifies their assignments. Even though the information from predictable mechanisms and terminal-cutoffs can each be manipulated by a dishonest designer without detection, in our main result we show that they nevertheless achieve transparency if used together. We suggest transparent environments for use in school admissions, single-object auctions and house allocation. We support the effectiveness of our solutions via a school admissions laboratory experiment.

JEL classification: C78, C73, D78, D82.

Keywords: Mechanism design, information, designer incentive-compatibility, dynamic mechanisms, cutoffs, school admissions experiment.

---

*We thank Tommy Andersson, Inacio Bó, Somouaoga Bonkoungou, Battal Doğan, Matthew Jackson, Bettina Klaus, Dorothea Kübler, François Maniquet, Alex Nesterov, Parag Pathak, Al Roth, Colin Sullivan, Camille Terrier, Ashutosh Thakur, William Thomson, and seminar and conference participants in Bern, Berlin, Budapest, Gothenburg, Lausanne, Lisbon, Málaga, Raleigh, Stanford, York and Zürich. We are grateful to Nina Bonge and Fidel Petros for help conducting experiments, and to Jennifer Rontganger and Christopher Eyer for copyediting. Hakimov and Raghavan acknowledge financial support from the Swiss National Science Foundation (projects #100018_189152 and #100018_162606, respectively).

†Department of Economics, Université de Lausanne, Switzerland and WZB Social Science Center Berlin. • rustamdjan.hakimov@unil.ch.

‡Department of Economics, Université de Lausanne, Switzerland • madhav.raghavan@unil.ch.

1 Introduction

Many allocation mechanisms suffer from a verifiability problem: they do not allow participants to check whether their assignments are correct. For instance, public school systems around the world use mechanisms to compute assignments of students to schools, based on information on their preferences and scores that is seldom made public. Without access to this information, students have no way of ascertaining whether their assigned school was the best they could get. This can lead to doubts about the legitimacy of their assignments, or even dissatisfaction with the admissions process itself.

Indeed, perceptions of non-transparency have limited the adoption of normatively desirable admissions mechanisms. The New Orleans Recovery School District in the US ditched the Pareto-efficient and strategy-proof top-trading-cycles (TTC) mechanism after just one year of use, likely because "TTC was perceived as difficult for participants to understand... why a child was assigned a seat over a child who was not" (Abdulkadiroğlu, Che, Pathak, Roth, and Tercieux, 2017). In France, an overhaul of the university admissions system followed protests against non-transparency of the existing system, based on the stable deferred acceptance (DA) mechanism.[1] The President of France acknowledged these concerns:

> We are using artificial intelligence to organise the access to schools for our students. That puts a lot of responsibility on an algorithm. A lot of people see it as a black box, they don't understand how the student selection process happens... [Y]ou have to create the conditions of... full transparency... [Otherwise] people will eventually reject this innovation.[2]

Non-transparency has also been the subject of much recent debate in auctions. In 2019, Google changed its Ad Manager auction mechanism from a second-price auction to a first-price auction, in order to "provide additional auction transparency to both publishers and advertisers" about how their bids were used and ad allocations were determined.[3] Similarly, the US Federal Communications Commission switched to using open auctions for spectrum in order to allow "bidders and other interested parties [to] verify that the rules are followed" (Cramton and Schwartz, 2000).

Intuitively, an assignment is 'verifiable' by a participant if she can eliminate the possibility of any other assignment. In a novel information-based framework, our first contribution is to identify two simple solutions to the verifiability problem, based on acquired information that participants can use to verify

[1] The proposed law: LegiFrance, 8 March 2018. https://www.legifrance.gouv.fr/affichTexte.do?cidTexte =JORFTEXT000036683777 (last accessed 15 November, 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 .

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

[3] "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 October 2020).

their assignments. In particular, we show how suitable public communications from the designer[4] allow participants to verify their assignments (Proposition 2 and Proposition 3) as do 'predictable' multi-stage mechanisms through participants' individualised and private 'experience' of them[5] (Proposition 1). This provides a formal basis for two ad hoc proposals for improving transparency in applications. On the one hand, our communications are based on 'cutoffs' (the minimum score required to be eligible for a school), which have been proposed as a way for students to check their eligibility at their assigned school, but also their ineligibility at preferred schools (Azevedo and Leshno, 2016; Dur and Morrill, 2018; Leshno and Lo, 2020). On the other hand, we use the experience of predictable multi-stage mechanisms to formalise why ascending auctions are perceived to be more transparent than sealed-bid equivalents (Cramton, 1998; Ausubel, 2004), essentially as the experience of multiple interactions with the mechanism allow participants to understand whether they have won or lost the auction even before learning the final outcome.

However, the information acquired by participants through these two sources is governed by the designer through her choice of communication and the mechanism, respectively. This leaves our solutions vulnerable to manipulation by a 'dishonest' designer who wishes to reach a different outcome. Concerns about the designer's motivations are not merely hypothetical: corruption was the primary reason behind the redesign of public procurement auctions for school supplies (Argentina) and telecommunications (Colombia) (Boehm and Olaya, 2006). Even if the designer is not actually corrupt, the suspicion that allocations might be manipulated can hinder the use of these mechanisms in practice, as evidenced by the reasons behind recent changes in auction formats at many major online ad exchanges.[6] In our second contribution, we take these concerns into account and propose a formal definition of a 'transparent' allocation environment, where the designer cannot cheat participants and produce a different allocation without the deviation being detectable by someone.[7]

Our main result is to show that transparency can be achieved by *simultaneously* utilising a predictable mechanism and communicating terminal-cutoffs (Theorem 1).[8] Remarkably, though these two solutions

[4]The public and common nature of communication from the designer is in contrast to the individualised and private communication between the auctioneer and bidders in Akbarpour and Li (2020).

[5]Participants experience the mechanism by observing the alternatives they are offered to choose from, and the actions they take in response. In principle, multi-stage mechanisms allow more opportunities for these observations, and so convey richer experiential information.

[6]"Big Changes Coming To Auctions, As Exchanges Roll The Dice On First-Price," 5 September 2017, https://www.adexchanger.com/platforms/big-changes-coming-auctions-exchanges-roll-dice-first-price/ (last accessed 10 November 2020).

[7]We take no position on the incentives of the designer, and do not model them directly. Instead, we explore her ability – for whatever reason – to change the allocation undetected. Detection by a participant involves her being able to tell if the designer used a different mechanism from what was promised, or otherwise communicated false information. In a sense, even if the designer does not deviate in this way, transparency is valuable as a 'full commitment' device to the true allocation.

[8]Our result holds under some mild restrictions on mechanisms that are nevertheless satisfied in a wide variety of applica-

provide little information to participants by themselves, and can moreover be individually falsified by the designer without detection, their presence together completely ties the designer's hands. This result has significant practical implications: for many applications, predictable mechanisms are easy to implement, and terminal-cutoffs are easy to compute.

Our main application is the widely-used 'student-proposing' DA (SPDA) procedure in school admissions. We confirm the intuition that the classical implementation of SPDA – with no communication and a one-stage mechanism that elicits students' preferences as rank-order lists – is not verifiable through either experience or communication.[9] We propose a transparent environment for SPDA (Corollary 1). We also apply our results to other canonical market design settings, finding transparent environments for 'serial dictatorship' (SD) in house allocation (Corollary 2), and the second-price and first-price auctions for single objects (Corollary 3 and Corollary 4).

Our theory is predicated on participants' using their acquired information to verify their assignments. An additional contribution in this paper is to provide an empirical measure of transparency, which we use to test whether this is indeed the case. We do this via a laboratory experiment based on the SPDA mechanism in school admissions. In a novel design feature, participants are told that assignments are randomly determined (in particular, not by the explained mechanism) half of the time. After learning their assignments, participants 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 assignment is random. Thus the correctness of decisions of whether or not to appeal provides a sense of whether participants can correctly spot when their assignments are random, and provides an experimental measure of the transparency of the environment.

We use four environments that vary the information provided to participants as treatments in the experiment. The first environment uses a one-stage SPDA mechanism with no cutoffs (DirNo). This is the classical (unverifiable) implementation of SPDA. In the second treatment, we provide participants experiential information by using a predictable (multi-stage) SPDA mechanism with no cutoffs (SeqNo). This environment is verifiable only through experience. In the third treatment, we use a one-stage SPDA mechanism, but provide cutoff information (DirCutoffs). This environment is verifiable only through communication. Finally, we combine the two by using a predictable SPDA mechanism and communicating cutoffs (SeqCutoffs), which is verifiable through experience and communication.

The proportion of correct decisions of whether or not to appeal is the highest in the SeqCutoffs envi-

tions. We assume that the designer is bound by 'non-wastefulness', in that she cannot leave a seat in a school unfilled, or an object unsold, if some participant wants it, and that the promised mechanism is 'monotonic in offers'.

[9]Another criticism of DA – though unrelated to our current exercise – is growing field evidence of the failure of truthful reporting as a dominant strategy in mechanisms (Hassidim, Marciano, Romm, and Shorrer, 2017; Shorrer and Sóvágó, 2017; Hassidim, Romm, and Shorrer, 2018; Rees-Jones, 2018). Here we abstract from strategic concerns for participants.

ronment (83%), as predicted. Moreover, information from experience (SeqNo, 64%) and communication (DirCutoffs, 65%) each has a positive and significant effect on correct decisions relative to the base case (DirNo, 52%), in which participants do no better than chance. This furthers our understanding of how each of our proposed solutions can improve transparency, and moreover that they are most powerful when used together. We also believe the experimental setting is of independent interest, both as an empirical tool to compare the relative transparency of potential environments in an application, and also to check whether a proposed environment improves observed transparency or not.

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 general centralised allocation settings.

There are, however, several studies on related issues. A close relation to our paper 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. Woodward (2020) studies different auction formats in detail, including multi-unit auctions, while we restrict our attention to single-unit auctions. However, our paper embeds such auctions in a general allocation setting that also covers other applications. Auditability is similar to verifiability in our paper, though it is weaker than transparency. In particular, the only constraint on public information disclosed by the designer is that it should not contradict any bidder's own information. We too assume this constraint, but we show that the designer can sometimes communicate 'false' public information that justifies an outcome different from that produced by the promised mechanism, yet does not contradict any bidder's information. Our transparency notion precludes this additional possibility: in a transparent environment, any false information published by the designer is outcome-irrelevant. Auditability is also studied by Pycia and Ünver (2020) in the context of Arrovian social welfare functions, though they do not consider questions of transparency.

Designer-incentive compatibility in auctions is also 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 and not of the environment. For instance, Akbarpour and Li (2020) show that the first-price sealed-bid auction is credible. But we show that it is not transparent unless the winning price is also (publicly) announced. Naor, Pinkas, and Sumner (1999) also consider the incentives of an 'auction issuer' who facilitates the computation of auctions. Bó and Chen (2019) and Boczoń and Wilson (2020) consider random mechanisms and the legitimacy benefits from conducting lotteries in public, such as by drawing balls from urns. However, we restrict our attention in this paper to deterministic mechanisms.

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; Troyan, 2019). Strategic simplicity is also studied in Börgers and Li (2019), where choices require limited strategic sophistication for participants. 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. However, we abstract from strategic issues of truthful revelation of preferences; transparency is an orthogonal requirement to strategy-proofness.

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 eligible for another. Cutoffs are a central feature of the communications from the designer in our model.

Our work also contributes to the literature on multi-stage mechanisms. The role of ascending auctions in providing relevant information to participants is studied for single units in Cramton (1998) and multiple units in Ausubel (2004, 2006). We support these arguments by formalising the role of experiential information. Kagel, Harstad, and Levin (1987) compare the performance of the second-price sealed bid and the ascending clock auction in the lab, showing that the participant tend to play closer to their dominant strategy in the ascending clock auction. Similar results are established for multi-unit auctions (Kagel and Levin, 2001). 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 (2020a) do so 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. While sequential implementations might have different motivations in this literature, we show that some of them can be used to improve transparency.

Another strand of the literature describes and empirically evaluates multi-stage mechanisms recently adopted by policymakers for college admissions in several countries, like Germany, Inner Mongolia (China), and Brazil (Gong and Liang, 2016; Bó and Hakimov, 2018; Grenet, He, and Kübler, 2019). The mechanisms in Inner-Mongolia and Brazil make information about students' chances at different schools available even before the final allocation is reached, emphasising the extensive use of intermediate cutoff communication in practice. Dur, Hammond, and Morrill (2018) evaluate strategic behaviour in a sequential mechanism used for allocating school seats in the Wake County Public School System. While in this mechanism the submission of preferences is done in one stage, students nevertheless observe some statistics about the actions taken by students who participated previously.

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

2 The Model

2.1 Preliminaries

Throughout, concepts are presented in ***boldface***, while definitions are underlined. There is a finite set of ***participants*** N and a finite set of ***alternatives*** X, which includes an unassigned or outside option, denoted $\emptyset \in X$. Each alternative $x \in X$ has a ***capacity*** $q_x \in \{1, \ldots, |N|\}$ which determines the maximum number of participants it can accommodate. We assume $q_\emptyset = |N|$ (the unassigned option has unrestricted capacity). An ***allocation*** is a vector $a \in X^N$ with $a_i \in X$ denoting the ***assignment*** of $i \in N$. An alternative $x \in X$ is unfilled (to capacity) in an allocation $a \in X^N$ if $|\{i \in N \mid a_i = x\}| < q_x$, and is overfilled in a if $|\{i \in N \mid a_i = x\}| > q_x$. Each application is associated with a set of ***feasible allocations***, denoted $\mathcal{A} \subset X^N$. Feasible allocations respect capacity constraints (no overfilled alternatives), and any other restrictions imposed in the specific application.

