School Choice with Appeals

CLAUDIA CERRONE
City St George's, University of London

YOAN HERMSTRÜWER
University of Zurich

JOSUÉ ORTEGA
Queen's University Belfast

More than fifty thousand families in England appeal their school assignments each year through a centralized and understudied procedure that improves the placement of roughly one in five claimants. This paper shows that appeals are not an institutional afterthought: they change how parents report preferences to the education authority and the quality of placements it generates. Using a theoretical model and a laboratory experiment, priority-based appeals raise truth-telling and student welfare under Immediate Acceptance while leaving Deferred Acceptance essentially unchanged. Probabilistic appeals also improve Immediate Acceptance, but they reduce the stability of Deferred Acceptance. Once the appeal stage is incorporated, Immediate Acceptance can support equilibrium outcomes that Pareto-dominate those produced by Deferred Acceptance, reshaping the DA–IA comparison and showing that the welfare ranking of matching mechanisms critically depends on the appeal institutions that follow them.

KEYWORDS. appeals, school choice, aftermarkets.

JEL CLASSIFICATION. D47, C78, C92, I28.

1. INTRODUCTION

School choice is the process through which students are assigned to schools via a centralized clearinghouse that asks and uses parents' preferences to determine placements. In systems as different as England, Northern Ireland, and New York City, the assignment procedure proceeds in three steps. First, parents rank the available schools. Second, the clearinghouse uses an algorithm to produce an allocation. Finally, parents may appeal each rejection they received from a school they ranked, and an independent panel decides which appeals to uphold.

Economists have profoundly shaped the first part of this process. A large literature studies how families should rank schools, which matching mechanisms a clearinghouse should run, and how priorities should enter centralized admissions, and this work has

Emails: claudia.cerrone@city.ac.uk, yoan.hermstruewer@ius.uzh.ch, j.ortega@qub.ac.uk.

We thank Sophie Bade, Estelle Cantillon, Yan Chen, Arcangelo Dimico, Arda Gitmez, Rustam Hakimov, Flip Klijn, Sander Onderstal, Bertan Turhan, M. Utku Ünver, and M. Bumin Yenmez for valuable comments. We are grateful to John Coldron, Chris Taylor, and Camille Terrier for helpful discussions. *Last edited on 21/08/2026.*

driven major reforms in school assignment systems around the world. Yet, the final appeal stage has received almost no attention. The omission matters. In England alone, tens of thousands of families appeal secondary-school placements each year, and appeals succeed often enough to reshape outcomes: roughly one in five appeals is upheld, and about one in seven families rejected from their first-choice school ultimately gains admission to a more preferred school through this route (Hunt, 2019).

In this paper, we explore how appeal rights change the economics of school choice. We first ask: how do appeals change how parents rank schools? And second, how do differences in the way parents rank schools affect the final student placement? These questions are not merely theoretical. Families and school officials already understand that the possibility of a later appeal affects how families should rank schools. In response to a parent asking how to order ambitious grammar-school applications and a safer alternative, one school-admissions forum advised:

> "*[Put] whichever Berks school you want ... higher than Burnham, so that you will then be able to appeal when you are refused a place.*"[1]

This advice captures the mechanism we study. The parent is not told to abandon a safe option. She is told to rank the ambitious school above it, so that rejection preserves an appeal right while the lower-ranked school remains a fallback. A rank-order list therefore serves two roles. It determines the initial assignment, and it determines which rejected schools can later be challenged. Appeals can change behavior before the assignment is made, as well as outcomes after it is made.

We model school choice with appeals as a two-stage process. In the first stage, students submit rank-order lists and either Immediate Acceptance (IA) or Deferred Acceptance (DA) produces an initial assignment. In the second stage, students may appeal to schools that they ranked above their assignment. This structure provides a stylized representation of the statutory process in England, as mandated by the nationwide regulations in the School Admission Appeals Code.[2] Our strict rule captures the mandatory correction of a violation of admissions criteria. A second, probabilistic, rule provides a simple representation of the uncertainty surrounding the subsequent balancing stage. It is motivated by the wide variation in observed appeal success rates, which suggests idiosyncratic interpretations of the Appeals Code across local authorities. Successful appeals add seats rather than remove students admitted in the first stage, consistent with the Code's treatment of the remedy as the admission of an additional child. The final allocation is therefore determined jointly by the assignment mechanism and the appeal rule.

[1] Source: https://www.elevenplusexams.co.uk/forum/11plus/viewtopic.php?t=44895. Accessed on July 22, 2026.

[2] School Admission Appeals Code 2022, paragraphs 3.5 to 3.10. The Code first requires the panel to determine whether the admission arrangements were unlawful or incorrectly applied and states that "the panel must uphold the appeal" when the child would otherwise have been offered a place. If the appeal is not upheld on those grounds, the panel must balance the family's case against the prejudice that admitting an additional child would cause the school. Department for Education statistics show substantial geographic variation in appeal outcomes, with success rates ranging from zero in some local authorities to above 40 percent in others in recent years.

This two-stage structure reflects a feature of institutional design that mechanism design typically abstracts away. The designer of a matching mechanism is implicitly the legislator or the administration: a unitary authority that sets the rules, applies them, and has the final word. Constitutional law separates these functions (Montesquieu, 1989, Locke, 1988, Hamilton et al., 2003) (see *Marbury v. Madison*, 5 U.S. (1 Cranch) 137 (1803)). Allocations made by public authorities are reviewable by an independent institution or branch of government – typically, the judiciary – that did not design the mechanism and can revise its outcomes. The mechanism is then no longer the whole institution: the allocation is produced by the (legislative or executive) mechanism and the (judicial) review process jointly, and the designer usually controls only the first.

Our first theoretical result establishes a sharp asymmetry in how parents' incentives are affected by appeals in the two assignment mechanisms. Strict appeals preserve the strategy-proofness of Deferred Acceptance. Since DA produces an assignment with no strict blocking pair relative to submitted preferences, the strict appeal stage has no priority violation to correct. Under Immediate Acceptance, the same rule changes incentives. A standard manipulation in IA is to skip applications to unfeasible schools because they harm chances at feasible second-choice schools that may be filled with lower-priority applicants by the time the second-choice school receives an application; this reasoning disappears once appeals are allowed, as a student can always overturn a priority violation through the appeal aftermarket. Appeals therefore make ambitious applications less risky. Although they do not make IA strategy-proof, they make it less manipulable in the profile-by-profile sense of Pathak and Sönmez (2013): every preference profile at which IA with strict appeals can be profitably manipulated is also manipulable under IA without appeals, while the converse fails at some profiles. The probabilistic rule adds a second channel that promotes truth-telling in IA: skipping a seemingly unattainable school throws away a lottery ticket through which the student might have gained admission.

Our main theoretical result concerns welfare under strategic behavior. Without appeals, Ergin and Sönmez (2006) show that every Nash equilibrium of the preference revelation game in Immediate Acceptance is weakly Pareto-dominated by truthful Deferred Acceptance. We show that this comparison can reverse once strict appeals are incorporated. For every school choice problem with strict priorities, Immediate Acceptance followed by strict appeals admits a Nash equilibrium whose outcome weakly Pareto-dominates truthful Deferred Acceptance, and makes at least one student strictly better off whenever the two outcomes differ. This is an equilibrium-existence result, not a claim that every equilibrium of IA with appeals dominates DA. Its implication is that the classical comparison between the two mechanisms depends on the review process that follows the initial assignment.

We test these predictions in a preregistered laboratory experiment with 1,022 participants across 59 independent matching groups. We vary both stages of the process. Participants face either IA or DA, followed by no appeals, strict appeals, or probabilistic appeals. The design lets us observe initial reports, first-stage assignments, appeal choices, and final assignments. We can therefore separate changes in behavior before any appeal is decided from the direct effect of successful appeals afterwards.

The reporting results closely follow the theoretical asymmetry. Under IA, strict and probabilistic appeals each raise truth-telling by approximately 11 percentage points, from 28 to 39 percent. Under DA, the corresponding effects are only 3 to 4 percentage points and are at most marginally significant. The meaning of a non-truthful report also differs sharply across mechanisms. Under DA, approximately 92 percent of non-truthful reports leave the participant's assignment unchanged, and every consequential deviation is harmful. Under IA, almost two thirds of non-truthful reports change the assignment, although harmful deviations are more common than beneficial ones. Appeals therefore reduce a behaviorally important strategic problem under IA, while most deviations observed under DA are assignment-irrelevant.

Appeals also substantially improve student welfare under IA. Strict appeals raise welfare under IA by approximately 0.3 ranks while leaving DA essentially unchanged. They increase the share of realized IA market outcomes that strictly Pareto-dominate truthful DA from 11 to 48 percent. The experiment does not reproduce a reversal in the sign of the average welfare ranking, because IA already performs better than DA without appeals in this environment. Instead, it confirms the comparative-static prediction of the theory: strict appeals improve IA much more than DA and move realized IA allocations sharply toward the Pareto region above truthful DA.

A decomposition of the welfare effect shows that the gains do not arise only because successful appeals expand school capacities. Under IA with strict appeals, approximately 60 percent of the observed rank improvement is already present in the first-stage assignment, before any appeal is decided. The remaining 40 percent accrues between the first-stage and final assignments. This chronological decomposition shows that appeal rights affect outcomes before as well as after claims are resolved, but it does not assign unique causal shares to reporting behavior and the appeal stage.

Participants use the appeal stage systematically, although imperfectly. They are more likely to pursue claims against more valuable schools and claims supported by strict priority. At the same time, they file many appeals that cannot succeed and leave approximately three in ten sure, payoff-improving claims unfiled. This incomplete use of appeals helps explain why strict appeals reduce, but do not fully eliminate, the number of blocking pairs in IA.

The standard of review matters. Under IA, strict appeals more than halve the number of strict blocking pairs and substantially increase the share of stable market outcomes while creating almost no waste. Under DA, strict appeals leave the stable first-stage assignment unchanged. Probabilistic appeals are less targeted. They improve IA, but under DA they can disturb an initially stable assignment, create vacancies, and generate new blocking pairs.

Our contribution is to show that a centralized assignment procedure cannot always be evaluated independently of the institutions that review its decisions. When families can challenge an assignment, the review rule changes both the final allocation and the incentives governing the initial report. The relevant object of analysis is therefore the

assignment procedure together with the appeal rule that follows it. More broadly, whenever one institution can revise an allocation produced by another, the design of the review process can alter the usual ranking of assignment mechanisms. Separation of powers, as a fundamental principle of constitutional law, is therefore not a detail that can be added after the fact to a mechanism a priori designed without it. Because families anticipate administrative or judicial review, its availability shapes the reports on which the initial assignment is built. A mechanism that performs well when its output is final may perform differently when that output is subject to independent review.

The remainder of the paper proceeds as follows. Section 2 discusses the related literature. Section 3 describes the institutional setting. Section 4 presents the model and theoretical results. Section 5 describes the experimental design, and Section 6 presents the results. Section 7 concludes.

2. Related Literature

Our work contributes to several strands of literature on school choice and aftermarkets. We organize our discussion around three interconnected themes.

Empirical literature. Taylor et al. (2002) show that appeals are more common in urban, more affluent, and higher-demand areas. Hunt (2019) finds that disadvantaged and minority pupils are less likely to secure a top-choice school through appeals, while parents who submit incomplete rank-order lists are 15% more likely to appeal. Successful appeals also tend to place students in less deprived schools with fewer disadvantaged peers, suggesting that the appeals process may disproportionately benefit advantaged families.

Models with appeals. Only a handful of papers explicitly examine the role of appeals. Kojima (2011) shows that allowing students to appeal their non-admission can compromise the stability and strategy-proofness of the DA mechanism. He introduces the concept of robust stability, arguing that if appeals are unrestricted, it becomes a weakly dominant strategy for students to forgo ranking schools initially and later appeal to their most preferred school with available spots (a result further generalized by Afacan (2012)). Our model builds on this framework but introduces a crucial distinction: students can only appeal to schools they initially ranked and were rejected from. This restriction, aligned with real-world systems in England and Northern Ireland, prevents the strategic behavior identified by Kojima and ensures that appeals serve as a secondary step rather than a complete circumvention of the initial matching process.

Afacan et al. (2017) propose a different appeals model with two key differences: (1) appealing is costly, and (2) instead of a post-assignment appeal process, parents preemptively report how much instability they are willing to tolerate. They introduce the concept of sticky stability, where priority violations up to a fixed threshold are permitted, ensuring that every stable mechanism is also sticky stable. Unlike Afacan et al. (2017), we do not model appeals as an ex-ante tolerance for instability, but as an ex-post institutional process that follows the initial matching.

Appeals also appear in the experiment of Hakimov and Raghavan (2026), who study transparency and verifiability in centralized admissions. There, appeal decisions serve as an experimental measure of whether participants can detect that the announced mechanism was not followed, rather than as a second-stage allocation rule. Our paper instead treats appeals as part of the mechanism itself and studies how alternative rules for upholding them affect reporting incentives and final outcomes.