Preferences for participant $i \in N$ are given by a reflexive, complete and transitive binary relation $\succeq_i$ over alternatives X. An alternative $x \in X$ is acceptable to i under $\succeq_i$ if $x \succeq_i \emptyset$, i.e., it is at least as good as being unassigned. A ***preference profile*** is denoted $\succeq = (\succeq_i)_{i \in N}$, and the set of all admissible profiles in the particular application is $\mathcal{R}$. ***Priorities*** for alternative $x \in X$ are given by a vector $v^x \in \mathbb{R}_+^N$, such that, for any $i \in N$, $v_i^x \in \mathbb{R}_+$ is the ***score*** of participant i. This is a cardinal representation of priorities; ordinal priorities, if required, can be recovered by ranking participants in descending order of score. A ***priority profile*** is denoted $v = (v_i^x)_{i \in N}^{x \in X}$, and the set of all admissible priority profiles in the application is denoted $\mathcal{V}$. Scores in our model are fixed and immutable, thus henceforth we fix a particular priority profile v^*. However, while the designer knows the entire profile v^*, we assume each participant initially knows only her score v_i^{*x} at each alternative $x \in X$. The ***domain of allocation problems*** is $\mathcal{P} \equiv (N, X, \mathcal{A}, \mathcal{R}, \mathcal{V})$. For simplicity, we often write an ***allocation problem*** simply as $P = (\succeq, v)$.

Information is used by both the designer and participants to refine the set of possible allocation problems. In our model, ***information*** identifies a subset of allocation problems and is given by $\phi \subseteq \mathcal{P}$. We say that a problem $P \in \phi$ is compatible with ϕ. In this framework, each component of a preference profile or priority profile is information, as it refines the set of problems to the subset of all problems containing that component. For instance, if a student i's information consists of her own preferences $\succeq_i$ and scores v_i^{*x} at each school x, the set of problems compatible with her information involves all possible preferences and scores of other students, fixing $\succeq_i$ and v_i^*. If $\phi = P$ for some $P \in \mathcal{P}$, we call ϕ full information as it precisely identifies an allocation problem. Let $\Phi \equiv 2^{\mathcal{P}}$ be the ***domain of information*** in the particular application. In general, ϕ contains more information than ϕ' if $\phi \subsetneq \phi'$, i.e., it further refines the set of compatible problems.

2.2 Information Acquired by Participants

We model two ways in which participants acquire information: 'communications' from the designer, and their 'experience' of the mechanism.

The ***communication*** from the designer is information denoted $M \in \Phi$, and is determined by a ***communication protocol*** $\mathcal{M}$. We assume communications are common, e.g., via announcements on public web sites, newspapers, etc.[10] This is our only restriction. The communication protocol can be described by a 'type', for instance the 'full disclosure' protocol that communicates all the designer's information to participants, or the 'all scores' protocol that reveals all scores, etc. The type could also allow for the communication of statistics derived from the computation of assignments. Some examples include the 'winning price' protocol in auctions, the 'number of participants assigned to each school' protocol in school admissions, and so on.

The other source of information for participants is through their experience of the mechanism. We formally define a mechanism shortly, but essentially we assume that participants take part in a mechanism as if they are sitting in a closed room, with no contact with either the designer or other participants. They interface with a mechanism by being offered some alternatives, and by taking actions (i.e., reporting preferences). These offered sets and actions, along with the rules of the mechanism, convey 'experiential information' to participants. Notably, multi-stage mechanisms, in which participants interact with the

[10] While this assumption is no doubt restrictive, it is quite plausible in the applications that we consider. Akbarpour and Li (2020) consider communication between the designer and a participant that is totally private and unobservable by other participants. In effect, we take the opposite approach. A more complicated analysis would allow for 'semi-private' communication between the designer and participants, but this is outside the scope of our current model.

mechanism multiple times, might convey more experiential information than one-stage mechanisms. For example, suppose a participant in an ascending clock auction with price increases of \$1 intervals observes a current price of \$30 and places a bid at that price. If she is subsequently asked to bid at \$31, the rules of the auction tell her that she cannot win the auction at any price of \$30 or less. Moreover, if she understands the auction mechanism, she infers that at least one other participant has placed a bid of \$30, which refines the set of allocation problems compatible with her information. These detailed inferences are generally not possible in a one-stage mechanism where participants submit a single bid and do not interact with the mechanism again.

By our closed room scenario, the information that this participant can infer from these observations is individualised and personal. In general, for a participant $i \in N$, this channel of information is called ***experiential information*** for i. For a mechanism g, this is denoted $\epsilon_i^g \in \Phi$. This requires us to be more precise about participants interact with the mechanism. In particular, we now address the challenges of formalising how mechanisms depend on information, how they convey experiential information to participants, and how multi-stage mechanisms are richer in this respect.

2.3 Mechanisms

Starting from Hurwicz (1973), a classical mechanism is generally represented as a function $g : \mathcal{P} \to \mathcal{A}$ which identifies a feasible allocation for each allocation problem. However, this 'static' definition hides some important details. In particular, a mechanism plays two roles: elicitation of preferences and computation of assignments. A deeper analysis of these roles is essential to our study of how to make the computation of assignments transparent.

The first role of a mechanism is to elicit private information from participants (in the form of reported preferences).[11] The static definition above assumes that the mechanism elicits full information on preferences, thus providing the designer with an allocation problem $P \in \mathcal{P}$ (since we assume the designer knows the priority profile). However, in many instances, the designer does not require full information on preferences to uniquely identify an allocation problem; it is often irrelevant how participants rank alternatives 'lower down' in their preferences. We formalise this intuition by representing a mechanism more generally as a correspondence $g : \Phi \to 2^{\mathcal{A}}$ that associates each information $\phi \in \Phi$ with a set of feasible allocations $g(\phi) \subseteq \mathcal{A}$. We assume, however (as in the classical definition), that g produces a unique feasible allocation for each full information, i.e., $g(P) = (g_i(P))_{i \in N} \in \mathcal{A}$ is uniquely determined for each $P \in \mathcal{P}$, and thus we can define $g(\phi) = \{a \in \mathcal{A} \mid g(P) = a \text{ for some } P \in \phi\}$.

A global constraint on the mechanisms we consider is 'non-wastefulness', where every alternative that

[11] The second role is the 'functional form' of the mechanism, which uses the relevant information on preferences and scores to compute a feasible allocation. We defer the analysis of this role to Section 5.

is preferred by some participant to her assignment is filled to capacity in the allocation determined by the mechanism. Formally, mechanism g is non-wasteful if, for any $P = (\succeq, v) \in \mathcal{P}$, there is no $i \in N$ and $x \in X$ such that $x \succ_i g_i(P)$ and x is unfilled in $g(P)$.[12] One justification for this assumption is that empty seats in schools, or unsold objects in an auction, are costly for the designer. Another is that such outcomes might be observable, and can lead to complaints from participants.

Two notions of eligibility are essential for our analysis. For a mechanism g, generic information $\phi \in \Phi$, and participant $i \in N$, the eligibility set of i under g for ϕ collects all assignments that are produced by g for i for problems in ϕ, i.e., $E_i^g(\phi) = \{x \in X \mid g_i(P) = x \text{ for some } P \in \phi\}$. The relation between information and eligibility sets is illustrated in Figure 1.

Figure 1: Mechanisms, information and eligibility

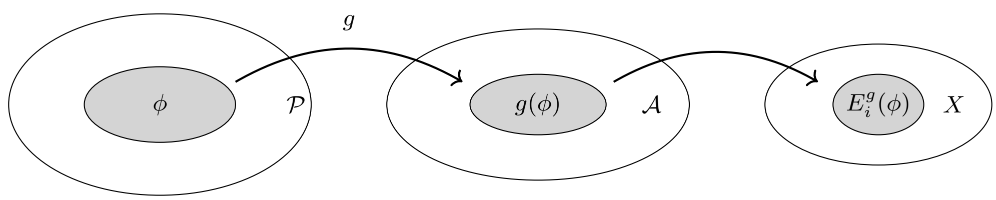

Notes: The mechanism g associates feasible allocations $g(\phi) \subseteq \mathcal{A}$ to information $\phi \in \Phi$. For a participant i, the eligibility set $E_i^g(\phi) \subseteq X$ is the set of alternatives assigned to i among allocations in $g(\phi)$.

An alternative is in the 'core eligibility' for a participant if she can 'claim' under g that alternative for some preferences compatible with the information, for any preferences of others that are also compatible with that information. Formally, given information ϕ, an alternative $x \in X$ is in the core eligibility for participant $i \in N$ (denoted $x \in C_i^g(\phi)$) if there exists $\succeq_i$ such that $g_i(P) = x$ for all $P \in \phi$ containing $\succeq_i$. The core eligibility set is always a subset of the eligibility set.

However, the classical definition of a mechanism is also silent on *how* preferences are elicited. In particular, the default 'static' assumption is that preferences are elicited in one-shot, i.e., each participant is required to submit rank-order lists (usually complete rankings) of preferences over alternatives. However, many popular mechanisms are in fact multi-stage, allowing participants to report preferences sequentially. This is the key difference, for example, between ascending auctions and the sealed-bid second-price auction. While the outcomes in both mechanisms might be the same in equilibrium, the 'dynamics' of the reporting of bids is substantially different. Capturing this intuition requires accommodating a multi-stage preference elicitation procedure in the definition of a mechanism. Our procedure has the following features.

Firstly, we assume the designer knows the priority profile v^* and needs to acquire preference information from participants to determine the allocation. We abstract from strategic issues related to the reporting

[12]Notice that non-wastefulness implies that no participant is ever assigned an alternative under g that she considers unacceptable. This is because the unassigned option has limitless capacity, by assumption. This property is variously known as 'individual rationality' or 'voluntariness'.

of preferences, and instead model reported preferences as 'actions', which involve ranking (subsets of) alternatives. Our mechanism proceeds in 'stages' (see Appendix A for a formal definition). In particular, each stage has the following components:

1. **Action:** The mechanism selects a set of active participants. Each active participant is offered a subset of alternatives. Each active participant takes an action which involves ranking at least one pair of offered alternatives.

2. **Evaluation:** The designer's information is updated to include all newly taken actions. If this information is sufficient to determine a unique feasible realised allocation, the designer tells each participant her realised assignment and the mechanism terminates. Otherwise, we go to the next stage.

We ensure that the designer acquires strictly more information in each stage, thus ensuring that mechanisms are well-defined and terminate finitely. Our definition allows for one-stage mechanisms as well as multi-stage ones. As can be seen above, mechanisms can differ across stages in three additional dimensions: (1) which participants are selected to be active in a stage; (2) which alternatives they are offered; and (3) which actions they are allowed to take. These differences might affect the experiential information they provide to participants.

3 The Verifiability Problem

An ***allocation environment*** is $\mathcal{E} = (g, \mathcal{M})$, where g is a mechanism and $\mathcal{M}$ is a communication protocol of a fixed type. The designer announces to participants that she will use a promised environment $\mathcal{E}^* = (g^*, \mathcal{M}^*)$. The true allocation in $\mathcal{E}^*$ is denoted a^*, and is determined by the procedure described in the previous section. However, when participant $i \in N$ observes that she is assigned some alternative x, how can she be sure that $x = a_i^*$? We call this the 'verifiability problem'.

In this section, we propose two intuitive ways to solve the verifiability problem by providing information to participants that allows them to eliminate the possibility of any other assignment. A participant's total information in $\mathcal{E}^*$ is given by $\epsilon_i^{g^*} \cap M^* \cap \phi_i^{*g^*}$, where $\epsilon_i^{g^*}$ is her experiential information from the mechanism g^*, M^* is the designer's communication, and $\phi_i^{*g^*}$ is her intrinsic information (her actions and scores, see Appendix A for details). Intuitively, her assignment x is 'justified' if any other assignment y would contradict her information. In our model, this corresponds to her eligibility set under her total information coinciding precisely with x, i.e., $E_i^{g^*}(\epsilon_i^{g^*} \cap M^* \cap \phi_i^{*g^*}) = x$. Thus an allocation a is justified for each participant in $\mathcal{E}^*$ if $E_i^{g^*}(\epsilon_i^{g^*} \cap M^* \cap \phi_i^{*g^*}) = a_i$ for each $i \in N$. A natural and basic transparency requirement is for the true allocation a^* to be justified in the promised environment $\mathcal{E}^*$.

Definition **1.** The promised environment $\mathcal{E}^* = (g^*, \mathcal{M}^*)$ is verifiable if the true allocation a^* in $\mathcal{E}^*$ is justified for each participant.

In a verifiable environment, participants acquire enough information such that any assignment other than their true assignments would contradict their information. By its definition, verifiability might be achieved through either experiential information $\epsilon_i^{g^*}$ ('verifiability through experience') or the designer's communication M^* ('verifiability through communication'). This unifies two increasingly popular ways in which transparency is being addressed in practice, namely, 'predictable' multi-stage mechanisms that convey rich experiential information to participants (Cramton, 1998; Ausubel, 2004), and 'terminal-cutoffs' that allow students to check their ineligibility at preferred schools (Azevedo and Leshno, 2016; Dur and Morrill, 2018; Leshno and Lo, 2020). We consider each of these in turn.

3.1 Verifiability through Experience

Consider the following three properties of a mechanism:

SS1: *Each* action is single-valued.

SS2: Participants are eligible under their experiential information for *only one* action at a time.

SS3: Each participant's assignment is the (unique) alternative in her eligibility set under their experiential information at the terminal stage of the mechanism.

A mechanism that satisfies SS1-SS3 is called a sequential single-valued mechanism (SSM) (formal definitions are in Appendix A). As a practical example, consider an ascending clock auction in which the designer raises the price in discrete intervals in set periods, as long as there are multiple bidders remaining. Each bidder's action in each stage of the auction is to either place a bid at that price or to drop out, and is thus single-valued (SS1). The auction terminates at the first price in which only one bidder remains. Notice that each bidder, even without observing others' actions, can predict her own assignment by observing her interactions with the mechanism and simply following the rules of the auction. If she chooses to drop out, then her only possible assignment is to be unassigned. On the other hand, if she bids, she will either win and pay a price corresponding to the previous period's bid (if the auction terminates at that stage), or otherwise will be asked to bid again. Thus she is eligible under experiential information for only one alternative at a time (SS2). Moreover, she gets what she expects to get at the terminal stage of the mechanism (SS3), even though she might not know the terminal stage in advance. This kind of 'predictability' is equivalent to verifiability through experience.[13]

[13] More complicated mechanisms than SSMs could still be predictable. For example, consider a mechanism in which each participant's action in each active stage consists of submitting multiple alternatives, but the rules of the mechanism convey

Proposition **1.** *The promised environment $\mathcal{E}^* = (g^*, \mathcal{M}^*)$ is verifiable through experience if and only if it is predictable.*

The proof is in Appendix B, and is based on the intuition that in a predictable mechanism, participants take actions that effectively involve one alternative at a time. Since non-wastefulness requires them to be offered the unassigned option in each stage in which they are active, any such stage could be the last stage in which they are asked to take an action. Since verifiability through experience requires each participant's eligibility set under her experiential information to be single-valued at the terminal stage, this effectively forces the condition onto *each* stage of the mechanism.

3.2 Verifiability through Communication

Our communications are based on the idea of the 'minimum value' of a eligibility criterion for a particular alternative. An 'eligibility criterion' for an alternative $x \in X$ under mechanism g^* is a statistic $\rho^{g^*,x}$ that induces a (weak) ordering over participants, i.e., for any $i, j \in N$, $\rho_i^{g^*,x} \geq \rho_j^{g^*,x}$ or vice versa. Since the mechanisms we consider always partition agents into eligible and ineligible in order to compute allocations, such an eligibility criterion always exists. However, the criterion could be simple or complicated, depending on the mechanism, and might not always be easy to compute. For instance, for some mechanisms, the eligibility criterion is just a number, e.g., bids in an auction, or students' scores at a school for the SPDA mechanism in school admissions (Azevedo and Leshno, 2016). For other mechanisms, the eligibility criterion could be more complicated and multidimensional (such as the matrix-based eligibility criterion proposed for the TTC mechanism in school admissions in Dur and Morrill (2018) and Leshno and Lo (2020)). We do not restrict the structure of the eligibility criterion further, except to stipulate that it is invoked by the mechanism g^* to determine eligibility using two conditions:

1. *Eligibility partition:* Suppose that, for information ϕ, there is an alternative $x \in X$ and two participants $i, j \in N$ such that i is eligible for x under ϕ but j is not. Then i has a higher value under the eligibility criterion for x than j. Formally, if $x \in E_i^{g^*}(\phi)$ but $x \notin E_j^{g^*}(\phi)$ then $\rho_i^{g^*,x} > \rho_j^{g^*,x}$.
2. *Monotonic eligibility:* Suppose that, for information ϕ, j is eligible for x under ϕ, and i has a higher eligibility criterion value for x. Then i is also eligible for x under ϕ. Formally, if $x \in E_j^{g^*}(\phi)$ and $\rho_i^{g^*,x} \geq \rho_j^{g^*,x}$ then $x \in E_i^{g^*}(\phi)$.

that she will be considered only for the top-ranked of these, and will otherwise be asked to submit a new list. An action is no longer necessarily single-valued, but all alternatives submitted except for the top-ranked one are essentially 'cheap talk' and are non-binding. Nevertheless, this mechanism is predictable, as each participant knows she will only be considered for the top-ranked of these alternatives in any stage. We propose a more general definition of predictability that incorporates such mechanisms in Appendix A. Our results hold for these mechanisms as well.

For an allocation $a \in X^N$ (not necessarily feasible) and an alternative $x \in X$, a 'cutoff' for x in a is the minimum eligibility criterion for x among participants assigned x in the allocation a. Formally, a cutoff for x in a is $c_x(a) = min\{\rho_i^{g^*,x} \mid a_i = x\}$.

3.2.1 Terminal-Cutoff Protocols

For an allocation $a \in X^N$, let $c(a) = (c_x(a))_{x \in X}$ be the (vector of) cutoffs. A communication protocol $\mathcal{M}^*$ is a terminal-cutoff protocol if $M^{*a} = c(a)$ for any $a \in \mathcal{A}$. Then:

Proposition **2.** *An environment $\mathcal{E}^* = (g^*, \mathcal{M}^*)$ is verifiable through communication if $\mathcal{M}^*$ is a terminal-cutoff protocol.*

Proof: Let $\mathcal{E}^* = (g^*, \mathcal{M}^*)$ be the promised environment. Let $a^* = g^*(\phi^*_{g^*})$ be the realised allocation. Let $\mathcal{M}^*$ be a terminal-cutoff protocol. Thus $M^* = c(a^*)$. Let $i \in N$ be a participant, and let $x >_i^* a_i^*$ be an alternative ranked higher in her actions. Since $\mathcal{E}^*$ is non-wasteful, x is filled in a^*. By implication $x \neq \emptyset$. By construction of cutoffs, i is ineligible for x under $c(a^*)$, and so there is no problem $P \in M^* \cap \phi_i^{*g^*}$ such that $g_i^*(P) = x$. Thus $E_i^{g^*}(M^* \cap \phi_i^{*g^*}) = a_i^*$. Since x and i were arbitrary, this is always true. Thus a^* is justified, proving that $\mathcal{E}^*$ is verifiable. ∎

4 Transparency

These solutions to the verifiability problem rely on the designer using the promised environment $\mathcal{E}^*$, thus providing each participant i with 'true' experiential information $\epsilon_i^{g^*}$ and 'true' communication M^*. But what if the designer uses a different environment $\mathcal{E} = (g, \mathcal{M})$ in order to reach a different allocation $a \neq a^*$? Could the designer do this and get away with it, in the sense that the false allocation a is nevertheless justified for each participant? The answer, unfortunately, is yes. We show through examples how an admissions designer can manipulate the experiential information for each participant through a choice of g (Subsection 6.2), or manipulate the commonly communicated information M through a choice of $\mathcal{M}$ (Subsection 6.3), to justify the false allocation a without being detected by any participant. We take no position on the incentives of the designer and do not model them directly. Instead, we ask if we can sufficiently 'tie the designer's hands' in terms of the environment she can use, so that she has no option but to reach the true allocation. Notice that an affirmative answer to this question is useful even if the designer is completely honest, as it amounts to 'full commitment' to the true allocation.

There are some natural limits on which environment the designer can use, and thus which allocation she can reach. As before, we require the used mechanism g to be non-wasteful, as a global constraint for the designer. But the change in environment from $\mathcal{E}^*$ to $\mathcal{E}$ must also be undetectable by participants. In one direction, violations can be easily detected if the used mechanism g does not resemble the promised mechanism g^* in offer sets and permitted actions (which are the parts of the mechanism observable by participants). In general, we say a mechanism g is indistinguishable from g^* if (1) permitted actions in both are the same; and (2) the 'dynamics' of offer sets in both are the same.[14] Violations can also be easily detected if the used communication protocol $\mathcal{M}$ does not resemble the the promised protocol $\mathcal{M}^*$. We say $\mathcal{M}$ is indistinguishable from $\mathcal{M}^*$ if (1) $\mathcal{M}$ and $\mathcal{M}^*$ are of the same type; and (2) the allocation a reached by g in $\mathcal{E}$ is justified if the true allocation a^* in $\mathcal{E}^*$ is justified. An environment $\mathcal{E} = (g, \mathcal{M})$ is plausible if g and $\mathcal{M}$ are indistinguishable from g^* and $\mathcal{M}^*$, respectively. In effect, the designer is constrained to use plausible environments, as any implausible environment is immediately detectable by some participant. Note also that $\mathcal{E}^*$ is trivially plausible. We can then be precise about what we mean by transparency: The promised environment $\mathcal{E}^*$ is 'transparent' if only the true allocation a^* is justified, *no matter which* plausible environment the designer actually uses.

Definition **2.** The promised environment $\mathcal{E}^*$ is transparent if a^* is the unique justified allocation in any plausible environment $\mathcal{E} = (g, \mathcal{M})$.

4.1 Main Result

Transparency entirely removes the designer's ability to change the allocation in an undetected way. If she uses an implausible environment, this can be detected directly by some participant. If she uses a plausible environment, she has to produce the true allocation a^*, since no other allocation can be justified. Intuitively, one might suppose that transparency is a greater informational burden, since it is a stronger objective than verifiability. Indeed, in Theorem 1 below, we show that the full disclosure protocol, which communicates the maximum possible information to participants, achieves transparency, since the designer cannot manipulate preferences and scores without being detected by some participant. However, we also show in Theorem 1 that transparency sometimes requires surprisingly little information to be acquired by participants. To show this, we require some additional restrictions.

Firstly, we make a behavioural assumption on participants: we assume that participants, when faced with the same offer set in g^* and in some other indistinguishable mechanism g, take the same permitted action in both. Formally, given an offer set Y, if the action in g^* for some agent $i \in N$ is $\geq_i^*$, then in any other mechanism g that is indistinguishable from g^* (i.e., with the same permitted actions), if the participant receives the same offer set Y at some stage, she takes the same action $\geq_i^*$ in that stage. We also restrict our attention to two widely-used classes of 'monotonic-offers' mechanisms, based on whether the size of offered sets of alternatives for a participant is increasing or decreasing across stages.[15] An

[14] We will be more precise about this when we introduce 'monotonic' mechanisms.

[15] A formal definition is in Appendix A. A mechanism is 'monotonically decreasing in offers' if each participant is offered,

environment $\mathcal{E} = (g, \mathcal{M})$ is monotonic-offers if g is monotonic-offers. While this no doubt restricts the class of mechanisms covered by our result, we note that most multi-stage mechanisms in our applications satisfy this property, such as many sequential implementations of the SPDA procedure in school admissions, or many multi-stage implementations of the first- or second-price auctions. Moreover, all direct non-wasteful mechanisms are trivially monotonic-offers.

We can now state our main result:

Theorem **1.** *Let $\mathcal{E}^* = (g^*, \mathcal{M}^*)$ be an environment. Then $\mathcal{E}^*$ is transparent if either:*

1. $\mathcal{M}^*$ is full disclosure.

2. $\mathcal{E}^*$ is monotonic-offers, and is verifiable through each of experience and communication.

The proof is in Appendix D. 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 necessary sequentiality of the used mechanism g (a consequence of predictability of g^* and indistinguishability of g from g^*) limits the designer in significant ways. The designer acquires information on preferences through actions in discrete stages in g. In deciding to deviate from the promised mechanism at any stage, she has to work with all possibilities of continuation actions of participants, 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 allocation. In particular, we show that there is always the risk that the deviation might lead to the final allocation being wasteful,[16] or that the communicated information no longer justifies the final allocation, thus making any deviation detectable to some participant.

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 participants with *even less* information than they acquire through predictable mechanisms and terminal-cutoffs. As such, the nature of the *minimum* information required by participants for transparency is an open question. Nevertheless, we believe our sufficiency result is useful because finding predictable mechanisms or terminal-cutoffs is straightforward in many applications, and thus transparency can be simply achieved. We demonstrate this by proposing transparent environments for use in several applications in Section 6 and Section 7.

in each stage in which she is active, all alternatives that she is eligible for under the designer's information. Intuitively, since the eligibility set for a participant shrinks as more information is acquired by the designer, so does the offer set across stages. A mechanism is 'monotonically increasing in offers' if each participant is offered all alternatives in her 'cumulative' core eligibility across previous stages under the designer's information. As the cumulative core eligibility set expands with the designer's information, so does the offer set.

[16] Importantly, this also prevents the designer from introducing dummy participants at any stage, as she has to account for the risk that this causes the real participants to 'drop out' of the mechanism (choose the unassigned option) when they should not, which could lead to waste. A similar intuition is used by Akbarpour and Li (2020).

5 Additional Results

We have shown in Theorem 1 that a predictable mechanism forms 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 allocation, or if it is costly to set up multiple interactions for participants. In this section, we again assume the designer uses the promised environment $\mathcal{E}^*$, and propose a stronger communication protocol that can be used even with a one-stage mechanism. We show that it achieves a stronger notion of verifiability (though falls short of transparency). This 'step-cutoff' protocol communicates cutoffs for each 'step' of the algorithm underlying the mechanism.

It is useful to define an algorithm for a one-stage mechanism, in order to be free of stage-dependent definitions; for certain sequential mechanisms, steps could correspond exactly to stages in the mechanism. Let g be a mechanism. Recall that the designer's initial information is $\phi^{g,0} = (\geq^0, v^*)$, and terminal information is $\phi^{*g} = (\geq^*, v^*)$. In our information-based representation, an ***algorithm*** for g is given by a sequence of information $\mathcal{U}_g = (u_0, \ldots, u_r) \in \Phi^r$ such that (1) $u_0 = \phi^{g,0}$; (2) each u_s is formed from u_{s-1} by adding at most one alternative in the action $\geq_i^*$ of any participant i selected in that step; and (3) $u_r = \phi^{*g}$. We call each $u_s \in \mathcal{U}_g$ a ***step*** in the algorithm for g. Define $a(u_s) \in X^N$ as the temporary allocation in step u_s formed by the vector of these single alternatives for each participant.