More recently, Decerf et al. (2024) study when a matching is incontestable, that is, appeal-proof for students who observe priorities and capacities but not other students' preferences or assignments. Their focus is on the informational foundations of a weakened notion of stability. Our paper asks a different question. Rather than characterizing when a final matching is appeal-proof, we model appeals as an actual institutional second stage and study how alternative uphold rules reshape incentives and final outcomes.

Immediate versus Deferred Acceptance. Our welfare result also contributes to the longstanding comparison between IA and DA. Under complete information and without appeals, every Nash equilibrium outcome of IA is weakly Pareto-dominated by truthful DA (Ergin and Sönmez, 2006). This ranking is not robust to uncertainty and cardinal welfare: IA can outperform DA under incomplete information (Miralles, 2009, Abdulkadiroğlu et al., 2011, Akyol et al., 2026). Structural studies using administrative data from Barcelona, Beijing, and Cambridge similarly find conditions under which IA yields higher estimated average welfare than DA, although the comparison can depend on assumptions about sophistication and beliefs (Calsamiglia et al., 2020, He, 2015, Agarwal and Somaini, 2018). Our paper identifies a different source of reversal: even under complete information, the appeal rule following the first-stage mechanism can overturn the classical Pareto comparison.

Beyond preference revelation. Our paper also connects to a broader literature that views school choice as a process extending beyond preference submission and the initial matching outcome. A first strand studies multi-stage admission procedures involving interviews, information acquisition, and other decisions before the matching takes place (Fu et al., 2024, Artemov, 2021, Chen and He, 2021, 2022, Maxey, 2024). A second strand studies aftermarkets and outside options, emphasizing that the assignment produced by a centralized clearinghouse need not coincide with the school or program a student ultimately attends (Akbarpour et al., 2022). Most closely related to our institutional perspective, Kapor et al. (2024) show that off-platform offers, declined placements, and frictional waitlists can alter final enrollment and generate externalities for other participants. Appeals similarly create a post-matching stage in which initial assignments may be revised. Unlike the aftermarket they study, however, appeals are governed by formal standards of review and are anticipated when families submit their rank-order lists; they therefore affect both final assignments and first-stage reporting incentives. Finally, our study relates to work allowing school capacities to adjust or be exceeded under specified conditions (Aygün and Bilgin, 2023, Afacan et al., 2024).

3. Institutional Setting

Since the 1998 Education Reform Act, both England and Northern Ireland have maintained clearly defined appeals procedures for school admissions (Taylor et al., 2002). When a child is refused a school place, parents receive a letter explaining how to appeal. The process is free and conducted by independent panels, with each school refusal requiring a separate appeal. However, the criteria for upholding appeals differ between England and Northern Ireland. In England, an appeal must be upheld if either:

1. The school's admissions criteria were not properly followed or do not comply with the school admissions code.
2. The parents' reasons for admission outweigh the school's justification for not accepting additional students.[3]

On the other hand, in Northern Ireland appeals are upheld only if:

1. The school did not apply, or misapplied, its admissions criteria.
2. The child would have been offered a place had the criteria been correctly applied.

This regional difference captures the distinction between our theoretical uphold rules: England's second criterion allows discretionary decisions based on weighing competing claims, while Northern Ireland's approach focuses purely on procedural compliance.

Appeals concerning infant classes, the first three years of primary school (Reception and Years 1 and 2, typically ages four to seven), are subject to a stricter legal standard when admitting an additional child would breach the statutory limit of 30 pupils per teacher. In 2024/25, 9.7% of appeals heard for infant classes were successful, compared with 28.3% for other primary classes.

Appeals operate at substantial scale and exhibit considerable geographic variation. For admissions at the start of the 2024/25 academic year, $14{,}584$ primary and $37{,}289$ secondary appeals were lodged in England. Of the appeals heard, 17.7% of primary and 19.9% of secondary appeals were successful. Variation across the 152 local authorities was much sharper: among authorities that heard at least one appeal, primary success rates ranged from zero to 66.7%, with 8 of 12 appeals successful in Westminster, while secondary success rates ranged from zero to 95.1%, with 39 of 41 appeals successful in Sunderland.

Outside the UK, the New York City High School match also combines centralized admissions with a formal appeals process. Successful appeals are processed using a variant of the top trading cycles mechanism (Morrill and Roth, 2024); in the 2003/04 academic year, approximately 5.7% of applicants, or $5{,}100$ out of $90{,}000$, lodged an appeal, about half of which were successful (Abdulkadiroglu et al., 2005). Our model more closely reflects the systems in England and Northern Ireland, where independent panels review admission decisions and successful appeals admit an additional student rather than displacing an incumbent, with substantial variation within local authorities. These institutional features map naturally onto our model and appeal rule design in the next Section.

[3] For England, see here. For Northern Ireland, see here. Accessed on August 29, 2024.

4. MODEL

We model school choice with appeals as a two-stage process. In the first stage, students submit preference rankings to a centralized mechanism, which produces an initial matching. In the second stage, students may appeal to schools they ranked above their initial assignment, and an uphold rule determines which appeals are successful. Successful appeals create additional seats rather than displacing initially assigned students.

The first stage is a standard many-to-one matching problem, defined by the following primitives:

- a set of students I,
- a set of schools S, each able to accept q_s students, with the school $s_\emptyset$ of capacity $q_{s_\emptyset} = \infty$ representing being unassigned,
- a strict preference $\succ_i$ of each student i over schools,
- a weak school priority $\trianglerighteq_s$ of each school s, so that some students may have the same priority at some school, and
- a single, strict tie-break order τ over students, which resolves priority ties, inducing a strict priority order.

We use $\succsim_i$ to denote the weak preference relation induced by $\succ_i$. A school choice problem is a tuple $(I, S, q, \succ, \trianglerighteq, \tau)$. Since all primitives except student preferences are fixed throughout, we represent a school choice problem simply by $\succ$.

A matching μ is a function that assigns each student to a school or to the outside option. The school assigned to student i in μ is denoted by μ_i, whereas the set of students assigned to s in μ is denoted by μ_s^{-1}. A matching is feasible if no school is assigned more students than its capacity.

A weak blocking pair at matching μ is a student-school pair (i, s) such that $s \succ_i \mu_i$ and $i \trianglerighteq_s j$ for some $j \in \mu_s^{-1}$. A strict blocking pair is a weak blocking pair in which the priority comparison is strict, i.e., $i \triangleright_s j$ for some $j \in \mu_s^{-1}$. A matching μ is non-wasteful if there is no student i and school s such that $s \succ_i \mu_i$ and $|\mu_s^{-1}| < q_s$. A matching is stable if it admits no strict blocking pair and is non-wasteful.

A (not necessarily feasible) matching ν weakly Pareto-dominates matching μ if, for every student i, $\nu_i \succsim_i \mu_i$, and Pareto-dominates it if additionally $\nu_j \succ_j \mu_j$ for some student j.

A mechanism M associates a feasible matching to every profile of preferences $\succ$, and we use $M_i(\succ)$ to denote the school to which student i is assigned under mechanism M. We allow submitted preference lists to be truncated in the model, with omitted schools treated as unacceptable. Due to their popularity, we focus on the immediate acceptance (IA, aka the Boston mechanism) and the student-proposing deferred acceptance mechanism (DA). The IA mechanism works in sequential rounds, as follows:

- In the first round, each student applies to her most preferred school. Each school considers the students who have listed it as their first choice and assigns seats of the school to these students one at a time, following their priority order, until there are no seats left or there is no student left who has listed it as her first choice.

- In each subsequent round k, each remaining student applies to her k-th preferred school. Each school that was already full automatically rejects all new applicants. For each school with seats still available, consider the students who have listed it as their k-th choice and assign the remaining seats to these students one at a time following their priority order, until there are no seats left or there is no student left who has listed it as her k-th choice.

DA differs from IA in that schools' acceptances are temporary. In each round, rejected students apply to their next preferred school. Each school then considers its current tentative holders together with its new applicants, keeps the highest-priority students up to capacity, and rejects the rest. $\texttt{IA}(\succ)$ and $\texttt{DA}(\succ)$ denote the allocation produced by IA and DA for the matching problem $\succ$, respectively.

The Appeals Process. We now describe the second-stage process. Given the publicly observable reported preferences $\hat{\succ}$ and an initial matching $M(\hat{\succ})$, the set of appealable schools for student i includes all schools ranked as more desirable by the student to which she was not admitted, i.e.,

$$A_i(\hat{\succ}, M) \coloneqq \{s \in S : s \,\hat{\succ}_i\, M_i(\hat{\succ})\}.$$

Thus, a student may appeal only to schools that she ranked as more preferred, i.e., schools that rejected her in the mechanisms we consider. This restriction is aligned with the rules in England and Northern Ireland. After the initial matching is announced, student i submits a set of appeals $a_i \subseteq A_i(\hat{\succ}, M)$ to an independent panel. The panel then applies an uphold rule, which determines which of these appeals succeed.

DEFINITION 1. An uphold rule r maps each submitted set of appeals a_i into a subset $r(a_i) \subseteq a_i$, representing the appeals that are upheld.[4]

Because appeals are free for families and hearings are conducted online, we assume that appealing is costless. We focus on rules in which the success of an appeal at a school s is independent of whether a student decided to appeal her rejection at a different school s'. Under these assumptions, it is a weakly dominant strategy for each student to appeal every school in her appealable set that she truly prefers to her first-stage assignment. We therefore assume throughout that

$$a_i = \{s \in A_i(\hat{\succ}, M) : s \succ_i M_i(\hat{\succ})\}.$$

If more than one appeal is upheld, the student enrols in her most preferred school, according to $\succ_i$, among the upheld appeals.

We consider two uphold rules in the main text, described in Table 1.

[4]The uphold rule depends on the reported preferences, priorities and initial matching too, but we make this dependence implicit to simplify notation.

TABLE 1. Uphold rules

Rule	An appeal by student i to school $s \in a_i$ is upheld if
Strict rule (r_1)	(i, s) is a strict blocking pair in $M(\hat{\succ})$.
Probabilistic rule (r_2)	(i, s) is a strict blocking pair in $M(\hat{\succ})$, or otherwise with independent probability $p \in (0, 1)$.

Rule r_1, our main focus, captures the priority-based component of school admission appeals in England.[5] Specifically, England's School Admission Appeals Code (Section 3.5) states:

> "*The panel must uphold the appeal...if it finds that the admission arrangements did not comply with admissions law or had not been correctly and impartially applied, and the child would have been offered a place if the arrangements had complied or had been correctly and impartially applied.*"

Rule r_2 is motivated by the substantial variation in appeal success rates across local authorities, which range from 0% to over 40%, as well as by evidence from our survey of appeal panel members suggesting considerable heterogeneity in how similar cases are assessed (described in Appendix F). In the main text, we will focus on r_1, postponing the analysis of r_2 and r_3 to Appendices B and C, respectively, highlighting only the main differences in the text.

We denote the resulting composite mechanism by $M \circ r$. Given $\hat{\succ}$, it first applies M, then applies the uphold rule r to the appeals generated from $M(\hat{\succ})$, and finally assigns each student either her first-stage school or, if at least one appeal is upheld, her most preferred upheld appeal according to $\succ_i$. To place mechanisms with and without appeals on the same footing, let r_0 denote the null uphold rule that rejects every appeal. Hence $M \circ r_0$ coincides with the baseline mechanism M without appeals.

Note that the final matching need not be feasible: schools may admit more students than their published quotas when multiple appeals are upheld. This is consistent with real-world evidence. In England in 2023/2024, 24% of secondary schools were over capacity, while during the same period, 25% of post-primary schools in Northern Ireland exceeded their quotas due to successful appeals.[6]

4.1 *Preliminary Observations*

We now present two straightforward but useful observations regarding the consequences of appeals. First, in our framework, every student becomes weakly better off by

[5] In Appendix C, we consider a similar rule r_3, in which a student can successfully appeal if they were not admitted to a school but a student with the same priority was admitted due to a tie-break. This rule, unlike r_1 and r_2, breaks DA's strategy-proofness.

[6] Source: England's Explore Education Statistics: https://explore-education-statistics.service.gov.uk/find-statistics/school-capacity/2024-25, and NI Education Authority FAQs: https://www.eani.org.uk/parents/pupil-applications-and-grants/admissions/admissions-support. Accessed on May 13, 2026.

adding appeals, and the outcome after appeals are upheld admits no appealable strict blocking pair involving a school that the student truly prefers to her final assignment.