In general, we impose no further constraints on the algorithm behind a mechanism.[17] A step-cutoff protocol communicates cutoffs for each step of the underlying algorithm. Formally, given an environment $\mathcal{E} = (g, \mathcal{M})$, a protocol $\mathcal{M}$ is a step-cutoff protocol if $M^g = (c(a(u_s)))_{u_s \in U_g}$. It is clear that step-cutoffs include terminal-cutoffs (at the final step of the mechanism), and thus step-cutoff protocols communicate more information than terminal-cutoff protocols.

We now show that such step-cutoffs achieve a stronger notion of transparency than verifiability. To motivate this notion, recall that several allocations could be justified under the promised mechanism by providing the corresponding terminal-cutoffs. For instance, if the designer promises the second-price auction, but communicates the highest bid as the winning price (i.e., the terminal-cutoffs for the first-price auction), no bidder observing this price can individually rule out the possibility that this was indeed

[17] However, note that, for some mechanisms, there are also restrictions on which participants are selected in a particular step. For instance, take a step $u_s \in U_g$. Since u_s is information, it can be used to compute eligibility sets for each participant. An 'eligibility-based' algorithm is such that a participant is selected in step u_{s+1} only if she is ineligible under g for all alternatives that have been included from her actions up to and including step u_s. For instance, the SPDA algorithm in school admissions selects students in a step who have been 'rejected' from each previous school to which they have applied, and the simulated ascending auction for the second-price auction selects bidders whose previous bids are exceeded by the current price.

the true winning price (see also Subsection 6.2).[18] To remove this possibility, a stronger transparency requirement for $\mathcal{E}^*$ is for the true allocation a^* to be the *only* allocation justified in $\mathcal{E}^*$.

Definition **3.** The promised environment $\mathcal{E}^*$ is strongly verifiable if the true allocation a^* is the *unique* justified allocation in $\mathcal{E}^*$.

A strongly verifiable environment is verifiable, though the converse is not true. Strong verifiability requires more information to be acquired externally by each participant, and is thus in general harder to achieve. In fact, we show that the mechanism dimension cannot be used to achieve strong verifiability at all. As we show in a school admissions example in Subsection 6.2, a student in a closed room cannot establish – through her experience alone – whether the mechanism she is taking part in is g^* or another similar mechanism. On the other hand:

Proposition **3.** *A monotonic-offers environment $\mathcal{E}^* = (g^*, \mathcal{M}^*)$ is strongly verifiable through communication if $\mathcal{M}^*$ is a step-cutoff protocol.*

The proof is in Appendix C. Essentially, step-cutoffs allow participants to infer more detailed information about the actions and scores of other participants, helping them to rule out the possibility of another justified allocation. However, we show in Subsection 6.3 that step-cutoffs are not sufficient for transparency, as they can be manipulated by the designer without detection.

6 Main Application: School Admissions

In this section, we present our main application: the student-proposing Deferred Acceptance (SPDA) procedure in school admissions. This section is also useful as an illustration of the main concepts we apply throughout the paper, namely, information, eligibility, terminal-cutoffs, and predictable mechanisms. We also use this section to demonstrate how the designer can manipulate information through experience and communication without being detected by any participant.

In a standard school admissions model, there is a finite set of students $N = \{1, 2, 3, \ldots\}$, and a finite set of schools $X = \{x, y, z, \ldots\}$, which includes the unassigned option $\emptyset$. For each school $x \in X$, an integer $q_x \in \{1, \ldots, N\}$ represents the capacity of school x, with $q_\emptyset = |N|$. A feasible allocation of students to schools respects capacities: $a \in X^N$ is a feasible allocation if $|\{i \in N \mid a_i = x\}| \leq q_x$ for each $x \in X$. Let $\mathcal{A}$ denote the set of feasible allocations. Each student $i \in N$ has (strict) preferences over schools, given by $\succeq_i$. A preference profile is $\succeq = (\succeq_i)_{i \in N}$. Each school $x \in X$ has a priority v^x, such that, for any $i \in N$,

[18] The second-highest bidder, for instance, does not know that she placed the second-highest bid, and might assume that there were two bidders with higher bids than hers. The winner also cannot detect that this is the outcome of an artificial bid.

$v_i^x \in \mathbb{R}_+$ is the score of student i. A priority profile is $v = (v^x)_{x \in X}$. Fixing the set of students and schools, a school admissions problem $P = (\succeq, v)$ is determined by a preference profile and a priority profile.

Example **1.** In this very simple illustrative example, let $N = \{1, 2, 3\}$ be a set of three students, and let $X = \{x, y, z\}$ be a set of three schools, each of unit capacity. The table on the left of Figure 2 gives an example of a school admissions problem, consisting of a preference profile and a priority profile. We illustrate the role of information with an example. Suppose Student 1 knows her own preferences and scores at each school, and knows nothing of other students' preferences and scores. This information is summarised in the table on the right of Figure 2.

Figure 2: A school admissions example

Complete information

Preferences			Scores			
$\succeq_1$	$\succeq_2$	$\succeq_3$	Student	v^x	v^y	v^z
x	y	y	1	85	70	75
y	z	x	2	55	45	90
z	x	z	3	40	80	60

Student 1's (partial) information

Preferences			Scores			
$\succeq_1$	$\succeq_2$	$\succeq_3$	Student	v^x	v^y	v^z
x			1	85	70	75
y			2			
z			3			

Notes: The table on the left represents complete information, i.e., a school admissions problem. A student's preferences over schools are such that a school in a higher row is preferred to a school in a lower row. Student scores at each school determine their priority at that school. For instance, Student 1 has a score of 85 at School x, and Student 3 has a score of 60 at School z, etc. The table on the right is represents Student 1's information, which in this example consists only of her own preferences and scores. Each way that she can suitably fill in the missing values in the tables generates an allocation problem compatible with her information.

It can be easily seen how this information generates a set of compatible problems, each of which is found by suitably 'filling in the blanks' of what Student 1 does not know about other students' preferences and scores. Any additional information that Student 1 might acquire (on others' scores at School x, say), would help her further refine the set of allocation problems. ■

6.1 School Admissions Mechanisms

The celebrated student-proposing deferred acceptance (SPDA) rule (Gale and Shapley, 1962; Abdulkadiroğlu and Sönmez, 2003) has a number of variants in terms of mechanisms. In a one-stage SPDA mechanism, all students are active and are offered all schools for which they are initially eligible, and permitted actions involve strictly ranking all offered schools (or at least all acceptable schools). The designer's information at the end of this stage is enough to determine a unique allocation via the SPDA algorithm, which works as follows: The set of students N is the 'proposing' side of the market. In the first step, an application is sent from each student $i \in N$ to her most-preferred school in X according to her actions $\geq_i$. Each school $x \in X$ tentatively accepts up to q_x of its highest-priority applicants (according to v^x), rejecting all other applications (if there are fewer than q_x applications, all are tentatively accepted). In each subsequent step, an application is sent from each rejected student in the previous step to her most-preferred school from among those that have not rejected her already, if any such schools remain in her actions. Each school x considers its revised set of applications, including the applications held from the previous steps, if any, and tentatively accepts up to q_x highest-priority of these, rejecting all others. The procedure terminates when no more applications are sent. As with all non-wasteful direct mechanisms, direct SPDA is monotonically decreasing in offers. Recall the school admissions problem in Example 1. If preferences are interpreted as students' actions in the direct mechanism, the SPDA algorithm proceeds in the following steps:

Step 1: Each student's application is sent to her top-ranked school. Thus Student 1 applies to School x while Students 2 and 3 apply to School y. School y thus receives more applications than its capacity.

Step 2: Student 2 has a lower score for School y than Student 3, and is thus rejected. Her application is sent to her next-ranked school, which is School z. There are no more rejections, and the final allocation assigns School x to Student 1, School y to Student 3 and School z to Student 2.

The experiential information conveyed by the one-stage SPDA mechanism is limited. There is only one stage in which any student is active, she is offered all schools, and is required to rank all those she considers acceptable. In particular, there is usually no way from her experience alone that she can determine which of those schools will be her assignment, as it depends also on the actions and scores of others. Notice that the table on the right of Figure 2 in effect summarises Student 1's information in the one-stage SPDA mechanism, since she acquires no useful experiential information from the mechanism. Each of the schools she lists in her actions could be her assignment, and so her eligibility set contains all three schools.

6.2 Predictable Mechanisms

Richer experiential information can be derived from multi-stage variants of SPDA. We highlight one SSM. In any stage, the students who are made active are precisely those who are ineligible for all schools contained in their previous actions, and this is conveyed to them by the rules of the mechanism. Each active student is offered all schools for which she is still eligible according to the designer's information, which makes the mechanism monotonically decreasing in offers. A permitted action involves selecting exactly one of these schools, and this school is appended to the bottom of the ranking that she has already indicated for schools contained in her previous actions. The mechanism terminates when each student is eligible for her last-chosen school, which is her assignment.

Notice that the experiential information for each student allows her to predict her assignment – either she will be assigned her last-chosen school (if the mechanism terminates at that stage), or she will be asked to pick again. However, the table on the right of Figure 3 illustrates how predictability is not enough for verifiability, as the designer can reach another allocation without violating any student's experience of the mechanism.

Figure 3: Predictable school admissions mechanisms

True SSM

Stage	Student		
	1	2	3
1:	x	~~y~~	y
2:		z	

'False' SSM

Stage	Student		
	1	2	3
1:	x	y	~~y~~
2:			~~x~~
3:			z

Notes: Predictable mechanisms require students to list one school at a time. The true SSM for SPDA correctly rejects Student 2 in the first stage, who then applies to School z in the second stage. In each stage, each student has exactly one school in her eligibility set. The table on the right illustrates how a 'false' SSM could be used, which is nevertheless experientially indistinguishable from the true SSM. In this mechanism, Student 3 is falsely rejected in the first stage, and rejected also in the second stage. The allocation reached is different from the SPDA allocation.

6.3 Verifiable Communications

We illustrate terminal-cutoffs and step-cutoffs with an example, and also show how they can be manipulated by the designer.

Example **2.** Let $N = \{1, 2, 3\}$ be three students, and let $X = \{x, y, z\}$ be three schools, each of unit capacity. Consider the designer's information generated by a direct SPDA mechanism in Figure 4. The true SPDA allocation is given in boxes, and a false allocation for this problem is underlined. The centre table of Figure 4 gives the terminal-cutoffs corresponding to each of these allocations. Notice that each student, if she is given the false assignment, has a single-valued eligibility set under the false terminal-cutoffs, which corresponds to her assignment. Thus multiple allocations could be justified, depending on which terminal-cutoffs are communicated. In this simple example, there should be only one step in the algorithm, since each student is assigned to her top-ranked school. Thus the true step-cutoffs would contain just one vector of cutoffs, which correspond to the terminal-cutoffs for the true allocation.

If students are told they will receive step-cutoffs, this information is enough for them to check that their assignments are uniquely part of the true allocation. However, the table on the right of Figure 4 shows how step-cutoffs can be manipulated by the designer to reach the false allocation. In particular,

Figure 4: Cutoffs in school admissions

The designer's information

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

True and false terminal-cutoffs

Allocation	School		
	x	y	z
(x, y, z)	40	45	50
(y, z, x)	70	65	55

False step-cutoffs

Step	School		
	x	y	z
1:	40	60	50
2:	70	60	55
3:	70	65	55

the cutoffs for School y and School z are falsified in the first step, and raised high enough to justify the rejections of both Students 2 and 3 in this step. In the second step, an application is sent from each of these students to her next school, and the cutoffs are updated. Notice that Student 1 is now ineligible for School x and is thus rejected in this step. In the third step, Student 1 applies to School y, which completes the false allocation. No student can individually detect that the designer has manipulated step-cutoffs in this way.[19] ■

However, by Theorem 1, we get the following corollary:

Corollary **1.** *A school admissions environment $\mathcal{E}^* = (g^*, \mathcal{M}^*)$ where g^* is an SSM and $\mathcal{M}^*$ is a terminal-cutoff protocol is transparent.*

7 Other Applications

7.1 Allocation of Public Housing

There is a finite set of families $N = \{i, j, k, \ldots,\}$ and a finite set of houses $X = \{1, 2, 3, \ldots\}$. Each house has unit capacity. Each family $i \in N$ has strict preferences over houses, given by $\succeq_i$. A feasible allocation assigns each family to at most one house, and vice versa: $a \in \; X^N$ is feasible if $a_i \neq a_j$ for all $i, j \in N$. The set of feasible allocations is $\mathcal{A}$. There is a common queue of families, given by a bijection $v^* : N \to \{1, \ldots, N\}$, that represents the priority for each house $x \in X$ (i.e. $v =^{*x}= v^*$ for each $x \in X$). A house allocation problem $P = (\succeq, v^*)$ consists of a profile of preferences and a queue.

[19]This manipulation of step-cutoffs is possible only if students do not know how many students there are in the application process. In particular, if Student 2 knew that there are only three students, then the manipulation of step-cutoffs in the first step could immediately be detected. But of course this is a very simple example, and such a situation could easily be embedded into a larger example. For instance, imagine a fourth student in this example, with a score of 11 in all schools. This student is unassigned under both true and false assignments. Then, even if all students know the number of students participating, manipulation of step-cutoffs is undetectable.