LEMMA 1. *For every mechanism M and every report profile $\hat{\succ}$, $(M \circ r_1)(\hat{\succ})$ weakly Pareto-dominates $(M \circ r_0)(\hat{\succ})$ with respect to true preferences, and admits no strict blocking pair (i, s) such that $s \in A_i(\hat{\succ}, M)$ and*

$$s \succ_i (M \circ r_1)_i(\hat{\succ}).$$

We postpone a straightforward proof to Appendix A. It is important to note that final outcomes need not be stable because the final matching can be wasteful: appeals may leave first-stage seats vacant that a student would prefer to her own, yet was unable to claim through an appeal because she had worse priority than all initially assigned students.

Lemma 1 implies an asymmetric effect of appeals in DA and IA. If DA is Pareto-dominated by a feasible matching, its outcome will remain inefficient after incorporating the appeals aftermarket. On the other hand, if IA generates appealable justified-envy violations involving schools that the affected students truly prefer, the appeals procedure removes those violations through capacity extensions, although the final outcome need not be stable because successful appeals may create vacancies elsewhere.

4.2 *Incentives*

We now focus on how appeals affect parents' incentives. It is well known that DA is strategy-proof whereas IA is not, which has long been a central argument in favor of the former. We show that the possibility of later appeals can modify these incentive properties, making IA less manipulable, simply because manipulations in IA are no longer needed in some problems where a more preferred school that could be obtained through a manipulation can now be obtained with a truthful report followed by a subsequent appeal.

Formally, fix a student i and a preference profile $(\succ_i, \succ_{-i})$. We say that student i can profitably manipulate the composite mechanism $M \circ r$ at $(\succ_i, \succ_{-i})$ if there exists a different preference $\hat{\succ}_i$ such that

$$(M \circ r)_i(\hat{\succ}_i, \succ_{-i}) \succ_i (M \circ r)_i(\succ_i, \succ_{-i}).$$

The composite mechanism is strategy-proof if no student has a profitable manipulation, for every possible preference $\succ_i$, manipulation $\succ_i'$ and preference profile $\succ_{-i}$. Mechanism N is *as manipulable as* mechanism M if, whenever some student can profitably manipulate M at some profile, that same student can profitably manipulate N at that profile. Mechanism N is *more manipulable* than mechanism M if, in addition, there exists a profile at which some student can profitably manipulate N but no student can profitably manipulate M (Pathak and Sönmez, 2013).

We first compare $\texttt{IA} \circ r_0$ and $\texttt{IA} \circ r_1$.

PROPOSITION 1. $\texttt{IA} \circ r_0$ *is more manipulable than* $\texttt{IA} \circ r_1$.

PROOF SKETCH. To show that every profitable manipulation of $\texttt{IA} \circ r_1$ is also a profitable manipulation of $\texttt{IA} \circ r_0$, fixing $\succ_{-i}$, note that under the manipulation student i obtains her improved school y either directly in the first stage or through an appeal. In the former case the same report is trivially profitable under $\texttt{IA} \circ r_0$. In the latter, the appeal succeeded because y admitted a student with lower priority than i; but then i, by ranking y first, would have been admitted to y in the first round under $\texttt{IA} \circ r_0$. Either way, i secures y under $\texttt{IA} \circ r_0$. By Lemma 1, her truthful $\texttt{IA} \circ r_1$ outcome is weakly preferred to her truthful $\texttt{IA} \circ r_0$ outcome. Since she strictly prefers y to the former, she also strictly prefers y to the latter. Example 1 below, with three students and three unit-capacity schools, illustrates how a manipulation opportunity (in parentheses for i_1) disappears in IA once appeals are incorporated.

EXAMPLE 1.

$\succ_{i_1}$	$(\succ'_{i_1})$	$\succ_{i_2}$	$\succ_{i_3}$	$\trianglerighteq_{s_1}$	$\trianglerighteq_{s_2}$	$\trianglerighteq_{s_3}$
s_1	(s_2)	s_1	s_2	i_2	i_1	·
s_2	(s_1)	s_3	s_3	i_1	i_3	·
s_3	(s_3)	s_2	s_1	i_3	i_2	·

Under truthful reporting, the first-stage IA outcome is

$$(i_1 - s_3,\ i_2 - s_1,\ i_3 - s_2).$$

Indeed, in round 1, i_1 and i_2 apply to s_1, and i_2 is accepted; i_1 is rejected. Student i_3 applies to s_2 and is accepted. In round 2, i_1 applies to s_2 but is rejected because s_2 is already full. In round 3, i_1 applies to s_3 and is accepted.

Under $\texttt{IA} \circ r_0$, student i_1 can profitably manipulate by reporting $s_2 \succ'_{i_1} s_1 \succ'_{i_1} s_3$. Then in round 1, student i_1 applies to s_2 and is accepted, i_2 applies to s_1 and is accepted, and i_3 applies to s_2 but is rejected. Student i_3 then applies to s_3 and is accepted. Hence the outcome is

$$(i_1 - s_2,\ i_2 - s_1,\ i_3 - s_3),$$

which is strictly better for i_1 than her truthful $\texttt{IA} \circ r_0$ assignment s_3.

Now consider $\texttt{IA} \circ r_1$ at the truthful profile. The first-stage outcome is again

$$(i_1 - s_3,\ i_2 - s_1,\ i_3 - s_2).$$

Student i_1 has a successful appeal to s_2, because i_1 ranks s_2 above her first-stage assignment s_3 and $i_1 \rhd_{s_2} i_3$. Thus, after appeals, i_1 obtains s_2. No student can profitably manipulate $\texttt{IA} \circ r_1$ at this profile. Students i_2 and i_3 obtain their top choices truthfully, so they cannot improve. Student i_1 obtains s_2 truthfully after appeals. Her only strictly better school is s_1. But i_2 ranks s_1 first and has higher priority than i_1 at s_1, so i_1 cannot obtain s_1 directly by any report and cannot successfully appeal to s_1. Therefore i_1 cannot profitably manipulate $\texttt{IA} \circ r_1$.

The full proof appears in Appendix A. □

While we have seen that appeals reduce the manipulability of the Immediate Acceptance mechanism, they do not make it strategy-proof. This is a potential concern because IA's efficiency under truthfulness disappears once we account for strategic behavior, so that in the no-appeals case the Nash equilibria of the corresponding preference revelation game become weakly Pareto-dominated by DA's truthful outcome (as shown by Ergin and Sönmez (2006))[7]. We will see shortly that this issue does not represent a concern in our setup.

Before showing that equilibrium play need not have an adverse welfare effect on IA with appeals, we note that appeals upheld through r_1 leave DA's strategy-proofness untouched (also under r_2, but not under r_3, as we show in Appendix C).

PROPOSITION 2. $\texttt{DA} \circ r_1$ *is strategy-proof.*

PROOF. Note that the DA outcome never admits a blocking pair under any reported preferences. Hence r_1 never upholds any appeal, and the composite mechanism coincides with DA, which is strategy-proof. □

4.3 *The Classical Ranking Can Reverse*

Proposition 1 shows that appeals reshape the incentive properties of standard mechanisms. We now show that appeals can also reverse the classical welfare comparison between DA and IA in the Nash equilibrium of their induced preference revelation games under complete information. Without appeals, a seminal result by Ergin and Sönmez (2006) shows that every Nash equilibrium of the IA game is weakly Pareto-dominated by the unique equilibrium in weakly dominant strategies of DA. But when we add appeals, not only does there exist a Nash equilibrium of the IA game that produces a weakly better outcome than when students report their preferences truthfully; in fact, there exists a Nash equilibrium that weakly Pareto-dominates (strictly so if DA is inefficient) the DA truthful outcome with or without appeals.

THEOREM 1. *For every school choice problem with strict priorities, there exists a Nash equilibrium* σ^* *of* $\texttt{IA} \circ r_1$ *such that*

$$(\texttt{IA} \circ r_1)(\sigma^*) \succsim \texttt{DA}(\succ).$$

Furthermore, if

$$(\texttt{IA} \circ r_1)(\sigma^*) \neq \texttt{DA}(\succ),$$

then at least one student is strictly better off under $(\texttt{IA} \circ r_1)(\sigma^*)$ *than under* $\texttt{DA}(\succ)$.

[7] Ortega (2026) quantifies IA's inefficiency loss in random markets when moving from truthful to strategic behavior.

PROOF. Our proof strategy is as follows: we construct an auxiliary one-shot mechanism, *Modified IA* (MIA), that runs like IA on truthful preferences but overbooks a rejected student whenever she outranks an admitted student, mimicking a successful strict appeal within the mechanism itself. We show that MIA weakly Pareto-dominates DA, using DA's stability, and then implement the MIA outcome as a Nash equilibrium of $\mathtt{IA} \circ r_1$: regularly admitted students rank their MIA school first, while overbooked students report truthfully and truncate at their MIA school, so that they are rejected in the first stage and recover it through an upheld appeal. The first stage of IA thus reproduces MIA's regular admissions and the appeal stage its overbooking, so the equilibrium outcome equals $\mathtt{MIA}(\succ) \gtrsim \mathtt{DA}(\succ)$.

Modified IA. The mechanism MIA proceeds in rounds, as IA does. In round k, every currently unassigned student applies to the k-th school on her true preference list. Admissions are permanent. At each school s, regular seats are filled only while fewer than q_s students have been regularly admitted to s. If s has remaining regular capacity, it admits the highest-priority current applicants up to the remaining regular capacity. Once q_s students have been regularly admitted to s, no further regular admissions occur. Any further current applicant i is admitted by overbooking if and only if $i \rhd_s j$ for some regularly admitted student j at s. In that case, s's effective capacity increases by one. Students admitted within the regular capacity are called *regularly admitted*; students admitted through the capacity-expansion rule are called *overbooked*. All other applicants are rejected and proceed to the next school on their list in the following round.

A useful consequence of this definition is that no student can be overbooked at a school she ranks first. If s is her first choice, she applies to s in round 1. If she is among the highest-priority applicants selected for the regular seats, she is regularly admitted. If she is not, then all regularly admitted students at s have higher priority than her, so she cannot trigger overbooking. Hence she is rejected. Let $\mathtt{MIA}(\succ)$ denote the outcome of this modified mechanism. We now present a useful property of MIA, which we call the Rejection Lemma.

LEMMA 2. *If student i is rejected from school s under $\mathtt{MIA}(\succ)$, then every regularly admitted student at s under $\mathtt{MIA}(\succ)$ has strictly higher priority at s than i.*

PROOF. Student i is rejected from s only if all regular seats at s have already been filled and i has lower priority than every regularly admitted student at s. Since regular admittees are never removed, and no further regular admittees can be added after the q_s regular seats are filled, every regularly admitted student at s under $\mathtt{MIA}(\succ)$ has strictly higher priority at s than i. □

By the definition of overbooking, if student i is overbooked at school s under $\mathtt{MIA}(\succ)$, then there exists a regularly admitted student j at s such that $i \rhd_s j$.

Having established the Rejection Lemma, we now show that MIA weakly Pareto-dominates DA.

LEMMA 3. *For every school choice problem with strict priorities,*

$$\texttt{MIA}(\succ) \gtrsim \texttt{DA}(\succ).$$

PROOF. Suppose, toward a contradiction, that some student is strictly better off under $\texttt{DA}(\succ)$ than under $\texttt{MIA}(\succ)$. Let

$$B = \{i \in I : \texttt{DA}_i(\succ) \succ_i \texttt{MIA}_i(\succ)\}.$$

By assumption, $B \neq \varnothing$.

For each $i \in B$, let $s_i = \texttt{DA}_i(\succ)$. Since $s_i \succ_i \texttt{MIA}_i(\succ)$, student i applied to s_i under $\texttt{MIA}$ before receiving her $\texttt{MIA}$-assignment, and was rejected. Let r_i be the round in which i applied to s_i under $\texttt{MIA}$. Let ρ_i be the round in which i receives her $\texttt{MIA}$-assignment, and set $\rho_i = \infty$ if i is unassigned under $\texttt{MIA}$. Then $r_i < \rho_i$.

By Lemma 2, every regularly admitted student at s_i under $\texttt{MIA}$ has strictly higher priority at s_i than i. Since i was rejected from s_i, the q_{s_i} regular seats at s_i had already been filled. Let $D(s_i)$ denote the set of these q_{s_i} regular admittees.

Since $i \notin D(s_i)$, $\texttt{DA}_i(\succ) = s_i$, and school s_i has only q_{s_i} seats under DA, at most $q_{s_i} - 1$ students from $D(s_i)$ can be assigned to s_i under $\texttt{DA}(\succ)$. Hence there exists some student $h \in D(s_i)$ who is not assigned to s_i under $\texttt{DA}(\succ)$.

Because $h \rhd_{s_i} i$ and i is assigned to s_i under $\texttt{DA}(\succ)$, stability of $\texttt{DA}(\succ)$ implies that $\texttt{DA}_h(\succ) \succ_h s_i$. Otherwise, (h, s_i) would be a blocking pair of $\texttt{DA}(\succ)$. But h's $\texttt{MIA}$-assignment is s_i, since h is regularly admitted to s_i under $\texttt{MIA}$. Therefore $\texttt{DA}_h(\succ) \succ_h \texttt{MIA}_h(\succ)$, so $h \in B$. Moreover, h was regularly admitted to s_i before or at the round in which i was rejected from s_i. Hence $\rho_h \leq r_i < \rho_i$.