The direct serial dictatorship mechanism SD^D works as follows: There is only one stage, all families are selected to be active, and are offered all houses, and permitted actions involve ranking all houses strictly in their actions at this stage. This mechanism is trivially monotonically decreasing in offers. At the end of this stage the designer's information is complete, and allows her to determine the unique feasible allocation and realised assignments, by assigning each family in order of the queue the highest-ranked house still available after earlier families in the queue have received their assignments. The non-transparency of an environment with this mechanism and no communication is immediate.

Verifiability through experience can be induced by the use of a suitable SSM. There are many SSMs to choose from. One particular SSM is denoted SD^A, and works as follows: A stage corresponds to selecting *exactly one* active family in order of their position in the queue. This family is offered all houses that have not already been selected by previous families, and the rules of the mechanism convey to the family that any house *not* offered has already been selected. This family selects precisely one of the offered houses, and this house is her assignment. It is easy to see that this mechanism is predictable, and thus the environment is verifiable through experience.[20] This mechanism is also monotonically decreasing in offers. An example illustrates cutoffs:

Example **3.** There are four families $N = \{1, 2, 3, 4\}$ and five houses $X = \{x, y, z, v, w\}$. Consider the direct mechanism SD^D. The actions for families are given in the table on the left of Figure 5. The queue over families is given in the table in the centre-left. The realised allocation is underlined.

Figure 5: A problem and the resulting allocation under SD

Actions			
s_1	s_2	s_3	s_4
x	x	z	y
y	y	v	z
z	w	w	x
v	v	x	w
w	z	y	v

Queue
1
2
3
4

Terminal-cutoffs	
House	Value
x	1
y	2
z	3
v	0
w	4

Step-Cutoffs					
Step	Value				
	x	y	z	v	w
1	1	0	0	0	0
2	1	2	0	0	0
3	1	2	3	0	0
4	1	2	3	0	4

The table on the centre-right of Figure 5 shows the minimum position in the queue required to be eligible for each house, which corresponds to the terminal-cutoff for each house.[21] A terminal-cutoff protocol communicates the equivalent of the table on the centre-right of Figure 5, and induces verifiability through communication. ■

[20] This mechanism is also strongly obviously strategy-proof (SOSP), in the sense of Pycia and Troyan (2019), showing that predictability and SOSP may be satisfied together in some cases.

[21] Note that using the rank in the queue inverts the cutoff requirement, since a *lower* cutoff value according to rank in the queue corresponds to a *higher* eligibility. The only exception is a cutoff of 0, which signifies that the house is unassigned.

Moreover:

Corollary **2.** *The house allocation environment $\mathcal{E} = (SD^A, \mathcal{M})$ where $\mathcal{M}$ is a terminal-cutoff protocol is transparent.*

7.2 Single-Object Auctions

There is a single seller, who owns a single unit of an indivisible object, and wishes to sell it to one among a finite set of buyers N, in return for a payment. For simplicity, we assume that the seller has a reservation price of zero; she is willing to sell the object at any positive price. The set of alternatives $X \subsetneq \mathbb{R}_+$ is a finite subset of real numbers such that $0 \in X$. The interpretation of an alternative $x \in X$ is the prospect of winning the object and paying the price x for it.[22] An allocation specifies which buyer wins the object, and the price she pays. Each buyer's assignment is whether she wins the object or not, and the price she pays. Feasible allocations identify a unique winner, and a payment in X for each buyer.

For simplicity, we assume that each buyer's preferences can be represented as a valuation for the object, or the 'maximum willingness to pay' in order to obtain the object. Thus buyer i's preferences are an element of X. A preference profile is a vector in X^N. For each buyer, a ***bid*** is an element of X. Priorities in this framework are bids, the interpretation being that a buyer with a higher bid for the object has a higher priority to win it. To sell the object, the seller conducts an auction. We consider the two most popular auctions: the second-price and the first-price auctions.

7.2.1 Second-Price Auction

A second-price auction (SPA), also known as the Vickrey auction (Vickrey, 1961), awards the object to the highest bidder, who pays a price corresponding to the highest losing bid (the second-highest bid), and all losing buyers pay nothing. A direct 'sealed-bid' SPA mechanism SPA^D involves one stage in which each buyer is selected to be active, is offered all prices in X, from which she selects a single price in her action. Assume there are never any ties in the maximum bid.[23] This mechanism is trivially monotonically decreasing in offers. The designer has enough information to determine the resulting allocation uniquely. It is easy to see that, absent any communication from the designer, the environment with the direct SPA^D mechanism is neither verifiable through experience nor through communication.

We identify a specific SSM for the Vickrey auction. This corresponds to an ascending clock auction, denoted SPA^A, which works as follows. The designer starts at the price of $p^0 = 0$, and in each stage k of

[22] The finiteness of X is to fit with our modelling assumptions of a finite set of alternatives, though it can be arbitrarily large.

[23] This is a restrictive assumption in general, and we make it only for simplicity of exposition. Otherwise, we would have to specify a tie-breaking rule to determine the winner in the case of ties in maximum bids.

the mechanism, she raises the price by the smallest increment in X. In stage k, the corresponding price p^k is offered to all remaining buyers in the auction, along with the option to drop out. Each buyer's action in that stage consists of selecting either this price (which is her bid at that stage), or to drop out.[24] Any buyer who drops out in a stage is never active in a future stage. If more than one buyer takes an action p^k, the designer continues to the next stage. Otherwise, the mechanism terminates with the unique allocation in which the sole remaining buyer in that stage wins the auction, and pays a price p^{k-1} corresponding to the price in the previous stage. It is straightforward to see that this mechanism is predictable, as each buyer's eligibility set under her experiential information, at each stage, consists precisely of one alternative – to be unassigned if she chooses to drop out, to win and pay the price p^{k-1} if this is the terminal stage of the mechanism, or to be asked to bid again. Thus an environment with the SPA^A mechanism is verifiable through experience (Proposition 1). This formalises the intuition that ascending auctions are perceived to be 'more transparent' to participants than sealed-bid equivalents (Cramton, 1998; Ausubel, 2004; Perry and Reny, 2005; de Vries, Schummer, and Vohra, 2007).

A terminal-cutoff protocol corresponds to communicating the winning price. It is easy to see that such a protocol induces verifiability through communication, since the true winning price allows each buyer to conclude whether she has won the auction or not. It is also easy to see why this protocol does not induce strong verifiability. If the designer were to communicate any other price between the highest bidder's bid and the second-highest bid, this too would lead to a justified allocation, as no buyer could individually detect that this was not indeed the second-highest bid.[25] However:

Corollary **3.** *A single-object auction environment $\mathcal{E} = (SPA^A, \mathcal{M})$ is transparent if $\mathcal{M}$ communicates the winning price.*

The intuition in this application is that, by committing to use an ascending auction, the designer is limited in terms of what she knows of buyer's valuations. In particular, she cannot artificially raise the bid beyond the price at which the second-highest bidder drops out, because of the risk that the highest bidder might also drop out. Predictability of the mechanism ensures then that no buyer wins the auction,

[24]To make this mechanism monotonically decreasing in offers, we could require all prices greater than or equal to p^k to be offered to each remaining buyer, and restrict permitted actions to selecting either the lowest offered price or to drop out.

[25]Step-cutoffs in this case correspond to a vector of prices, such that each element of the vector is a price at which some bidder has dropped out of the auction. If step-cutoffs are communicated truthfully, only the true allocation can be justified, as the price at which the penultimate bidder dropped out of the auction is communicated truthfully, and is the true winning price. It is also possible to see how step-cutoffs can be manipulated – the designer can in some instances artificially raise the price at which the penultimate bidder was supposed to have dropped out. As long as there is at least one price at which two or more bidders dropped out, this is undetectable by bidders (if each bidder dropped out at a unique price, then the step-cutoffs must include each of these prices in order to be plausible, but then an artificially elevated dropout price involves introducing an extra dropout price, which can be detected.)

which violates non-wastefulness. The winning price is thus constrained to be the price at which the second-highest bidder drops out, and so the sole justified allocation is the true allocation.

7.2.2 First-Price Auction

A first-price auction (FPA) awards the object to the highest bidder, who pays her bid as the winning price, and all losing buyers pay nothing. As before, a direct mechanism FPA^D works like SPA^D, except of course for a different winning price. If there is no communication, this environment is unverifiable through experience or communication.

A particular SSM for the first-price auction is denoted FPA^A, and is the following variant of a descending auction: In the first stage, the designer starts with the highest price in X. Each buyer is selected to be active, and is offered this price, along with an option to 'continue'. Each buyer's action consists of selecting one of these options. If each buyer chooses to continue, the designer goes to the next stage, in which the offered price is reduced by the smallest decrement in X. If instead a buyer's action is a bid at that price, the auction terminates, and the unique allocation consists of awarding the object to this buyer at this price.[26] It is easy to see that this mechanism is predictable, and so an environment with this mechanism is verifiable through experience (Proposition 1).

As in the case of the second-price auction, terminal-cutoffs correspond to the winning price, and achieve verifiability through communication. Interestingly, step-cutoffs communicate exactly the same information, thus the winning price also makes the environment strongly verifiable. However, transparency requires each buyer to be able to conclude that the designer indeed used the first-price auction.

Corollary **4.** *The single-object auction environment* $\mathcal{E} = (FPA^A, \mathcal{M})$ *where* $\mathcal{M}$ *communicates the winning price is transparent.*

The intuition behind this result is that, by committing to use a descending auction, the designer is constrained in what she knows about bids. The risk for the designer from lowering the price 'too much' in any stage is that more than one buyer may place a bid at that price. By predictability, they should each win the auction at that price, which is not feasible. Thus, to determine the maximum price at which a unique buyer places a bid, the designer has to lower the price in the smallest possible decrements, and by our assumption on unique maximal bids, there is a unique first stage in which the auction terminates, which uniquely determines the winning price that corresponds to the true allocation.[27]

[26] Under our assumptions of unique maximal bids, there can never be more than one buyer who first chooses a price. Moreover, this mechanism can be made monotonically increasing in offers by requiring all previous prices to be offered in each stage.

[27] It can be seen that the unique maximal bid assumption is not crucial, but saves having to specify a tie-breaking rule when there is more than one maximal bid.

8 A School Admissions Experiment

In this section, we present a laboratory experiment based on the SPDA rule in school admissions. Recall that our theory assumes rational participants who can fully and accurately process communicated and experiential information through the promised mechanism into determining their eligibility set. As our primary motivation is to address the lack of transparency in practice, it is essential to check whether the participants can indeed do this processing. Our experiment is designed as a proof of concept of the usefulness of this information for transparency, and not only to test our theoretical results.

In our experimental design, we opt for honesty of the designer. Thus, the promised mechanism is always run, but a random draw might determine the assignment. In our setup, verifiability is enough to spot random assignments. We test separately for the effect of experiential information and communicated information on 'observed transparency.' Experiential information is generated by the use of a predictable mechanism, which theoretically ensures verifiability. The complexity of decision-making arises from the need to understand the rules of the promised mechanism and infer eligibility sets 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 eligibility. We also use a theoretically transparent environment, which provides both sources of information. Our experimental setting is a novel empirical tool to compare the relative transparency of any environment. Thus, it can be used to test the observed transparency of newly proposed environments before putting them into practice. In this sense, it is of independent interest.

8.1 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 'sequential' mechanism, where participants apply to schools one by one (Echenique, Wilson, and Yariv, 2016; Bó and Hakimov, 2018; Klijn, Pais, and Vorsatz, 2019). The other dimension is on the communication protocol. The first protocol conveys no information to participants. The other is a step-cutoff protocol where we provide the cutoff grades of each school at each step of the underlying DA algorithm.[28] The four environments described below correspond to our four treatments in

[28] 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 allocation). This is further discussed when we provide the exact description of step-cutoffs in the case of random

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 it holds, and students who are not held at any school are left unmatched. This is the unique realised allocation. 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 allocation 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 verifiable through communicated information (Proposition 3). Thus, this treatment tests the ability of participants to use communicated step-cutoffs to refine their eligibility sets under SPDA.

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

Each student applies to schools sequentially, one at a time. After each step, a student learns through the rules of the mechanism whether she was temporarily accepted or rejected 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

allocations.

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.[29]

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

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

Every student participates in SeqNo, but the cutoff grades of each school are revealed at each step. Moreover, after the procedure terminates, students also receive the table of step-cutoff grades for each school. This environment is verifiable through information received from both experience and communication, and is also transparent.[30]

8.2 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. In all treatments, subjects received 31 Swiss Francs (CHF) if they were assigned 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 assignment would be determined not by the explained SPDA procedure but instead randomly. This probability was common across all treatments. After the allocation was determined, subjects were asked to take a decision on whether or not to appeal. Payoffs were constructed to make appealing optimal if the assignment was determined randomly. If a subject

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

[30] 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.

appealed and the appeal was incorrect (the allocation was determined by the explained procedure), she incurred a cost of 6 CHF, which was deducted from her allocation payoff. If a subject appealed and the appeal was correct (the allocation was determined randomly), her total payoff for the round was 40 CHF, replacing whatever allocation payoff she might have otherwise received. If a subject decided not to appeal, she received her allocation 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 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 assigned 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.[31] 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).[32] 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 only of their own grades. We believe that this too approximates real-life informational conditions. After each round, subjects received feedback only about their own assigned school. Additionally, in the DirCutoffs and SeqCutoffs treatments, they observed the table of step-cutoffs grades of all schools. After submitting an appeals decision, they learned whether their decision was correct, as well as their final payoff for the round.

The random allocation for each round was pre-generated and fixed for all subjects. It never coincided with the student-optimal stable allocation. 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 allocation 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 allocation was determined randomly in a round in one session (a 'random round'), then the allocation was determined through the explained procedure

[31] The details of all markets are presented in Appendix F.

[32] Note that payoffs were designed to make appealing optimal if the assignment 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 appeals decisions.