This shows that from any $i \in B$, we can construct another student $h \in B$ with $\rho_h < \rho_i$. If all students in B had $\rho_i = \infty$, this argument would produce a student in B with finite ρ_h, a contradiction. Hence B contains at least one student with finite ρ_i. Choose $i^* \in B$ with minimal finite ρ_{i^*}. Applying the argument above to i^* produces $h \in B$ with finite $\rho_h < \rho_{i^*}$, contradicting the minimality of ρ_{i^*}. Therefore $B = \varnothing$, and $\texttt{MIA}(\succ) \gtrsim \texttt{DA}(\succ)$. □

Constructing the report profile. Having established that MIA weakly Pareto dominates DA, we proceed to construct the equilibrium strategies that deliver our result.

Partition students according to their status under $\texttt{MIA}(\succ)$:

- D: regularly admitted students;
- O: overbooked students;
- U: unassigned students.

Define the strategy profile σ^* as follows:

- If $i \in D$ and $\texttt{MIA}_i(\succ) = s$, then i ranks s first, followed by the remaining schools in her true order.

- If $i \in O$ and $\texttt{MIA}_i(\succ) = s$, then i reports truthfully down to s, truncating her list after s.
- If $i \in U$, then i reports truthfully.

The general construction permits truncation. The laboratory interface requires complete rank-order lists, but the particular experimental environment supports the relevant equilibrium with complete lists, as shown in Section 5.

Implementation. We show that

$$(\texttt{IA} \circ r_1)(\sigma^*) = \texttt{MIA}(\succ).$$

Fix a school s, and let D_s denote the set of students regularly admitted to s under MIA. Every student in D_s ranks s first under σ^*, so every student in D_s applies to s in round 1 of IA. We claim that the round-1 admits at s under $\texttt{IA}(\sigma^*)$ are exactly the students in D_s.

First suppose that $|D_s| < q_s$. Then no student outside D_s ranks s first under σ^*. To see this, suppose some student $i \notin D_s$ also ranks s first under σ^*. If $i \in D$, then i ranks her own MIA-school first, so she cannot rank s first unless she belongs to D_s, a contradiction. If $i \in O$ or $i \in U$, then s is also her first choice under $\succ_i$. Hence i applied to s in round 1 under MIA.

Since $|D_s| < q_s$, school s never filled its regular capacity under MIA. Hence a student applying to s in round 1 could not have been rejected from s, and would have been regularly admitted, again a contradiction.

Therefore the only round-1 applicants to s under σ^* are the students in D_s, and they are admitted.

Now suppose that $|D_s| = q_s$. Consider any student $i \notin D_s$ who also applies to s in round 1 under σ^*. There are three cases.

First, suppose $i \in O$ and $\texttt{MIA}_i(\succ) = s$. This is impossible. If s were i's first-ranked school under σ_i^*, then s would also be i's first choice under $\succ_i$. But no student can be overbooked at her first-choice school under MIA, as observed above.

Second, suppose $i \in O$, $\texttt{MIA}_i(\succ) = t \neq s$, and $s \succ_i t$. Then i was rejected from s under MIA. By Lemma 2, every student in D_s has strictly higher priority at s than i.

Third, suppose $i \in U$ and s is her first choice. Then i was rejected from s under MIA. Again, by Lemma 2, every student in D_s has strictly higher priority at s than i.

Thus every student in D_s has higher priority at s than every additional round-1 applicant to s under σ^*. Since $|D_s| = q_s$, the students in D_s are exactly the round-1 admits at s.

Now consider $i \in O$, and let $\texttt{MIA}_i(\succ) = s$. Under σ_i^*, student i applies to all schools she truly prefers to s, and then to s. At every school $t \succ_i s$, student i was rejected under MIA. By Lemma 2, every regular admittee at t has strictly higher priority than i. Since the round-1 IA admits at t are exactly the students in D_t, student i is rejected from every such t in the IA first stage.

At s, the regular seats are held by D_s, so i is rejected in the first stage. Her list is truncated at s, so she applies to no further school. Since i is overbooked at s under $\texttt{MIA}$, there exists $j \in D_s$ such that $i \rhd_s j$.

This student j occupies s in the IA first stage. Therefore (i,s) is a strict blocking pair of the first-stage outcome, and r_1 upholds i's appeal. Hence i obtains s, exactly as under $\texttt{MIA}$.

Finally, consider $i \in U$. Student i was rejected from every school under $\texttt{MIA}$. By Lemma 2, every regular admittee at every school t has strictly higher priority than i. Since the first-stage IA seats at t are filled by the students in D_t, student i is rejected everywhere in the first stage and has no successful strict appeal. Hence i remains unassigned, as under $\texttt{MIA}$.

We conclude that $(\texttt{IA} \circ r_1)(\sigma^*) = \texttt{MIA}(\succ)$, as desired.

σ^* is a Nash equilibrium. Fix a student i, and consider any school t such that $t \succ_i \texttt{MIA}_i(\succ)$. Student i was rejected from t under $\texttt{MIA}$. By Lemma 2, every regular admittee of t has strictly higher priority at t than i. Under σ^*_{-i}, these regular admittees rank t first and fill the first-stage IA seats at t, independently of i's report. Since IA admissions are permanent, i cannot obtain t directly, regardless of when she applies to t. Nor can i obtain t through a strict appeal. The first-stage occupants of t are precisely the students in D_t, all of whom have strictly higher priority at t than i. Thus (i,t) is not a strict blocking pair, and i's appeal to t fails. Therefore no deviation allows i to obtain a school she strictly prefers to $\texttt{MIA}_i(\succ)$. Since i was arbitrary, σ^* is a Nash equilibrium of $\texttt{IA} \circ r_1$.

Combining the implementation result with Lemma 3, we obtain

$$(\texttt{IA} \circ r_1)(\sigma^*) = \texttt{MIA}(\succ) \succsim \texttt{DA}(\succ).$$

If $(\texttt{IA} \circ r_1)(\sigma^*) \neq \texttt{DA}(\succ)$, then Pareto dominance together with strict preferences implies that at least one student is strictly better off. This proves the theorem. □

Theorem 1 is an equilibrium-existence result: it shows that once strict appeals are available, IA admits an equilibrium whose outcome weakly Pareto-dominates the truthful DA outcome. It does not say that IA with appeals is always better than DA, nor that every equilibrium of IA with appeals has this property. This distinction is important. The point is not that IA becomes uniformly superior once appeals are introduced. Rather, the theorem shows that appeals change the comparison between mechanisms in a fundamental way. In the standard one-stage model, IA is strategically fragile relative to DA's efficiency. With appeals, however, some of the risk that generates "defensive" ranking under IA is shifted to the second stage. A student who is rejected from a preferred school may still obtain it if the rejection creates a strict priority claim. Thus appeals act as partial insurance against one of IA's central strategic costs.

Together, Proposition 1 and Theorem 1 suggest that appeals are not a small institutional detail added after the mechanism has done its work. They change both the feasible final outcomes and the strategic environment in which families submit rankings.

Consequently, a model that compares DA and IA without appeals may miss an important part of the institutional trade-off.

In Appendix B, we extend the reversal result to the probabilistic rule r_2. Let $\mu = \texttt{DA}(\succ)$, and let each student report her true ranking truncated immediately after her assignment under μ. For every school-choice problem with strict priorities and every $p \in (0,1)$, this DA-cutoff profile is an sd-Nash equilibrium of $\texttt{IA} \circ r_2$, and its induced lottery sd-Pareto-dominates the degenerate truthful DA outcome. The equilibrium notion is ordinal: no unilateral deviation yields a stochastically dominant lottery. Probabilistic appeals therefore preserve the reversal in stochastic-dominance terms while turning otherwise unsuccessful applications into lottery tickets.

4.4 *Hypotheses*

The theoretical analysis yields predictions concerning reporting behavior, student welfare, and the number of blocking pairs.

HYPOTHESIS 1. [Truth-telling (Propositions 1 and 2)].

(1.A) Truth-telling is more frequent under $\texttt{IA} \circ r_1$ *than under* $\texttt{IA} \circ r_0$*, and more frequent under* $\texttt{IA} \circ r_2$ *than under* $\texttt{IA} \circ r_0$*.*

(1.B) The effect of strict and probabilistic appeals on truth-telling is smaller under DA than under IA.

HYPOTHESIS 2. [Efficiency (Theorem 1)]. *Strict appeals improve student welfare more under IA than under DA. The same directional comparison holds in stochastic-dominance terms under the probabilistic rule.*

The welfare prediction is based on the equilibrium constructed in Theorem 1. It is therefore an equilibrium-selection prediction rather than a claim that every equilibrium of IA with appeals dominates the truthful DA outcome.

HYPOTHESIS 3. [Blocking pairs (Lemma 1)]. *Strict appeals reduce the incidence of strict blocking pairs under IA. They do not improve this outcome under DA, which already produces an assignment without strict blocking pairs.*

We do not formulate an analogous hypothesis for probabilistic appeals. Unlike the strict rule, the probabilistic rule does not guarantee the elimination of blocking pairs.

5. EXPERIMENT

To test our hypotheses, we conduct a laboratory experiment in the spirit of the seminal school-choice experiment of Chen and Sönmez (2006).[8] Participants take part in a simulated school-choice market consisting of seven students and three schools, each with two seats. Preferences and priorities are as follows:

[8] For a comprehensive review of the experimental literature on school choice, see Hakimov and Kübler (2021).

$\succ_{i_1}$	$\succ_{i_2}$	$\succ_{i_3}$	$\succ_{i_4}$	$\succ_{i_5}$	$\succ_{i_6}$	$\succ_{i_7}$	$\trianglerighteq_{s_1}$	$\trianglerighteq_{s_2}$	$\trianglerighteq_{s_3}$
s_3	s_3	s_2	s_1	s_1	s_2	s_3	i_6, i_7	i_4	i_4, i_5
s_2	s_2	s_1	s_3	s_2	s_3	s_2	i_5	i_1, i_2, i_3, i_5	i_3, i_6
s_1	s_1	s_3	s_2	s_3	s_1	s_1	i_2, i_3, i_4	i_6, i_7	i_2, i_7
$\cdot$	$\cdot$	$\cdot$	$\cdot$	$\cdot$	$\cdot$	$\cdot$	i_1	$\cdot$	i_1

Comma-separated students have equal underlying priority, with ties broken in favor of the lower-indexed student. Participants receive 15, 10, or 5 points when assigned to their first, second, or third choice, respectively, and 0 points when unmatched.

We use a 2×3 factorial between-subjects design. The first factor varies the first-stage assignment mechanism between immediate acceptance (IA) and deferred acceptance (DA). The second factor varies the appeal rule between no appeals (r_0), strict appeals (r_1), and probabilistic appeals (r_2). Under r_1, an appeal is upheld only when the appellant has a strict priority claim. Under r_2, strict priority claims are upheld with certainty, while all other appeals are upheld independently with probability $p = 0.1$. The six treatments are therefore

$$\begin{array}{lll} \mathtt{DA} \circ r_0, & \mathtt{DA} \circ r_1, & \mathtt{DA} \circ r_2, \\ \mathtt{IA} \circ r_0, & \mathtt{IA} \circ r_1, & \mathtt{IA} \circ r_2. \end{array}$$

Participants are randomly assigned to one of the six treatments for the duration of the experiment. The experiment proceeds for ten periods. In each period, sessions of either 14 or 28 participants were partitioned into seven-person markets. Within each market, the seven student types were assigned one-to-one at random. Participants were rematched across periods within their session. Because rematching occurred across the entire session, each session constitutes one independent observation (matching group), and all standard errors are clustered at the matching group level. We collected between 9 and 12 independent observations per treatment and 59 independent sessions in total.