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

In random rounds, if subjects observe a final step-cutoff for their assigned school different from their grade at that school, the appeal decision is trivial. To address this concern, we modified the cutoffs in the final step for random rounds. The final cutoff of the random assignment showed the student's grade (so as to match the allocation), and the cutoff was randomly generated at the school corresponding to the correct allocation.[33] Participants were informed about this modification in 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 allocation 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 gets the assured payoff of 31 CHF. But 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 assigned to the second most-preferred school, down to Row 8, which corresponds to assignment 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

[33] 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.

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.[34] Since subjects played against sims 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.[35] On average, the experiment lasted 70 minutes, and the average earnings per subject were 38 CHF, including a show-up fee of 10 CHF.

At the beginning of the experiment, subjects were given printed instructions (see Appendix F). 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.

8.3 Experimental Results

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

8.3.1 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[36] measures the degree of transparency of the environment, as the more transparent the environment, the easier it is 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 6:

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

[34] We invited 26 subjects to every session in order to have 24 subjects per session, but some sessions had a high rate of no-shows.

[35] 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 5% significance level.

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

of correct appeal decisions under DirCutoffs (65%) and SeqNo (64%).

Figure 6: Proportion of correct appeal decisions

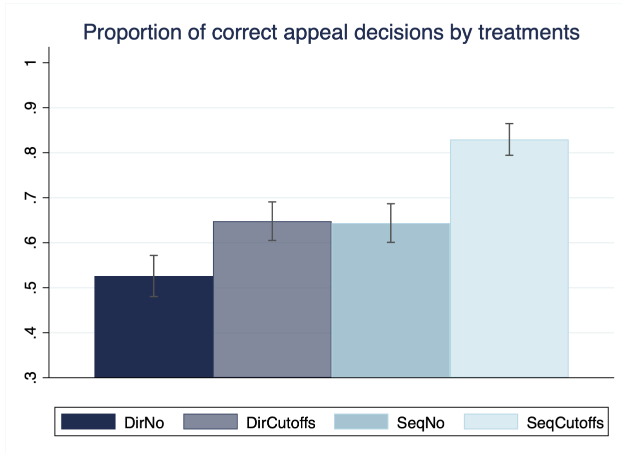

The proportion of 52% correct appeal decisions in DirNo corresponds effectively to a random chance of a correct decision. This evidence supports our hypothesis regarding the lack of transparency of the direct DA mechanism with no feedback. Once step-cutoff grades are provided to subjects, the proportion increases significantly to 65%. This difference suggests that subjects are at least partially able to use step-cutoffs to judge their assignment'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 assignments' legitimacy.

Our theoretical predictions find support in terms of treatment differences, but not the levels. Note that verifiability in our design should lead to 100% correct appeal decisions.[37] 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 participant's grade always corresponds to the announced allocation. Similar mistake rates are also observed in SeqNo, suggesting that many subjects did not understand the relation between their assignments 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 the last five rounds for each treatment. There is a significant increase in the

[37]In fact, there were 29 of 1,860 outcomes in the experiment when the random outcome was the same as the allocation outcome. This might make the identification of manipulation impossible, though not necessarily. 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, while excluding these observations would not qualitatively change the results.

proportion between the first five and the last five rounds of both the DirCutoff and SeqNo 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 (1)	DirCutoffs (2)	SeqNo (3)	SeqCutoffs (4)	1=2 p-val.	1=3 p-val.	1=4 p-val.	2=3 p-val.	2=4 p-val.	3=4 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 the 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 allocation from both sources, but signals are followed only when both are present. Note also that there is no significant difference across rounds in SeqCutoffs, with the rate remaining stable at around 83 per cent in the first five and last five rounds.

8.3.2 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 assignment payoff, the more likely a subject is to appeal.*
3. *Subjects are more likely to make correct appeal decisions when they play truthfully.*

Figure 7 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.

Figure 7: Proportion of correct appeal decisions by optimal decision

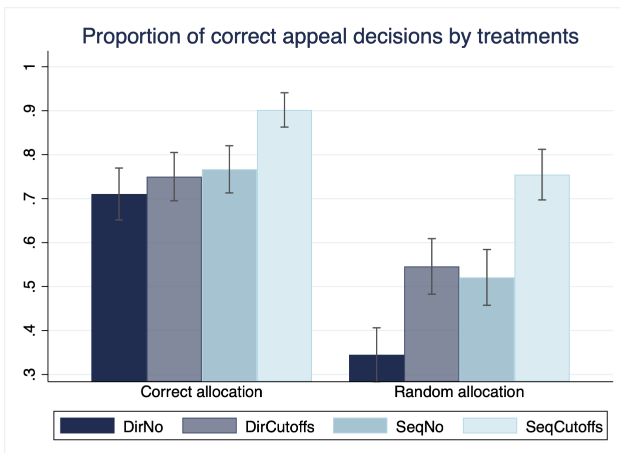

What explains the aversion to appeals? It is useful to have a benchmark. If there is no further information to suggest whether the allocation is random or not, appeal decisions should be driven purely by subjects' risk preferences based on 50% probability of random allocation. The tradeoff that subjects face in appeal decisions depends on the rank of the school that they received. Thus, for each decision, we can construct a decision of appeal based purely on a 50% probability of a successful appeal. For instance, if a subject is assigned 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 allocation was generated randomly with a 50% probability. Similar choices for lower ranked assignments 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, DirCutoffs, 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 allocation being generated by the mechanism, and not randomly. This might be driven by confirmation bias when subjects try to rationalise the allocation through the prism of the mechanism.

- Appeal-aversion is explained by the loss aversion of subjects with the allocation 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).[38] 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 assigned school, the less likely participants are to appeal correctly (the coefficient for the rank of the assigned school 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 assignment, the more likely the appeal.

Third, on average, subjects who play 'truthfully'[39] 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 assignments.

We present the results regarding the optimality of mechanism strategies, stability, and efficiency in Appendix E. The main takeaway is that transparency is not necessarily aligned with reporting preferences

[38] 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.

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

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 assignment		-.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. Models (5) and (6) restrict the sample to random rounds. Dummy for random is 1 in random rounds and 0 otherwise. Rank of assignment is the rank of the resulting allocation in the true preferences of the participant. Dummy for truthful is 1 if all participants played the truthful strategy (see Appendix E for an explanation). 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 assignments. * $p < 0.10$, **$p < 0.05$, *** $p < 0.01$. Standard errors are clustered on the individual level and are presented in parentheses.

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

9 Conclusion

This paper identifies verifiability as a desirable attribute of centralised allocation environments, in terms of being able to convince participants of the legitimacy of their assignments. Moreover, it formulates a stronger notion of transparency as precluding the designer from changing the allocation undetected. We show how achieving verifiability and transparency depends not only on the appropriate communication of information from the designer to participants but also on the new idea of the experience induced by a mechanism. We believe this provides a fundamentally new direction for research in mechanism design. Our experimental results support the idea that the two sources of information through experience and communication improve observed transparency, and especially so when provided together. The experimental setup can be of independent interest as a tool to compare the relative transparency of alternative proposals for mechanisms.

We have identified sufficient conditions for transparency. The simple transparent implementations of popular mechanisms that we identify provide blueprints for policymakers who seek a way to improve transparency in those applications. However, an open theoretical question is whether transparency is possible to achieve if participants acquire even less information. In particular, what would be the *minimum* information required for transparency? A related open question is to formulate a way to compare applications in terms of the informational requirements for transparency.

Many important applications remain open. The TTC mechanism, for instance, is an immediate objective, given its theoretical relevance and relative non-transparency. Our results show that transparency is possible for TTC. While predictability appears straightforward (perhaps by using a 'pick-an-object' mechanism (Bó and Hakimov, 2020b)), verifiability through communication might be more challenging. A starting point is Dur and Morrill (2018) and Leshno and Lo (2020), who identify the (multidimensional) structure of cutoffs for TTC. There are other applications where cutoffs can be complicated in structure – for instance multi-unit auctions with interdependent values. Determining these cutoffs for such applications is an open question, as is the empirical test of whether they can be used by participants.

Another prominent mechanism that might benefit from this analysis is the efficiency-adjusted deferred acceptance mechanism (EADAM) (Kesten, 2010). Based on obtaining consent from students to waive potential priority violations, EADAM generates welfare improvements to SPDA. However, the final allocation produced by EADAM might contain instances where some students are assigned schools for which their grades do not meet the original SPDA cutoff. Our framework does not suggest a transparency solution for EADAM. For instance, a sequential implementation might lead to some students being assigned to schools that had rejected them previously, and thus might not be predictable (or monotonic-offers).

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.

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 L. Chen (2019): "Designing Heaven's Will: Lessons in Market Design from the Chinese Imperial Civil Servants Match," Working paper.

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.

Boczoń, M. and A. J. Wilson (2020): "Goals, constraints and transparent assignment: A field study of the UEFA Champions League," Working paper.

Boehm, F. and J. Olaya (2006): "Corruption in public contracting auctions: the role of transparency in bidding process," *Annals of Public and Cooperative Economics*, 77, 431–452.

Börgers, T. and J. Li (2019): "Strategically simple mechanisms," *Econometrica*, 87, 2003–2035.

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.

de Vries, S., J. Schummer, and R. V. Vohra (2007): "On ascending Vickrey auctions for heterogeneous objects," *Journal of Economic Theory*, 132, 95–118.

Dur, U., R. G. Hammond, and T. Morrill (2018): "Identifying the harm of manipulable school-choice mechanisms," *American Economic Journal: Economic Policy*, 10, 187–213.

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.

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 (2019): "Decentralizing Centralized Matching Markets: Implications from Early Offers in University Admissions," Working paper.

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. and D. Kübler (2020): "Experiments on centralized school choice and college admissions: a survey." *Experimental Economics*.

Hassidim, A., D. Marciano, A. Romm, and R. I. Shorrer (2017): "The mechanism is truthful, why aren't you?" *American Economic Review*, 107, 220–24.

Hassidim, A., A. Romm, and R. I. Shorrer (2018): "'Strategic' behavior in a strategy-proof environment," Available at SSRN 2784659.

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

Hurwicz, L. (1973): "The Design of Mechanisms for Resource Allocation," *American Economic Review*, 63, 1–30, papers and Proceedings of the Eighty-fifth Annual Meeting of the American Economic Association (May 1973).

Kagel, J. H., R. M. Harstad, and D. Levin (1987): "Information impact and allocation rules in auctions with affiliated private values: A laboratory study," *Econometrica*, 1275–1304.

Kagel, J. H. and D. Levin (2001): "Behavior in multi-unit demand auctions: experiments with uniform price and dynamic Vickrey auctions," *Econometrica*, 69, 413–454.

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," Accepted, *Review of Economic Studies*.

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.

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.

Perry, M. and P. J. Reny (2005): "An Efficient Multi-Unit Ascending Auction," *Review of Economic Studies*, 72, 567–592.

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.

Rees-Jones, A. (2018): "Suboptimal behavior in strategy-proof mechanisms: Evidence from the residency match," *Games and Economic Behavior*, 108, 317–330.

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

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

Shorrer, R. I. and S. Sóvágó (2017): "Obvious mistakes in a strategically simple college admissions environment," Tinbergen Institute Discussion Paper 2017-107/V.

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

Troyan, P. (2019): "Obviously strategy-proof implementation of Top Trading Cycles," *International Economic Review*, 60, 1249–1261.

Vickrey, W. (1961): "Counterspeculation, auctions, and competitive sealed tenders," *The Journal of Finance*, 16, 8–37.

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 Formal Definitions

A.1 Mechanisms: General Definition

The following class of information will be useful to define when a mechanism 'terminates'. Let $\phi \in \Phi$ be such that $g(P) = g(P')$ for all $P, P' \in \phi$, i.e., ϕ is such that the mechanism g produces the same feasible allocation for all compatible problems. We define a class of information called the unique allocation frontier for g, denoted $\Phi^1_g = \{\phi \in \Phi \mid g(P) = g(P') \text{ for all } P, P' \in \phi\}$, which is the collection of such information with unique feasible allocations produced by g. The unique allocation frontier captures the intuition that a mechanism g might not require full information (i.e., how each participants ranks all alternatives) in order to uniquely determine a feasible allocation; it only requires information in Φ^1_g.

Recall that the designer knows the priority profile v^*. In this paper, we abstract from strategic issues related to the reporting of preferences, and instead model reported preferences as 'actions'. For participant $i \in N$, an ***action*** is denoted by a reflexive and transitive (not necessarily complete) binary relation $\geq_i$ over X. Each mechanism has a set of permitted actions consistent with the nature of preferences in the application. For instance, if admissible profiles in the application consist of strict preferences over alternatives, then permitted actions involve ranking (subsets of) alternatives strictly.[40] Let $Y(\geq_i) \subseteq X$ be the set of alternatives ranked in an action $\geq_i$. We say that an action $\geq'_i$ 'extends' action $\geq_i$ if it ranks more alternatives, preserves strict rankings of already ranked alternatives, and strictly ranks at least one new pair of alternatives. Formally, $\geq'_i$ extends $\geq_i$ if (1) $Y(\geq'_i) \supsetneq Y(\geq_i)$; (2) for all $x, y \in Y(\geq_i)$, $x >_i y \implies x >'_i y$, and (3) there exist $x, y \in Y(\geq'_i)$ such that $x >'_i y$ and either $\{x, y\} \setminus Y(\geq_i) \neq \emptyset$ or $x \geq_i y$ and $y \geq_i x$. An action that extends another contains more information on the participant's preferences. An ***action profile*** is denoted $\geq = (\geq_i)_{i \in N}$.