Participants have complete information about the preferences of all student types and the priorities of all schools. They do not, however, observe the rank-order lists submitted by the other participants. Each participant submits a complete rank-order list over the three schools; truncation and ranking the outside option are not permitted. Under truthful reporting, the relevant first- and second-stage assignments are

$$\begin{aligned} \mathtt{DA}(\succ) = (\mathtt{DA} \circ r_1)(\succ) &= (s_1 : \{i_6, i_7\};\ s_2 : \{i_1, i_2\};\ s_3 : \{i_4, i_5\};\ s_\emptyset : \{i_3\}), \\ \mathtt{IA}(\succ) &= (s_1 : \{i_4, i_5\};\ s_2 : \{i_3, i_6\};\ s_3 : \{i_2, i_7\};\ s_\emptyset : \{i_1\}), \\ (\mathtt{IA} \circ r_1)(\succ) &= (s_1 : \{i_4, i_5\};\ s_2 : \{i_1, i_3, i_6\};\ s_3 : \{i_2, i_7\}). \end{aligned}$$

Thus, strict appeals leave the truthful DA assignment unchanged in this market, whereas under IA they allow i_1 to obtain a seat at s_2. For this preference profile, truthful reporting is a Nash equilibrium of the preference revelation game induced by $\mathtt{IA} \circ r_1$. Students $i_2, \ldots, i_7$ obtain their first-choice schools, while i_1 obtains her second choice, s_2, through a successful strict appeal. The only school that i_1 strictly prefers to s_2 is s_3, whose two seats are filled in the first round by i_2 and i_7, both of whom have strictly higher priority

than i_1. No report can therefore give i_1 a strictly better final assignment. This equilibrium is supported by complete rank-order lists and hence does not rely on truncation.

In the appeal treatments, participants observe the first-stage assignment and may appeal to any school that they ranked above their assignment and from which they were rejected. Filing an appeal carries no monetary fee but requires completing a short typing task. This small non-monetary friction is identical under the strict and probabilistic appeal rules. It makes each appeal an active and observable choice without changing either the substantive grounds on which it is evaluated or its probability of success. This feature deliberately differs from the theoretical benchmark. In the experiment, participants instead choose whether to appeal and, when several appeal opportunities are available, which claims to pursue. Appeal take-up is therefore an additional behavioral outcome of the experiment.

Because participants are rematched across the entire 14- or 28-person session, each session (matching group) constitutes one independent observation. We cluster all standard errors at the session level.

Experimental earnings are denominated in points and converted at the rate 1 point $= £0.05$. At the end of the experiment, two of the ten periods are randomly selected for payment, and participants receive the sum of their earnings in those periods together with a £5 show-up fee. Before entering the simulated school-choice market, participants must correctly answer a set of comprehension questions. After the ten periods, they complete a short incentivized risk-elicitation task, a questionnaire concerning their appeal decisions in the appeal treatments, and a sociodemographic survey.

We collected between 9 and 12 independent observations per treatment, i.e., 168 participants in each treatment, except for $\texttt{IA} \circ r_2$, which included 182 participants. The experiment therefore involved 1,022 participants across 59 independent matching groups. Following the session records, we exclude 25 observations recorded after six participants had exited their sessions. The resulting analysis sample contains $10{,}195$ participant-period observations.

The experiment received ethics approval from the institutional review board at City St George's, University of London, and was preregistered on the Open Science Framework (https://osf.io/q92de/overview). The experiment was conducted at EssexLab at the University of Essex.

6. Results

We begin by presenting results on individual-level outcomes (truth-telling and appeals) and then discuss market-level outcomes (efficiency and stability). We conclude with a discussion of the trade-offs observed in all treatments.

6.1 *Truth-telling*

We define truth-telling as exactly submitting the induced preference ranking. Figure 1 reveals a clear asymmetry in how appeals affect truth-telling under the two mechanisms. Without appeals, truth-telling is nearly identical under DA and IA, at 0.29 and 0.28, respectively. This baseline pattern is striking from a theory point of view, given DA's

strategy-proofness, but it is in line with other laboratory experiments where substantial manipulation has been observed under DA, both among student subjects and parents (Freer et al., 2026, Cerrone et al., 2024). Appeals substantially increase truth-telling under IA. The percentage of truthful reports rises from 28 without appeals to 39 under both strict and probabilistic appeals. The corresponding increases under DA are considerably smaller, from 29 without appeals to 32 under strict appeals and 33 under probabilistic appeals.

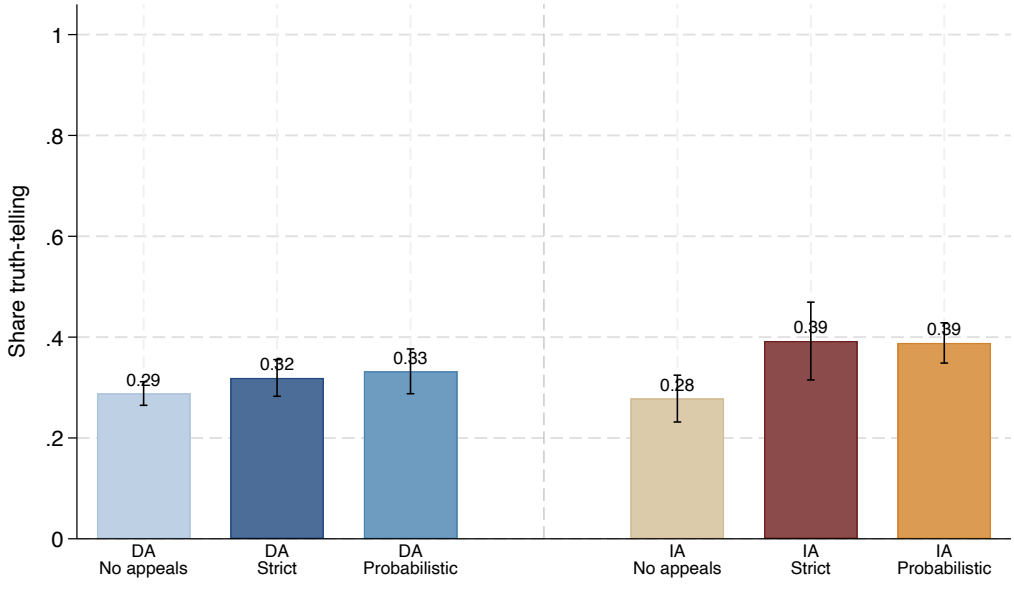


FIGURE 1. Truth-telling by treatment.
Bars show the share of participants submitting the complete induced preference ranking. Whiskers are 95% confidence intervals clustered at the matching-group level.

To test these treatment differences more formally and account for changes over repeated play, including learning, Table 2 reports average marginal effects from logit regressions with progressively richer controls. Specification (1) includes treatment indicators only. Specification (2) adds round effects. Specification (3) adds participant-type effects, and Specification (4) adds age, gender, and risk aversion. Standard errors are clustered at the matching-group level.

Under IA, strict and probabilistic appeals increase truth-telling by approximately 11 percentage points – an effect that is highly significant and remains nearly unchanged across specifications. Under DA, the estimated increases are only 3 to 4 percentage points and are at most marginally significant.

Panel C directly compares the appeal effects across mechanisms. The effect of strict appeals is 8 to 9 percentage points larger under IA than under DA, while the effect of

TABLE 2. Impact of appeals on truth-telling.

	Specifications			
	(1)	(2)	(3)	(4)
Panel A. IA treatments				
IA strict	0.11***	0.11***	0.11***	0.12***
	(0.04)	(0.04)	(0.04)	(0.04)
IA probabilistic	0.11***	0.11***	0.11***	0.11***
	(0.03)	(0.03)	(0.03)	(0.03)
Panel B. DA treatments				
DA strict	0.03	0.03	0.03	0.03*
	(0.02)	(0.02)	(0.02)	(0.02)
DA probabilistic	0.04*	0.04*	0.04*	0.04*
	(0.02)	(0.02)	(0.02)	(0.02)
Panel C. Difference in appeal effects between IA and DA				
Strict appeals	0.08*	0.08*	0.08*	0.09*
	(0.05)	(0.05)	(0.05)	(0.05)
Probabilistic appeals	0.07*	0.07*	0.07*	0.07*
	(0.04)	(0.04)	(0.04)	(0.04)
Round	No	Yes	Yes	Yes
Participant type	No	No	Yes	Yes
Demographics + risk att.	No	No	No	Yes
Observations	10,195	10,195	10,195	10,185
Matching groups	59	59	59	59

*** $p < 0.01$; ** $p < 0.05$; * $p < 0.10$. Coefficients are reported as average marginal effects from logit regressions. Standard errors are clustered at the matching-group level. Truth-telling equals one when the participant submits the complete induced preference ranking. The reference categories in Panels A and B are the corresponding mechanisms without appeals. Panel C reports (IA appeal − IA no appeals) − (DA appeal − DA no appeals), so positive values indicate that appeals increase truth-telling more under IA. Each observation is a participant-round. Specification (4) excludes ten participant-rounds with missing demographic or risk-attitude data. Matching groups are experimental sessions. Demographic controls are gender and age. Risk att. refers to risk attitudes.

probabilistic appeals is approximately 7 percentage points larger. Both differences are marginally significant across specifications.

This asymmetry is consistent with the mechanism identified by the theory. Under IA, appeals reduce the risk associated with ranking an ambitious school first because an unjustified rejection may be corrected afterwards. Under DA, truthful reporting is already a dominant strategy, leaving less scope for appeal rights to change initial reports.

RESULT 1. *The data strongly support Hypothesis 1.A and provide partial support for Hypothesis 1.B. Under IA, both appeal rules increase truth-telling by approximately* 11 *percentage points. Under DA, the effects are only* 3 *to* 4 *percentage points and are at most marginally significant.*

As a robustness check, we perform the same analysis when only the first reported school was the same as the induced first preference. The same qualitative conclusions emerge (see Appendix D).

Consequential non-truthful reports. Exact truth-telling treats every departure from the induced ranking alike, even when it has no effect on the participant's assignment. We therefore replace each non-truthful report with the participant's induced ranking, hold all other reports fixed, and rerun the first-stage mechanism. We classify a report as inconsequential if the assignment is unchanged, beneficial if the observed report yields more points, and harmful if it yields fewer points.[9]

The contrast between the mechanisms is prononuced. Under DA, only about 9 percent of non-truthful reports are consequential without appeals, 8 percent under strict appeals, and 7 percent under probabilistic appeals. Every consequential DA report is harmful. Low truth-telling under DA therefore largely reflects reports that do not affect assignments. Under IA, by contrast, almost two thirds of non-truthful reports are consequential in every treatment. Yet participants exploit IA's manipulability poorly: among consequential reports, about 40 percent are beneficial and 60 percent are harmful.

Appeals mainly affect whether participants misreport at all. They increase truth-telling under IA, but do not make the remaining non-truthful reporters more successful at manipulating the mechanism.

6.2 *Appeal take-up*

We next examine when participants make use of the appeal stage and whether they direct appeals toward claims that can improve their payoff.

Appeal take-up is high in all four treatments. Participants nevertheless file more appeals under DA than under IA, averaging between 6 and 7 appeals over ten rounds under DA, compared with about 4 under IA. This does not reflect a greater willingness to appeal under DA. Rather, DA produces worse first-stage assignments and therefore almost twice as many school-specific appeal opportunities, approximately 11 per participant compared with 6 under IA.

When having the opportunity to file an appeal, participants are actually more likely to appeal under IA. Under strict appeals, at least one appeal is filed in 67 percent of eligible participant-rounds under DA and 82 percent under IA. Under probabilistic appeals, the corresponding rates are 75 percent and 80 percent. Thus, DA generates more appeals overall only because it creates more opportunities to appeal.

[9] We evaluate the counterfactual at the first stage because a different report generally changes the set of available appeals. A final-outcome comparison would require assumptions about counterfactual appeal choices and, under the probabilistic rule, unobserved appeal draws.

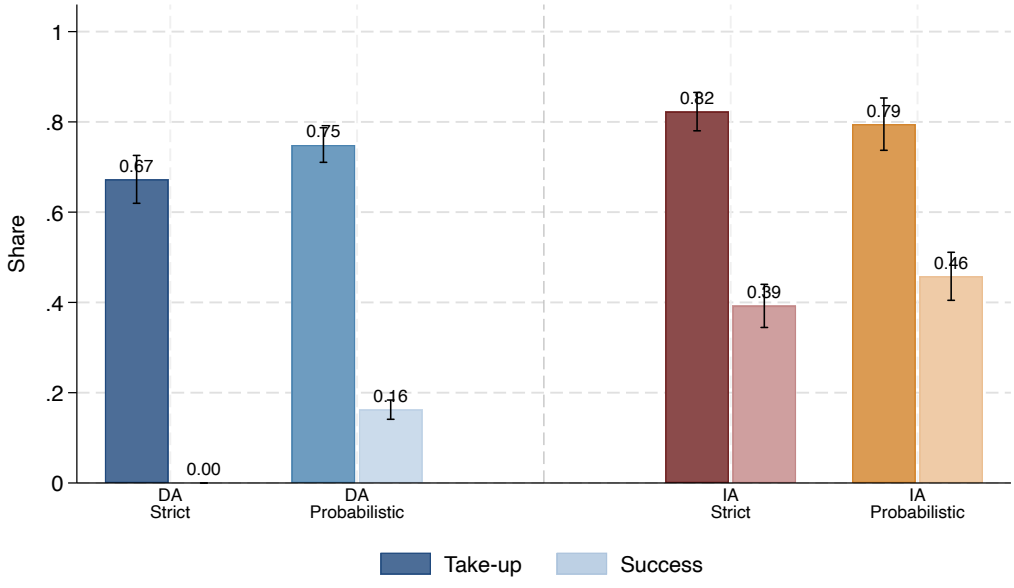


FIGURE 2. Appeal take-up and success by treatment.
Dark bars show the share of eligible participant-rounds in which at least one appeal is filed. Light bars show the share of appealing participant-rounds in which at least one appeal succeeds. Whiskers are 95% confidence intervals clustered at the matching-group level.

The central difference between DA and IA is whether those appeals eventually work. Under strict appeals, no appeal succeeds under DA because it creates no strict priority violation, whereas at least one appeal succeeds in 39 percent of appealing participant-rounds under IA. Under probabilistic appeals, the corresponding success rates are 16 percent under DA and 46 percent under IA. Appeals under DA therefore rely entirely on the random channel, while IA frequently creates priority claims that the appeal stage can correct. Overall, IA produces fewer appeals, but substantially more effective ones.[10]

These aggregate rates do not reveal whether participants select carefully among the appeal opportunities available to them. Table 3 therefore examines which features of an appeal predict whether it is lodged. Column (1) is estimated at the participant-round level and considers whether at least one appeal is filed. Column (2) is estimated at the school-specific opportunity level and considers whether a particular rejection is appealed.

The first-stage outcome is the strongest predictor of appeal use. Relative to receiving the true first-choice school, receiving the second choice raises the probability of filing an appeal by 0.31, receiving the third choice raises it by 0.49, and remaining unassigned

[10]Success is measured at the participant-round level and equals one when at least one filed appeal succeeds. Because a participant may file several probabilistic appeals, each with success probability 0.1, the probability that at least one succeeds can exceed 10 percent.

TABLE 3. Determinants of appeal use.

	Any appeal	Appeal opportunity taken
DA, probabilistic appeals	0.08**	0.05
	(0.03)	(0.03)
IA, strict appeals	0.11***	0.03
	(0.04)	(0.04)
IA, probabilistic appeals	0.08**	0.05
	(0.04)	(0.04)
First-stage assignment: second choice	0.31***	0.23**
	(0.09)	(0.09)
First-stage assignment: third choice	0.49***	0.40***
	(0.09)	(0.08)
First-stage assignment: unassigned	0.51***	0.44***
	(0.10)	(0.08)
Strict priority claim	−0.01	0.07***
	(0.02)	(0.03)
Number of appeal opportunities	0.11***	·
	(0.03)	
Appeal target: second choice	·	−0.05***
		(0.01)
Appeal target: third choice	·	−0.10***
		(0.02)
Round fixed effects	Yes	Yes
Observations	2,912	5,725
Matching groups	41	41
R^2	0.16	0.03

*** $p < 0.01$; ** $p < 0.05$; * $p < 0.10$. Coefficients are reported from linear probability models. Standard errors are clustered at the matching-group level. The reference category is DA with strict appeals. Column (1) contains one observation for each participant-round with at least one appeal opportunity. Column (2) contains one observation for each school-specific appeal opportunity. Only the four appeal treatments are included, so the sample contains 41 matching groups. Matching groups are experimental sessions. First-stage assignments and appeal targets are ranked using the induced true preferences. In Column (1), strict priority claim indicates that at least one available appeal is supported by a strict priority violation.

raises it by 0.51. Each additional appeal opportunity raises the probability of filing at least one appeal by 0.11. Participants therefore respond strongly to how poorly they fare in the first stage.

They also respond to the quality of individual claims. Conditional on being offered a particular appeal opportunity, a strict priority claim is 7 percentage points more likely to be pursued. Appeals to the true second- and third-choice schools are respectively 5 and 10 percentage points less likely to be lodged than appeals to the true first choice. Once

these features are held fixed, the treatment indicators have little effect at the school-specific opportunity level. The aggregate differences across treatments therefore arise principally from the kinds of appeal opportunities the mechanisms generate.

Appeal use is nevertheless imperfect. Under DA with strict appeals, none of the $1{,}845$ appeal opportunities can succeed, yet participants pursue $1{,}083$ of them. Under IA with strict appeals, participants pursue 166 of 233 sure, payoff-improving claims, but also lodge 465 appeals with no chance of success. Under IA with probabilistic appeals, they pursue 178 of 256 sure gains and 519 of 785 positive-value lottery tickets. Participants are more likely to pursue probabilistic appeals as their expected payoff increases, but leave approximately three in ten sure gains unclaimed.[11]

6.3 *Efficiency*

We now test the paper's central welfare prediction. As Figure 3 shows, strict appeals improve IA while leaving DA essentially unchanged. Our primary outcome is the rank of the final assignment in the participant's induced true preferences; rank 1 denotes the most preferred school, and unassigned participants are coded as rank 4. We also report welfare in experimental points in Appendix D.

IA delivers better average ranks than DA in every treatment. This is not a reversal in the sign of the observed ranking, because IA already outperforms DA without appeals. The relevant comparative static is that strict appeals substantially widen IA's advantage. The average-rank gap between DA and IA rises from 0.41 without appeals to 0.68 with strict appeals, an increase of 0.27 rank positions (matching-group-clustered standard error 0.07; $p < 0.001$).

Appeals also reduce the probability of remaining unassigned under IA. The percentage of unassigned subjects falls from 14 without appeals to 7 with strict appeals and 6 with probabilistic appeals. Under DA, the corresponding rate remains 14 with strict appeals and falls only to 12 with probabilistic appeals. Part of the improvement in mean rank under IA therefore comes from assigning students who would otherwise remain unmatched.

The level comparison differs from the no-appeals benchmark embedded in Hypothesis 2. IA already produces better assignments than DA without appeals, with mean ranks of 2.33 and 2.74, respectively. Appeals therefore do not reverse the observed ranking between DA and IA. Instead, they widen an IA advantage that is already present in

[11] Appendix E imposes the full-take-up appeal strategy assumed by the model while holding submitted rank-order lists and first-stage assignments fixed. The results in the main text are preserved and sharpened. Under strict appeals, IA's average assigned rank improves from 2.06 to 2.01, while DA remains unchanged at 2.74; the DA–IA rank gap therefore widens from 0.68 to 0.73. Under probabilistic appeals, average rank improves from 2.04 to 1.94 under IA and from 2.61 to 2.55 under DA, so IA's rank advantage again widens. Full take-up also raises the share of IA market-rounds that strictly Pareto-dominate truthful DA from 0.48 to 0.60 under strict appeals and from 0.46 to 0.63 under probabilistic appeals. These additional efficiency gains are accompanied by substantially greater procedural stability, which rises to 0.94 under strict appeals and 0.84 under probabilistic appeals. Under the probabilistic rule, full take-up therefore makes IA both more rank-efficient and more procedurally stable than DA.

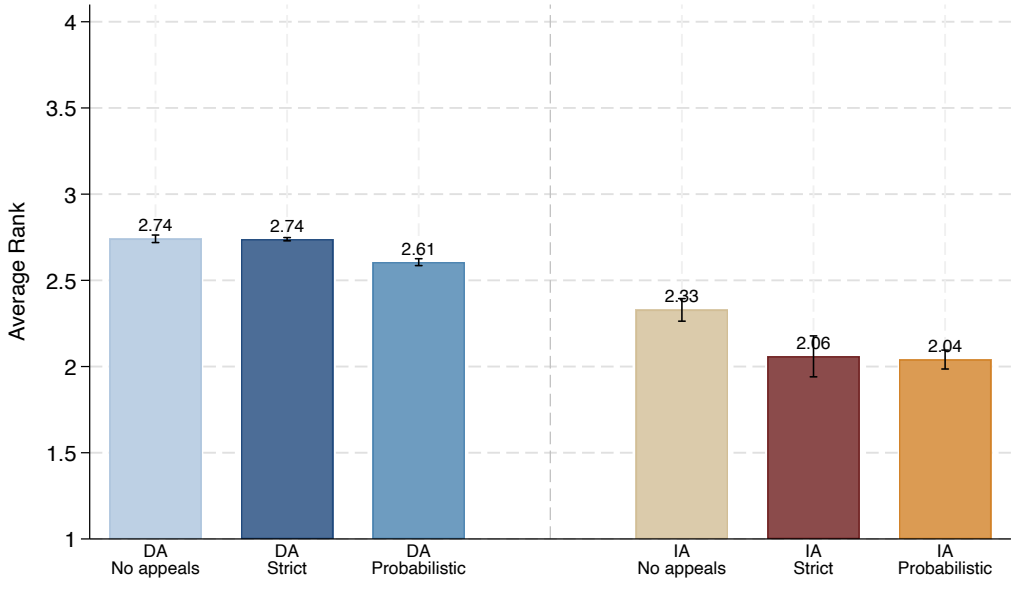


FIGURE 3. Average assigned rank by treatment.
Assignments are evaluated against induced true preferences. Unassigned participants are coded as rank 4, and lower values indicate better assignments. Whiskers are 95% confidence intervals clustered at the matching-group level.

the experiment. This efficiency-improving effect is substantially stronger under strict and probabilistic appeals, because appeals recover part of the inefficiency obtained under IA without appeals while leaving DA broadly unchanged.

Table 4 examines the robustness of these comprisons and reports adjusted differences in expected assigned rank from ordered logit models. The specifications follow those in Table 2. Specification (1) includes treatment indicators only. Specification (2) adds round effects. Specification (3) adds participant types. Specification (4) adds demographics and risk attitudes. Standard errors are clustered at the matching-group level. The regression estimates closely match the descriptive differences. Under IA, strict and probabilistic appeals improve expected rank by approximately 0.27 and 0.29 rank points, respectively. Under DA, strict appeals have no detectable effect, while probabilistic appeals improve expected rank by approximately 0.11 to 0.12 rank points. The estimates are stable across specifications.

Panel C shows that appeals significantly widen the rank gap between DA and IA. In specification (4), which includes most controls, strict appeals increase the gap by 0.27 rank points, while probabilistic appeals increase it by 0.18 rank points. The strict rule produces the sharper contrast because it improves IA without generating the random gains that probabilistic appeals also create under DA.

TABLE 4. Impact of appeals on assigned rank.

	Specifications			
	(1)	(2)	(3)	(4)
Panel A. IA treatments				
IA strict	-0.27^{***}	-0.27^{***}	-0.27^{***}	-0.27^{***}
	(0.07)	(0.07)	(0.07)	(0.07)
IA probabilistic	-0.29^{***}	-0.29^{***}	-0.29^{***}	-0.29^{***}
	(0.05)	(0.05)	(0.04)	(0.04)
Panel B. DA treatments				
DA strict	0.00	0.00	0.00	0.00
	(0.01)	(0.01)	(0.01)	(0.01)
DA probabilistic	-0.12^{***}	-0.12^{***}	-0.11^{***}	-0.11^{***}
	(0.01)	(0.01)	(0.01)	(0.01)
Panel C. Change in rank gap between DA and IA relative to no appeals				
Strict appeals	0.26^{***}	0.26^{***}	0.27^{***}	0.27^{***}
	(0.07)	(0.07)	(0.07)	(0.07)
Probabilistic appeals	0.16^{***}	0.16^{***}	0.18^{***}	0.18^{***}
	(0.05)	(0.05)	(0.05)	(0.05)
Round effects	No	Yes	Yes	Yes
Participant-type effects	No	No	Yes	Yes
Demographics + risk att.	No	No	No	Yes
Observations	10,195	10,195	10,195	10,185
Matching groups	59	59	59	59

*** $p < 0.01$; ** $p < 0.05$; * $p < 0.10$. Coefficients are reported as adjusted differences in expected assigned rank from ordered logit models. Lower values indicate better assignments, and unassigned participants are coded as rank 4. The reference categories in Panels A and B are the corresponding mechanisms without appeals. Panel C reports (DA appeal − IA appeal) − (DA no appeals − IA no appeals), so positive values indicate that appeals widen IA's rank advantage. Standard errors are clustered at the matching-group level. Each observation is a participant-round. Specification (4) excludes ten participant-rounds with missing demographic or risk-attitude data. Matching groups are experimental sessions. Demographic controls are gender and age. Risk att. refers to risk attitudes.

Pareto dominance relative to truthful DA. Theorem 1 is an equilibrium-existence result: strict appeals allow IA to support an equilibrium outcome that weakly Pareto-dominates truthful DA, and strictly Pareto-dominates it whenever the two outcomes differ. Observed laboratory play need not coincide with the equilibrium constructed in the proof. We therefore use a direct allocation-level counterpart of the theorem. For every complete seven-person market-round, we compare the realized final assignment with the truthful DA benchmark. We say that the realized assignment weakly Pareto-dominates this benchmark if every participant is weakly better off according to her induced true preferences, and strictly Pareto-dominates it if, in addition, at least one participant is strictly better off.