A 'stage' in our elicitation procedure is indexed by k.[41] Let g be a mechanism. Let $\geq^{g,0}$ be an initial

[40] In general, a mechanism might require participants to strictly rank all alternatives, or to bifurcate the set of alternatives into an 'acceptable' and 'unacceptable' subset, or to strictly rank pairs of alternatives. Survey mechanisms, asking participants yes/no questions, are not covered by this definition.

[41] A short note on notation: in our description of 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. For example, we use $\phi^{g,k}$ to refer to the designer's information obtained precisely in Stage k, and $\phi^{g,\{k\}} = \bigcap_{k' \leq k} \phi^{g,k'}$ to refer to the cumulative

'empty' action profile. Let $\phi^{g,0} \equiv (\geq^{g,0}, v^*)$ be initial information for the designer. If $\phi^{g,0} \in \Phi^1_g$, the designer already has enough information to uniquely determine an allocation under g. Otherwise:

Stage k, $k \geq 1$:

1. **Action:** The mechanism selects active participants $Act^{g,k} \subseteq N$ with $Act^{g,k} \neq \emptyset$. Each active participant $i \in Act^{g,k}$ is offered a set of alternatives $Y_i^{g,k}$ which contains at least one alternative she has not ranked before, i.e., $Y_i^{g,k} \setminus Y(\geq_i^{g,\{k\}}) \neq \emptyset$. Each active participant selects some of these offered alternatives and takes a permitted action that extends previous actions. Formally, active participant $i \in Act^{g,k}$ takes a permitted action $\geq_i^{g,k}$ such that: (1) $Y(\geq_i^{g,k}) \subseteq Y_i^{g,k}$; and (2) $\geq_i^{g,\{k\}}$ extends $\geq_i^{g,\{k-1\}}$. In particular, she ranks at least one new pair of alternatives strictly. Set $\geq_i^{g,k} = \emptyset$ for all inactive participants, and let $\geq^{g,k}$ be the reported preference profile in this stage.

2. **Evaluation:** The designer's information is updated to $\phi^{g,\{k\}} = (\geq^{g,\{k\}}, v^*)$ that includes all actions taken in this stage. If $\phi^{g,\{k\}} \notin \Phi^1_g$, the mechanism goes to Stage $k+1$. Otherwise $\phi^{g,\{k\}} \in \Phi^1_g$, and the unique ***realised allocation*** is $a^* = g(\phi^{g,\{k\}})$. The designer communicates the ***realised assignment*** a_i^* to each $i \in N$, and the mechanism terminates.

The procedure is well defined and terminates in a finite number of stages, since the designer acquires strictly more information in each stage and (since participants and alternatives are finite) eventually acquires information $\phi_g^* \in \Phi^1_g$ at the unique allocation frontier of g. This definition allows for 'direct' one-stage mechanisms as well as multi-stage ones. Across stages, mechanisms can differ in three additional dimensions: (1) which participants are selected to be active in a stage; (2) which alternatives they are offered; and (3) which actions are permitted.

Let K be the terminal stage of g. Set $\geq_i^{*g} \equiv \geq_i^{g,\{K\}}$ for each $i \in N$ as her cumulative actions in the mechanism, set $\phi_i^{*g} \equiv (\geq_i^{*g}, v_i^*)$ as her intrinsic information (consisting of her actions and scores) and $\phi_g^* \equiv (\geq^{*g}, v^*) = \phi^{g,\{K\}}$ as the designer's terminal information.

A.2 SSMs and Predictable Mechanisms

Formally, SSMs are defined by the following three properties:

SS1: *Each* action is single-valued. For any $i \in N$ and any stage k such that $i \in Act^{g^*,k}$, $|Y(\geq_i^{g^*,k})| = 1$.

SS2: Participants are eligible under their experiential information for *only one* action at a time. For each $i \in N$, any two stages k, k' with $k < k'$, and alternatives $x, y \in X$ such that $x \in Y(\geq_i^{g^*,k})$ and $y \in Y(\geq_i^{g^*,k'})$, $x, y \in E_i^{g^*}(\epsilon_i^{g^*,\{k'\}}) \implies x = y$.

information of the designer up to and including Stage k.

SS3: Each participant's assignment is the (unique) alternative in her eligibility set under their experiential information at the terminal stage of the mechanism. Let K be the terminal stage of the mechanism g^*. Then, for each $i \in N$, $a^*_i \in E^{g^*}_i(\epsilon^{g^*,\{K\}})$.

SSMs are a subclass of 'predictable' mechanisms, a concept we now formalise. Let $\mathcal{E} = (g^*, \mathcal{M}^*)$ be the promised environment. We say that the mechanism g^* is predictable if each participant is eligible under her experiential information at each stage for only one alternative contained in her actions up until that stage, and moreover this alternative at the terminal stage of the mechanism is her assignment in a^* reached by g^*. Let k be a stage in g^*, and let $i \in Act^{g^*,k}$ be an active participant in this stage. Recall that her experiential information up to and including this stage is $\epsilon_i^{g^*,\{k\}}$. Moreover, the alternatives contained in her actions up to and including this stage are $Y(\geq_i^{g^*,\{k\}})$. Thus the set of alternatives in her actions for which she is eligible under her experiential information is the intersection of these two sets, i.e., $E^{g^*}_i(\epsilon_i^{g^*,\{k\}}) \cap Y(\geq_i^{g^*,\{k\}})$.

Formally, a mechanism g^* is predictable if:

1. For any k in g^* and any $i \in Act^{g^*,k}$, the intersection of these sets is single-valued, i.e., $|E^{g^*}_i(\epsilon_i^{g^*,\{k\}}) \cap Y(\geq_i^{g^*,\{k\}})| = 1$; and

2. if K is the terminal stage of g^*, then the unique alternative in the intersection of these sets at that stage is her realised assignment, i.e., $E^{g^*}_i(\epsilon_i^{g^*,\{K\}}) \cap Y(\geq_i^{g^*,\{K\}}) = a^*_i$.

The environment $\mathcal{E} = (g^*, \mathcal{M}^*)$ is predictable if g^* is predictable. It should be easy to see that SSMs are predictable in this sense.

A.3 Monotonic-Offer Mechanisms

A mechanism g^* is monotonic-offers if it is monotonically increasing or decreasing in offers, where:

1. g^* is monotonically decreasing in offers if, in any stage, an active participant is offered all alternatives for which she is eligible under the designer's information at that stage. Formally, for any stage k in g^* and any active participant $i \in Act^{g^*,k}$, $Y_i^{g^*,k} = \{x \in X \mid x \in E^{g^*}_i(\phi^{g^*,\{k-1\}})\}$. Since the eligibility set is itself monotonically decreasing in size, offers are thus also monotonically decreasing in size.

2. g is monotonically increasing in offers if each active participant is offered all alternatives in her core eligibility under the designer's information at that stage and all previous stages. Formally, for any stage k in g^* and any active participant $i \in Act^{g^*,k}$: $x \in C^{g^*}_i(\phi^{g^*,l})$ for some $l \leq k$ implies $x \in Y_i^{g^*,k}$. Offers are thus monotonically increasing in size.

B Appendix B: Proof of Proposition 1

It is immediate that if $\mathcal{E}^*$ is predictable, it is also verifiable through experience (and thus verifiable in general), since participants are ineligible for all alternatives except one in their cumulative actions, and this single alternative is not only their assignment, but is also part of the true allocation. We now show the converse. Let $\mathcal{E}^*$ be verifiable through experience. Then g^* is sequential (direct mechanisms, unless they are constant mechanisms, cannot be verifiable through experience). In sequential mechanisms, since non-wastefulness requires each participant to be offered the unassigned option in each active stage, any such stage could be the last stage in which a participant is asked to take an action. Thus, verifiability through experience requires $|E_i^{g^*}(\epsilon_i^{g^*,\{k\}}) \cap Y(\succeq_i^{g^*,\{k\}})| = 1$ for any k in g^* and any $i \in Act^{g^*,k}$, since participant i might not be required to take another action. Moreover, if K is the terminal stage of g^*, then $E_i^{g^*}(\epsilon_i^{g^*,\{K\}}) \cap Y(\succeq_i^{g^*,\{K\}}) = a_i^*$. Together, these two conditions imply that $\mathcal{E}^*$ is predictable.

C Appendix C: Proof of Proposition 3

Let $\mathcal{E}^* = (g^*, \mathcal{M}^*)$ be monotonic-offers. Let $a^* = g^*(\phi^*_{g^*})$ be the realised allocation. Let $\mathcal{M}^*$ be a step-cutoff protocol. Thus $M^{g^*} = (c(a(u_s)))_{u_s \in U_{g^*}}$ is the collection of cutoffs for each temporary allocation in the steps in U_{g^*}. Since $\mathcal{M}^*$ contains terminal-cutoffs, it is verifiable by Proposition 2, and so a^* is justified. We show that a^* is the unique justified allocation in $\mathcal{E}^*$. Let $i \in N$, and let $x \in X$ such that $x >_i^* a_i^*$. Since $\mathcal{E}^*$ is non-wasteful, x is filled in a^*. By implication $x \neq \emptyset$.

Case 1: g^* is monotonically decreasing in offers, thus the eligibility set under cutoffs is weakly increasing in size across steps. In particular, there is a first step u_t in U_{g^*} where i is ineligible for x under the cutoff for x at that step, i.e., $x \notin E_i^{g^*}(c_x(a(u_t)))$. Since $\mathcal{E}^*$ is non-wasteful, x is filled in the temporary allocation $a(u_t)$ at this step. By construction, there are participants $J \subseteq N$ with $|J| = q_x$ who are assigned x in this temporary allocation who are eligible for x given $c_x(a(u_t))$. Thus $\rho_j^{g^*,x} > \rho_i^{g^*,x}$ for each $j \in J$. Given this inference, there is no problem $P \in c(a(u_t))$ compatible with this information, and no feasible allocation $a' \in \mathcal{A}$ with $a'_i = x$, such that $g^*(P) = a'$. Since eligibility under cutoffs is weakly increasing in steps, this is also true at the terminal step of the mechanism, and so there is no problem $P \in M^*$ and allocation $a' \in \mathcal{A}$ such that $g^*(P) = a'$. Since x is arbitrary, any allocation a in $\mathcal{E}^*$ that is justified is such that $a_i = a_i^*$. Since i is arbitrary, this is always true, and so the unique justified allocation in $\mathcal{E}^*$ is a^*. Thus $\mathcal{E}^*$ is strongly verifiable.

Case 2: g^* is monotonically increasing in offers, thus the eligibility set under cutoffs is weakly decreasing in size across steps. In particular, there is a first step u_t in U_{g^*} where i is ineligible for x at that step and all future steps, i.e., $x \notin E_i^{g^*}(c_x(a(u_{t+k})))$ for all $k = 0, 1, 2 \ldots$. Pick any of these k. Since $\mathcal{E}^*$ is non-wasteful, x is filled in $a(u_{t+k})$. By construction, there are participants $J \subseteq N$ with $|J| = q_x$ who are eligible for x under $c_x(a(u_{t+k}))$. By definition of g^*, x is in the core eligibility for each such agent at this step, i.e., $x \in C_j^{g^*}(c(a(u_{t+k})))$ for each $j \in J$. Thus $\rho_j^{g^*,x} > \rho_i^{g^*,x}$ for each $j \in J$. Given this inference, there is no problem $P \in c(a(u_{t+k}))$ and no feasible allocation $a' \in \mathcal{A}$ with $a'_i = x$ such that $g^*(P) = a'$. Since k is arbitrary, this is true for all such k, and there is no problem $P \in M^*$ and allocation $a' \in \mathcal{A}$ such that $g^*(P) = a'$. Since x is arbitrary, any justified allocation a in $\mathcal{E}^*$ is such that $a_i = a_i^*$. Since i is arbitrary, this is always true, and so the unique justified allocation in $\mathcal{E}^*$ is a^*. Thus $\mathcal{E}^*$ is strongly verifiable.

D Appendix D: Proof of Theorem 1

Let $\mathcal{E}^* = (g^*, \mathcal{M}^*)$ be the promised environment. Let $a^* = g^*(\phi^*_{g^*})$ be the true allocation. Suppose the designer uses a plausible environment $\mathcal{E} = (g, \mathcal{M})$ to reach an allocation $a \in \mathcal{A}$. If $\mathcal{M}^*$ is full disclosure, then $\mathcal{M}$ is full disclosure, by indistinguishability. Moreover, $M = M^*$, otherwise some participant can spot that her preferences or scores were not accurately represented. Then $\phi^*_g = \phi^*_{g^*}$, and as $\phi^*_{g^*} \in \Phi^1_{g^*}$, it follows that a^* is the unique justified allocation, and $\mathcal{E}^*$ is transparent.

Let $\mathcal{E}^*$ be monotonic-offers and verifiable through each of experience and communication. By Proposition 1, $\mathcal{E}^*$ is predictable. Under the mechanism g, the designer acquires information $\phi^*_g \in \Phi^1_g$, such that $a \equiv g(\phi^*_g)$. To show that $\mathcal{E}^*$ is transparent, we must show that $a = a^*$. Recall that the plausibility of $\mathcal{E}$ has the following implications:

1. The realised allocation a is non-wasteful with respect to any problem $P \in \phi^*_g$.
2. The environment $\mathcal{E}$ is indistinguishable from $\mathcal{E}^*$. Thus offer sets are monotonic and permitted actions in g are as in g^*. Also, $\mathcal{M}$ is the same type as $\mathcal{M}^*$, and a is justified in $\mathcal{E}$. In particular, for each $i \in N$: (1) $a_i \in E_i^{g^*}(M \cap \phi_i^{*g})$; and (2) for each $x >_i^{*g} a_i$: $x \notin E_i^{g^*}(M \cap \phi_i^{*g})$.