Without appeals, observed IA strictly Pareto-dominates truthful DA in 27 market-rounds (11 percent). With strict appeals, this frequency rises to 112 of 235 (48 percent). Strict appeals therefore increase the frequency of strict Pareto dominance by 36 percentage points (the one-point discrepancy is due to rounding; standard error 4.6 percentage points; $p < 0.001$, with standard errors clustered at the session level. This is a close empirical counterpart of Theorem 1. The theorem identifies a change in the outcomes that IA can sustain in equilibrium; the experiment shows that strict appeals move realized IA assignments sharply toward the Pareto region above truthful DA.

When do the rank improvements emerge? Appeal rights can affect final assignments at two points. Anticipating the appeal stage may change the rank-order lists participants submit and therefore the assignment produced by the first-stage mechanism. Filed and upheld appeals may then generate a further improvement after the first-stage assignment has been announced. We therefore ask how much of the final rank difference is already present before appeals are resolved and how much accrues subsequently.

For a first-stage mechanism $M \in \{\texttt{DA}, \texttt{IA}\}$ and an appeal rule $r \in \{r_1, r_2\}$, let $\bar{R}^{\text{first}}_{M,r}$ and $\bar{R}^{\text{final}}_{M,r}$ denote average assigned rank after the first-stage and final assignments, respectively. Both are evaluated against induced true preferences, with unassigned participants coded as rank 4. Lower values therefore indicate better assignments. Under no appeals, $\bar{R}^{\text{final}}_{M,r_0} = \bar{R}^{\text{first}}_{M,r_0}$, so

$$\underbrace{\bar{R}^{\text{final}}_{M,r_0} - \bar{R}^{\text{final}}_{M,r}}_{\text{total rank improvement}} = \underbrace{\bar{R}^{\text{first}}_{M,r_0} - \bar{R}^{\text{first}}_{M,r}}_{\text{first-stage difference}} + \underbrace{\bar{R}^{\text{first}}_{M,r} - \bar{R}^{\text{final}}_{M,r}}_{\text{subsequent appeal gain}} .$$

This identity is a chronological accounting decomposition rather than a causal mediation analysis. The first-stage difference measures how much of the treatment difference is already present before any appeal is resolved. In our design, this difference can operate only through changes in submitted rank-order lists and the assignments they generate. The subsequent appeal gain measures the improvement from the first-stage to the final assignment among participants assigned to the appeal treatment. It is evaluated at the reports those participants submitted when appeals were available and is non-negative by construction. The decomposition therefore identifies when the observed rank improvement emerges, but it does not provide a unique causal allocation between reporting behavior and the appeal stage.

Table 5 shows that the rank improvement under IA emerges both before and after appeals are resolved. Strict appeals improve final assigned rank by 0.270 positions. Of this improvement, 0.163 positions are already present in the first-stage assignment, while the appeal stage generates a further gain of 0.107 positions. Thus, approximately 60 percent of the observed improvement is visible before any appeal is resolved.

Probabilistic appeals produce a similar pattern under IA. They improve final assigned rank by 0.288 positions, of which 0.167 positions are present at the first stage and 0.121 accrue subsequently. Under both appeal rules, more than half of IA's final rank improvement is therefore already visible in the first-stage assignment.

TABLE 5. Decomposition of improvements in average assigned rank.

Treatment	Total rank change	First-stage change	Subsequent appeal change
DA, strict appeals	0.00	0.00	0.00
	(0.01)	(0.01)	(0.00)
DA, probabilistic appeals	0.14***	0.02	0.12***
	(0.01)	(0.01)	(0.01)
IA, strict appeals	0.27***	0.16**	0.11***
	(0.07)	(0.06)	(0.01)
IA, probabilistic appeals	0.29***	0.17***	0.12***
	(0.04)	(0.04)	(0.01)

Positive entries indicate reductions in average assigned rank. Each observation is a participant-round, and matching groups are experimental sessions. Ranks are evaluated against induced true preferences, with unassigned participants coded as rank 4. The total improvement compares final rank under the indicated appeal rule with final rank under the same first-stage mechanism without appeals. The first-stage difference makes the corresponding comparison before appeals are resolved. The subsequent appeal gain is the within-treatment reduction in average rank between the first-stage and final assignments. The three point estimates in each row satisfy the accounting identity exactly. Standard errors, clustered at the matching-group level, are in parentheses. *** $p < 0.01$, ** $p < 0.05$, * $p < 0.10$.

The contrast with DA is pronounced. Strict appeals leave both first-stage and final rank essentially unchanged because DA produces no strict priority claim for the appeal stage to correct. Probabilistic appeals improve final rank by 0.136 positions, but almost all of this improvement, 0.117 positions, arises after DA has run. The first-stage difference of 0.019 positions is small and statistically insignificant.

RESULT 2. *The data strongly support Hypothesis 2. Strict appeals improve IA substantially more than DA: they improve mean assigned rank under IA by* 0.27 *positions while leaving DA unchanged, thereby widening IA's rank advantage by* 0.27 *positions. They also raise the share of realized IA outcomes that strictly Pareto-dominate truthful DA from* 11 *to* 48 *percent. Approximately* 60 *percent of the observed IA rank improvement is already present before appeals are resolved. Probabilistic appeals produce the same directional pattern.*

6.4 *Stability*

We finally examine whether IA's efficiency gains come at the cost of stability. The strict appeal rule is defined using the underlying priority classes, not the tie-break used by the first-stage mechanism. Accordingly, a strict blocking pair requires the student to have strictly higher underlying priority than an assigned student. Equal-priority students who lose the tie-break do not form strict blocking pairs. We treat wastefulness separately. A market-round is stable if it contains neither a strict blocking pair nor a wasted-seat claim.

Our main measure evaluates stability against submitted preferences. This is the natural procedural benchmark because both appeal eligibility and the panel's decision are

based on the rank-order list submitted in the first stage. We also report stability against induced true preferences. The latter is a more demanding benchmark because a school that a participant places below her assignment cannot subsequently be claimed on appeal.

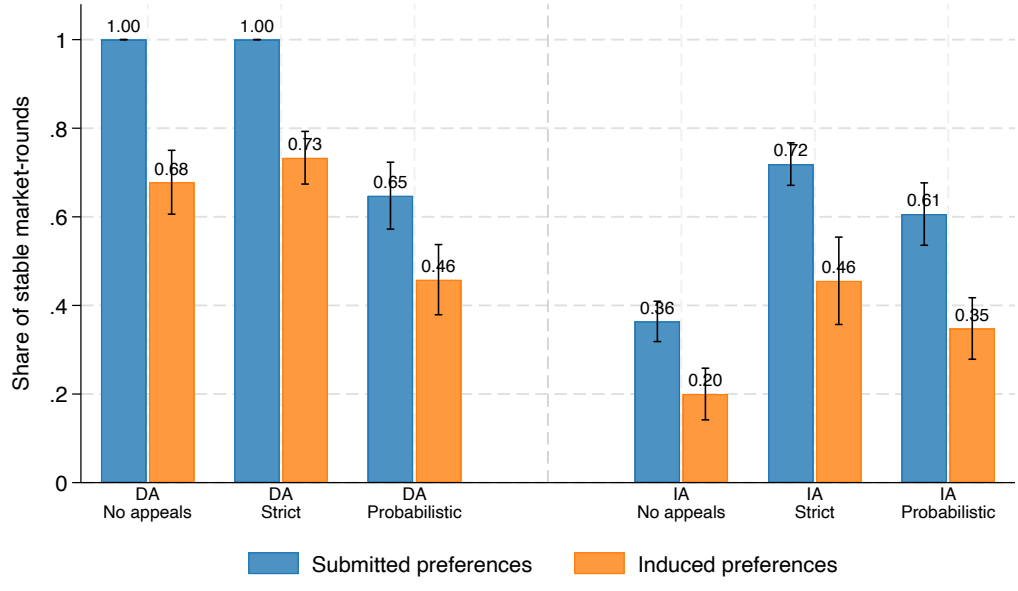


FIGURE 4. Stability by treatment.

Bars report the share of stable market-rounds, evaluated against submitted and induced true preferences. A market-round is stable if it contains neither a strict blocking pair nor a wasted-seat claim. Equal-priority tie-break claims are not counted as strict blocking pairs. Whiskers are 95% confidence intervals clustered at the matching-group level.

Figure 4 shows that strict appeals substantially improve IA's procedural stability. Without appeals, IA produces 0.9 strict blocking pairs per market-round, and the stable share is only 0.4. Strict appeals reduce strict blocking pairs to 0.3, generate only 0.02 wasted-seat claims per market-round, and raise the stable share to 0.72. Probabilistic appeals also improve IA, but less cleanly. They reduce strict blocking pairs to 0.4, generate 0.08 wasted-seat claims, and raise the stable share to 0.6.

These findings align closely with the logic of the appeal rules. Strict appeals target precisely the priority violations generated by IA and create almost no waste. The experiment deliberately departs from the theoretical benchmark by requiring participants to choose which claims to pursue, whereas the model assumes that every profitable appeal is submitted. Appendix E shows that imposing the model's appeal-stage strategy sharpens the main findings. Holding submitted rank-order lists and first-stage assignments fixed, IA's remaining strict blocking pairs fall from 0.33 to 0.04 per market-round, while

its procedurally stable share rises from 0.72 to 0.94. Thus, the residual instability observed in the experiment does not reflect a failure of the appeal rule. When appeal-stage behavior follows the theoretical benchmark, strict appeals eliminate almost all of IA's procedural instability while preserving its substantial efficiency advantage. The small remaining instability reflects that the empirical measure is evaluated against submitted preferences and also requires non-wastefulness, whereas Lemma 1 concerns appealable strict claims involving schools that the student genuinely prefers.

The DA comparison is the reverse. DA without appeals contains no strict blocking pair and no waste when evaluated against submitted preferences. Strict appeals leave this outcome unchanged, so every market-round remains stable. Probabilistic appeals instead disturb an initially stable assignment. They generate 0.3 strict blocking pairs and 0.8 wasted-seat claims per market-round, reducing the stable share to 0.7. The random channel can admit a student without a priority claim and leave her first-stage seat vacant. It therefore creates precisely the instability and waste that strict appeals avoid.

Table 6 reports regressions at the market-round level. The specifications include treatment indicators and round fixed effects. Standard errors are clustered at the matching-group level.

The regressions confirm the descriptive pattern. Under IA, strict appeals reduce strict blocking pairs by 0.5 per market-round and increase the probability of stability by 0.3. Probabilistic appeals also improve IA, but their effects are smaller and they create more waste. Under DA, strict appeals have no effect on any outcome. Probabilistic appeals instead increase both strict blocking pairs and wasted-seat claims and reduce stability by 0.3.

Stability is lower when assignments are evaluated against induced true preferences. Under IA, the stable share rises from 0.20 without appeals to 0.5 with strict appeals and 0.4 with probabilistic appeals. Under DA, it is 0.7 without appeals, 0.7 with strict appeals, and 0.5 with probabilistic appeals.

RESULT 3. *Strict appeals more than halve IA's strict blocking pairs and double its procedurally stable share from* 0.36 *to* 0.72*, while generating almost no waste. Probabilistic appeals also improve IA, but create more waste and leave more blocking pairs. Under DA, strict appeals leave the stable first-stage outcome unchanged, whereas probabilistic appeals create both blocking pairs and waste.*

6.5 *Trade-offs*

The results do not produce a complete ranking of the mechanisms. Instead, the consequences of appeals depend on the first-stage assignment rule. IA creates priority violations that appeals can correct, whereas DA begins from a procedurally stable assignment and therefore leaves strict appeals little to do.

Figure 5 summarizes rank-efficiency, truth-telling, and procedural stability. Each dimension is normalized by its across-treatment maximum, so that higher values indicate better performance. The figure is intended only as a visual comparison and does not constitute an aggregate welfare index.

TABLE 6. Impact of appeals on stability.

	Outcomes		
	Blocking pairs	Wasted-seat claims	Stable market-round
	(1)	(2)	(3)
Panel A. IA treatments			
IA strict	-0.52^{***}	0.03^{***}	0.34^{***}
	(0.06)	(0.01)	(0.03)
IA probabilistic	-0.45^{***}	0.08^{***}	0.23^{***}
	(0.06)	(0.01)	(0.04)
Panel B. DA treatments			
DA strict	0.00	0.00	0.00
	(0.00)	(0.00)	(0.00)
DA probabilistic	0.31^{***}	0.78^{***}	-0.33^{***}
	(0.05)	(0.09)	(0.04)
Round effects	Yes	Yes	Yes
Observations	1,435	1,435	1,435
Matching groups	59	59	59

*** $p < 0.01$; ** $p < 0.05$; * $p < 0.10$. Coefficients are reported from linear regressions. Column (3) is a linear probability model. The reference categories are the corresponding mechanisms without appeals. All outcomes are evaluated against submitted preferences. A strict blocking pair requires a strict difference in the underlying priority classes; equal-priority tie-break claims are excluded. A market-round is stable if it contains neither a strict blocking pair nor a wasted-seat claim. Each observation is a complete seven-person market-round; market-rounds affected by participant exits are excluded. Matching groups are experimental sessions. Standard errors are clustered at the matching-group level.