We will show that these properties ensure that g is 'essentially equivalent' to g^*, in that g cannot materially alter from g^* either the selection of active participants in a stage or the sets of alternatives offered to them. Given our behavioural assumption that participants take the same permitted action when faced with the same offer set, this means the designer acquires information $\phi^*_g = \phi^*_{g^*}$. As a result, the only justified allocation in $\mathcal{E}$ is a^*.

Since $\mathcal{E}^*$ is predictable, the indistinguishability of g implies that there is a single alternative in each participant's actions for which she is eligible under her experiential information at any stage. That is, $|E_i^{g^*}(\epsilon_i^{g,\{k\}}) \cap Y(\geq_i^{g,\{k\}})| = 1$ for each stage k in g and each $i \in Act^{g,k}$. Set $b_i^k \in X$ as this unique alternative, and call it her 'temporary assignment' at this stage. For each $j \notin Act^{g,k}$, her temporary assignment carries over from the last previous stage in which she was active, and is $\emptyset$ otherwise. The 'temporary allocation' at stage k in g collects these temporary assignments, and is given by $b^k \equiv (b_i^k)_{i \in N}$. Note that b^k is not necessarily feasible. However, at the terminal stage K in g, predictability ensures that $b^K = a$.

Case 1: g^* is monotonically decreasing in offers: We first show that offer sets in g are exactly as in g^*.

Lemma **1.** *In any stage in g, any active participant is offered all alternatives in her eligibility set under the designer's information at that stage that she has not already ranked. Formally, for any $i \in Act^{g,k}$ for a stage k in g, $x \in Y_i^{g,k}$ for any $x \in E_i^{g^*}(\phi^{g,\{k-1\}})$ such that $x \notin Y(\succeq_i^{g,\{k-1\}})$.*

Proof: Suppose not. Then there is a stage k in g, a participant $i \in Act^{g,k}$ and an alternative $x \in E_i^{g^*}(\phi^{g,\{k-1\}})$, such that $x \notin Y(\succeq_i^{g,\{k-1\}})$ and $x \notin Y_i^{g,k}$. Then $x \notin Y_i^l$ for all $l > k$ (x cannot be offered to i in any future stage, since g^* is monotonicall decreasing in offers and g is indistinguishable from g^*).

Suppose x is not filled in the temporary allocation b^{k-1}. We show that this could lead to non-wastefulness. Since x has not yet been ranked by i, the designer does not know i's preferences over x. Thus there is a problem $P = (\succeq, v^*) \in \phi^{g,\{k-1\}}$ where x is top-ranked in $\succeq_i$, and moreover x is unacceptable to any participant j who has not ranked x in any previous stage (this is possible because the designer does not know j's preferences over x as she has not ranked it yet). P is compatible with the designer's information $\phi^{g,\{k-1\}}$. By non-wastefulness, none of these participants j can be assigned x as the outside option is preferred by each of them and is unfilled by definition. Thus $a_j \neq x$ for each such j. Since x is top-ranked by i and is unfilled, non-wastefulness implies that $a_i = x$. However, since x cannot be offered to i in any future stage, it cannot be ranked by her, and then predictability implies $a_i \neq x$. This 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 could lead to a communication that does not justify the reached allocation. Since i is eligible for x under the designer's information, it follows that there is at least one $j \in J$ such that $\rho_i^{g^*,x} \geq \rho_j^{g^*,x}$. However, since the designer does not know i's preferences over x, there is a problem $P = (\succeq, v^*) \in \phi^{\{k-1\}}$ in which (1) x is top-ranked in $\succeq_i$ and no alternative that has not yet been offered to her is acceptable to her; and (2) for each $j \neq i$, there is no other alternative that has not yet been offered to them that is acceptable to her. P is compatible with the designer's information $\phi^{g,\{k-1\}}$. Thus in effect k is the terminal stage of the mechanism. By predictability, $a_j = x$ for all j such that $b_j^{k-1} = x$. Since a is justified in $\mathcal{E}$, j is eligible for x according to the communication M, i.e., $x \in E_j^{g^*}(M \cap \phi_j^{*g})$. But M is common, and $\rho_i^{g^*,x} \geq \rho_j^{g^*,x}$, and so $x \in E_i^{g^*}(M \cap \phi_i^{*g})$. In particular, $E_i^{g^*}(M \cap \phi_i^{*g}) \neq a_i$, which violates verifiability through commmunication. ■

By Lemma 1, any active participant in the first stage in g is offered all alternatives, since none are filled in $b^0 = \emptyset$. In any subsequent stage $l > 1$ with $i \in Act^{g,l}$, $x \notin Y_i^{g,l}$ only if x is filled in b^{l-1} and it is not in her eligibility set, i.e., $x \notin E_i^{g^*}(g, \phi^{\{l-1\}})$. This proves that offer sets in g are exactly as in g^*.

It remains to be shown that the designer cannot materially alter the selection of active participants in any stage either. We say that i is 'rejected' from x in g if there is a stage k such that she is temporarily assigned x in b^{k-1}, is made active in stage k, and is not offered x in stage k. That is, i is rejected from x in g if there is a stage k in g such that $b_i^{k-1} = x$, $i \in Act^{g,k}$, and $x \notin Y_i^{g,k}$. Let k be a stage in g and let the temporary allocation at $k-1$ be b^{k-1}.

Case 1.1: Let $x \in X$ be unfilled or exactly filled in b^{k-1}. We show that no participant can be rejected from x in g. Similar to the logic in Lemma 1, there is a problem $P = (\succeq, v^*) \in \phi^{g,\{k-1\}}$ compatible with the designer's information in which x is ranked higher in $\succeq_i$ than any alternative she has not ranked so far, and any participant j who has not ranked x so far ranks it below the unassigned option. By non-wastefulness, $a_i = x$, which is not possible as x cannot be offered to i in a later stage, and predictability ensures that $a_i \neq x$, which is a contradiction.

Case 1.2: Let $x \in X$ be overfilled in b^{k-1}. There is no feasible allocation that assigns each participant this alternative (due to capacity constraints), so the designer must reject some participants. Rejected participants in g^* are exactly those participants temporarily assigned x in b^{k-1} who are excess to capacity and have the lowest scores. Call this set $B(b^{k-1})$. We show that $B(b^{k-1})$ is exactly the set of participants who are rejected from x in g. As in Case 1, the designer cannot reject 'extra' participants, leaving x unfilled, because of the possibility of waste. Suppose the designer rejects some i in g who does not have one of the lowest scores, i.e., $i \notin B(b^{k-1})$. By predictability, i cannot be assigned x at a later stage, and so $a_i \neq x$. Moreover, $\rho_i^{g^*,x} \geq \rho_j^{g^*,x}$ for all $j \in B(b^{k-1})$ by construction of $B(b^{k-1})$. There are two possibilities:

1. All agents assigned x in b^{k-1} retain this assignment in the final allocation a, i.e., $a_j = x$ for all $j \in B(b^{k-1})$. Since a is justified in $\mathcal{E}$ by indistinguishability, it follows that for any $j \in B(b^{k-1})$ such that $a_j = x$, $x \in E_j^{g^*}(M \cap \phi_j^{*g})$. Since $\rho_i^{g^*,x} \geq \rho_j^{g^*,x}$, it follows that $x \in E_i^{g^*}(M \cap \phi_i^{*g})$. Thus $x \in E_i^{g^*}(M \cap \phi_i^{*g})$, which violates verifiability through communication.

2. Some agent j in $B(b^{k-1})$ is rejected from x at a future stage. Then, by predictability, $a_j \neq x$. As before, there is a problem $P = (\succeq, v^*) \in \phi^{\{k-1\}}$ compatible with the designer's information where x is ranked higher by i than anything she has not yet ranked, and no other participant yet to rank x finds it acceptable. Since P could be the true allocation problem, this creates waste, as x is unfilled in a and i prefers x to her assignment (which cannot be x by predictability). This is a contradiction.

Thus, both possibilities lead to a contradiction, and so rejections in g are the same as in g^*. As the selection of active participants in g is therefore essentially the same as in g^* (the only possible difference between g and g^* is the stage in which these rejections are performed), and the offer sets in both mechanisms are the same, our behavioural assumption on actions implies that participants take the same actions in both mechanisms. Thus, $b^K = a^*$ where K is the terminal stage of g. Moreover, by predictability, it follows that $b^K = a$. Thus $a = a^*$, proving the result.

Case 2: g^* is monotonically increasing in offers: Recall that the indistinguishability of g from g^* implies that if a participant is offered some alternatives in a stage in g, she is offered these alternatives in all future stages in which she is active.

Lemma **2.** *In any stage of g, an alternative cannot be offered to a participant for whom it is not in her core eligibility at that stage. Formally, for any stage k in g, and any active participant $i \in Act^{g,k}$, $x \notin C_i^{g^*}(\phi^{g,\{k-1\}}) \implies x \notin Y_i^{g,k}$.*

Proof: Suppose for contradiction that there is a stage k in g, an active participant $i \in Act^{g,k}$, and an alternative $x \notin C_i^{g^*}(\phi^{g,\{k-1\}})$, such that $x \in Y_i^{g,k}$.

Claim **1.** The alternative x is unfilled in b^{k-1}.

Proof: Suppose not. Then there is a problem $P = (\succeq, v^*) \in \phi^{g,\{k-1\}}$ compatible with the designer's information in which each participant j such that $b_j^{k-1} = x$ ranks x above any alternative she has not yet ranked, and so does participant i. By predictability, $E_j^{g^*}(\epsilon_j^{g,\{k\}}) \cap Y(\succeq_j^{g,\{k\}}) = x$. Monotonicity of g^* and indistinguishability of g from g^* ensures that x is offered to each such j in all stages $l > k$ in which she is active. Since no alternative not already offered to and ranked by each such j is acceptable, predictability ensures that $a_j = x$ for each such j. Moreover, since x is top-ranked by i in $\succeq_i$ to all alternatives not yet offered and ranked by her, predictability ensures that $a_i = x$. This violates capacity constraints, as more participants are assigned x in a than its capacity. ■

Claim 1 establishes that x is unfilled in b^{k-1}. Since $x \notin C_i^{g^*}(\phi^{g,\{k-1\}})$, by definition there is some set of participants B such that $x \in C_j^{g^*}(\phi^{g,\{k-1\}})$ for each $j \in B$, such that j has not yet ranked x. Moreover, $\rho_j^{g^*,x} \geq \rho_i^{g^*,x}$ for each $j \in B$. There is a problem $P = (\succeq, v^*) \in \phi^{g,\{k-1\}}$ compatible with the designer's information in which each j ranks x higher than any alternative she has not yet ranked, and so does participant i. By predictability, x is the assignment of i in the temporary allocation at stage k. If x is offered to each $j \in B$ at some stage, it follows that $a_j = x$ for each $j \in B$, which violates capacity constraints. Thus there is some participant $j \in B$ who is not offered x at any stage. Thus, by predictability, $a_j \neq x$. Since a is justified in $\mathcal{E}$, and since $\rho_j^{g^*,x} \geq \rho_i^{g^*,x}$, $x \in E_j^{g^*}(M \cap \phi_j^{*g})$, which violates verifiability through communication, as $a_j \neq x$. ■

By Lemma 2, each alternative x is offered only to participants for whom it is in their core eligibility at that stage. In particular, in the first stage of g, each active participant is offered some subset of her initial core eligibility, and in subsequent stages is offered all other such alternatives, before they are offered to other participants. Thus there is a stage in which each participant has been offered all alternatives in her core eligibility, before any of them is offered to any other participant. Under our behavioural assumption, her action in this stage coincides with her action in g^* at the corresponding stage. Since it is without loss of generality that previously offered alternatives continue to be offered, it follows that offers are monotonically increasing and consistent with g^*. Moreover, non-wastefulness ensures that each alternative is offered to some participant as long as it is unfilled. By Lemma 2, these offers are in the order induced by core eligibility. Thus the selection of active participants is also as in g^*.

As the selection of active participants in g is the same as in g^*, and the offer sets in both mechanisms are essentially the same, our behavioural assumption on actions implies that participants take the same cumulative actions in both mechanisms. Thus $b^K = a^*$ where K is the terminal stage of g. Moreover, by predictability, it follows that $b^K = a$. Thus $a = a^*$, proving the result.

E Appendix E: 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 the 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,[42] and under SeqNo and SeqCutoffs she follows the ***straightforward strategy*** of applying to the best school among the ones that are still available (restricted either by previous rejections or by too-high intermediate cutoff grades).

Table 4 presents the proportion of truthful strategies by treatments in total and for the first five and the last five rounds of the experiments separately. First, the proportion of truthful strategies is quite

[42]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.

Table 4: Proportions of truthful strategies

	DirNo (1)	DirCutoffs (2)	SeqNo (3)	SeqCutoffs (4)	1=2 p-val.	1=3 p-val.	1=4 p-val.	2=3 p-val.	2=4 p-val.	3=4 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.

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%.[43] 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 (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

[43] 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.

the student-optimal-stable match.[44] 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 (1)	DirCutoffs (2)	SeqNo (3)	SeqCutoffs (4)	1=2 p-val.	1=3 p-val.	1=4 p-val.	2=3 p-val.	2=4 p-val.	3=4 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.

F Online Appendix: Experiment Instructions

F.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 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.

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

- 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.

F.2 Treatment-Specific Instructions

F.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

F.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

F.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 exam grades 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

F.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 exam marks 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

F.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 new grades, and

  – generate new 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.