Within IA. Both appeal rules improve truth-telling, efficiency, and stability. Relative to no appeals, truth-telling rises from 0.28 to approximately 0.39, while mean true-preference rank improves from 2.33 to approximately 2.05. Procedural stability rises from 0.36 to 0.72 with strict appeals and 0.61 with probabilistic appeals. The two appeal rules therefore deliver similar efficiency and truth-telling gains, but strict appeals leave fewer blocking pairs and substantially less waste.

Between IA and DA. IA with strict appeals performs substantially better than DA with strict appeals on efficiency and truth-telling, but DA remains fully procedurally stable. Strict appeals nevertheless close more than half of IA's initial stability deficit. Probabilistic appeals produce a different comparison: they improve IA along all three dimensions but undermine DA's principal advantage, reducing its procedural stability from 1.00 to 0.65.

The mechanism and the appeal rule must therefore be evaluated jointly. Strict appeals are particularly valuable under IA because they correct genuine priority violations while creating almost no waste. Under DA they leave the original assignment essentially

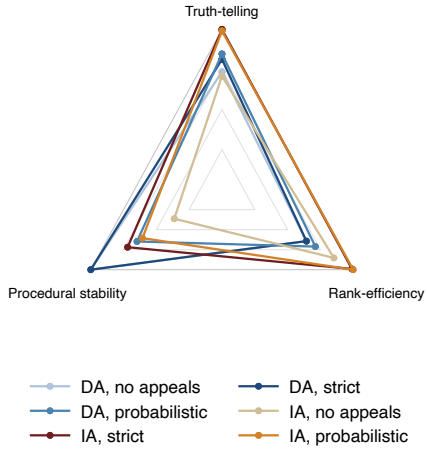


FIGURE 5. Efficiency, truth-telling, and stability across treatments.
Rank-efficiency is based on assigned rank evaluated against induced true preferences. Stability is the share of market-rounds containing neither a strict blocking pair nor a wasted-seat claim, evaluated against submitted preferences. Each axis is normalized by its across-treatment maximum. The normalization is descriptive and does not assign welfare weights to the three outcomes.

unchanged. Probabilistic appeals are less targeted and can destabilize an otherwise stable DA outcome.

7. CONCLUSION

Appeals are part of the allocation mechanism, not an afterthought. Strict priority-based appeals leave Deferred Acceptance largely unchanged but make Immediate Acceptance less manipulable, improve student assignments, and reduce priority violations. Much of the experimental improvement appears before any appeal is decided because appeal rights change submitted rankings. Probabilistic appeals also improve Immediate Acceptance, but can create waste and instability after Deferred Acceptance. Assignment mechanisms should therefore be designed and evaluated jointly with the appeal rules that follow them.

REFERENCES

ABDULKADIROĞLU, A., Y.-K. CHE, AND Y. YASUDA (2011): "Resolving conflicting preferences in school choice: The "Boston mechanism" reconsidered," *American Economic Review*, 101, 399–410. [6]

ABDULKADIROGLU, A., P. PATHAK, AND A. ROTH (2005): "The New York City high school match," *American Economic Review*, 95, 364–367. [7]

AFACAN, M. O. (2012): "Group robust stability in matching markets," *Games and Economic Behavior*, 74, 394–398. [5]

AFACAN, M. O., Z. H. ALIOGULLARI, AND M. BARLO (2017): "Sticky matching in school choice," *Economic Theory*, 64, 509–538. [5]

AFACAN, M. O., U. DUR, AND M. VAN DER LINDEN (2024): "Capacity design in school choice," *Games and Economic Behavior*, 146, 277–291. [6]

AGARWAL, N. AND P. SOMAINI (2018): "Demand analysis using strategic reports: An application to a school choice mechanism," *Econometrica*, 86, 391–444. [6]

AKBARPOUR, M., A. KAPOR, C. NEILSON, W. VAN DIJK, AND S. ZIMMERMAN (2022): "Centralized school choice with unequal outside options," *Journal of Public Economics*, 210, 104644. [6]

AKYOL, E., I. HAFALIR, AND A. MIRALLES (2026): "Bayesian school choice: welfare comparison of immediate acceptance and deferred acceptance mechanisms," *Mathematics of Operations Research*, forthcoming. [6]

ARTEMOV, G. (2021): "Assignment mechanisms: Common preferences and information acquisition," *Journal of Economic Theory*, 198, 105370. [6]

AYGÜN, O. AND G. E. BILGIN (2023): "Placement with assignment guarantees and semi-flexible capacities," *unpublished*. [6]

CALSAMIGLIA, C., C. FU, AND M. GÜELL (2020): "Structural estimation of a model of school choices: The boston mechanism versus its alternatives," *Journal of Political Economy*, 128, 642–680. [6]

CERRONE, C., Y. HERMSTRÜWER, AND O. KESTEN (2024): "School Choice with Consent: An Experiment," *The Economic Journal*, 134, 1760–1805. [21]

CHEN, Y. AND Y. HE (2021): "Information acquisition and provision in school choice: an experimental study," *Journal of Economic Theory*, 197, 105345. [6]

——— (2022): "Information acquisition and provision in school choice: a theoretical investigation," *Economic Theory*, 74, 293–327. [6]

CHEN, Y. AND T. SÖNMEZ (2006): "School choice: an experimental study," *Journal of Economic theory*, 127, 202–231. [18]

DECERF, B., G. HAERINGER, AND M. VAN DER LINDEN (2024): "Incontestable Assignments," *arXiv preprint arXiv:2401.03598*. [6]

ERGIN, H. AND T. SÖNMEZ (2006): "Games of school choice under the Boston mechanism," *Journal of Public Economics*, 90, 215 – 237. [3, 6, 13]

FREER, M., T. KLEIN, AND J. ORTEGA (2026): "Experimental School Choice with Parents," *arXiv preprint arXiv:2603.24615*. [21]

FU, C., D. W. HAHM, AND M. PARK (2024): "Information, match quality, and the design of student-school assignment mechanisms," *unpublished*. [6]

HAKIMOV, R. AND D. KÜBLER (2021): "Experiments on centralized school choice and college admissions: a survey," *Experimental Economics*, 24, 434–488. [18]

HAKIMOV, R. AND M. RAGHAVAN (2026): "Improving transparency and verifiability in school admissions: Theory and experiment," *Management Science*, 72, 2867–2894. [6]

HAMILTON, A., J. MADISON, AND J. JAY (2003): *The Federalist Papers*, New York: Signet Classics. [3]

HE, Y. (2015): "Gaming the Boston school choice mechanism in Beijing," . [6]

HUNT, E. (2019): "Fair access to schools? The impact of the appeals and waiting list system," Tech. rep., Education Policy Institute. [2, 5]

KAPOR, A., M. KARNANI, AND C. NEILSON (2024): "Aftermarket frictions and the cost of off-platform options in centralized assignment mechanisms," *Journal of Political Economy*, 132, 2346–2395. [6]

KOJIMA, F. (2011): "Robust stability in matching markets," *Theoretical Economics*, 6, 257–267. [5]

LOCKE, J. (1988): *Two Treatises of Government*, Cambridge: Cambridge University Press. [3]

MAXEY, T. (2024): "School choice with costly information acquisition," *Games and Economic Behavior*, 143, 248–268. [6]

MIRALLES, A. (2009): "School choice: The case for the Boston mechanism," in *International conference on auctions, market mechanisms and their applications*, Springer, 58–60. [6]

MONTESQUIEU, C. D. S. (1989): *The Spirit of the Laws*, Cambridge: Cambridge University Press. [3]

MORRILL, T. AND A. E. ROTH (2024): "Top trading cycles," *Journal of Mathematical Economics*, 112, 102984. [7]

ORTEGA, J. (2026): "A note on the strategic vulnerability of the Boston mechanism in random markets," *Social Choice and Welfare*, 1–12. [13]

PATHAK, P. A. AND T. SÖNMEZ (2013): "School admissions reform in Chicago and England: Comparing mechanisms by their vulnerability to manipulation," *American Economic Review*, 103, 80–106. [3, 11]

TAYLOR, C., S. GORARD, AND J. FITZ (2002): "Market frustration? Admission appeals in the UK education market," *Educational Management & Administration*, 30, 243–260. [5, 7]

SUPPLEMENTARY APPENDIX A: PROOFS OMITTED IN MAIN TEXT.

A.1 *Proof of Lemma 1.*

PROOF. Let

$$\mu^0 = M(\hat{\succ}) \qquad \text{and} \qquad \mu = (M \circ r_1)(\hat{\succ}).$$

Weak Pareto dominance with respect to true preferences is immediate. Each student either keeps her first-stage assignment under μ^0 or is reassigned through an upheld appeal to a school she truly prefers to it.

We now show that μ admits no strict blocking pair (i,s) such that $s \in A_i(\hat{\succ}, M)$ and $s \succ_i \mu_i$. Suppose, toward a contradiction, that such a pair exists. Since $\mu_i \succsim_i \mu_i^0$, we also have $s \succ_i \mu_i^0$. Therefore $s \in a_i$, so student i appealed to s.

Because (i,s) is a strict blocking pair of μ, there exists $k \in \mu_s^{-1}$ such that $i \rhd_s k$. If $k \in (\mu_s^0)^{-1}$, then (i,s) was already a strict blocking pair of the initial matching μ^0, and r_1 upheld student i's appeal to s.

Now suppose $k \notin (\mu_s^0)^{-1}$. Then k was assigned to s through a successful appeal. Hence, by definition of r_1, there exists an initial occupant $j \in (\mu_s^0)^{-1}$ such that $k \rhd_s j$. Since priorities are transitive and $i \rhd_s k$, we have $i \rhd_s j$. Thus (i,s) was already a strict blocking pair of μ^0, and again r_1 upheld student i's appeal to s.

In either case, i's appeal to s was upheld. Since final reassignment selects the student's most preferred school according to $\succ_i$ among her upheld appeals, $\mu_i \succsim_i s$, contradicting $s \succ_i \mu_i$. □

A.2 *Proof of Proposition 1.*

We provide the full proof that we only sketched in the main text.

PROOF. We first prove containment. Fix a preference profile $(\succ_i, \succ_{-i})$ and suppose student i can profitably manipulate $\texttt{IA} \circ r_1$ at this profile. Let $\hat{\succ}_i$ be a profitable deviation and write

$$x = (\texttt{IA} \circ r_1)_i(\succ_i, \succ_{-i}) \qquad \text{and} \qquad y = (\texttt{IA} \circ r_1)_i(\hat{\succ}_i, \succ_{-i}).$$

Thus

$$y \succ_i x.$$

Since r_1 only adds successful appeals to the first-stage IA outcome and never worsens a student's assignment,

$$x \succsim_i (\texttt{IA} \circ r_0)_i(\succ_i, \succ_{-i}).$$

We show that student i can also obtain y under $\texttt{IA} \circ r_0$ by a suitable misreport.

If student i already obtains y in the first-stage IA outcome under $(\hat{\succ}_i, \succ_{-i})$, then the same report $\hat{\succ}_i$ gives student i school y under $\texttt{IA} \circ r_0$. Therefore

$$(\texttt{IA} \circ r_0)_i(\hat{\succ}_i, \succ_{-i}) = y \succ_i x \succsim_i (\texttt{IA} \circ r_0)_i(\succ_i, \succ_{-i}),$$

so $\hat{\succ}_i$ is a profitable manipulation of $\texttt{IA} \circ r_0$.

Suppose instead that student i obtains y only through a successful appeal under r_1. Let μ be the first-stage IA outcome under $(\hat{\succ}_i, \succ_{-i})$. Since i successfully appeals to y,

there is a regular-seat occupant j of y in μ such that

$$i \triangleright_y j.$$

Let $\hat{\succ}_i'$ be any report that ranks y first. We claim that student i obtains y in the first stage of IA under $(\hat{\succ}_i', \succ_{-i})$.

Let A_y^1 be the set of students other than i who rank y first. This set depends only on $\succ_{-i}$. If fewer than q_y students in A_y^1 have higher priority at y than i, then i is among the q_y highest-priority round-1 applicants to y when she ranks y first. Hence i is admitted to y in round 1.

So suppose, by way of contradiction, that at least q_y students in A_y^1 have higher priority at y than i. Then, independently of i's report, the q_y regular seats of y are filled in round 1 by students who all have higher priority than i. Since IA assignments are permanent, every regular-seat occupant of y in the first-stage outcome μ has higher priority than i. This contradicts the existence of a regular-seat occupant j of y with

$$i \triangleright_y j.$$

Therefore fewer than q_y students in A_y^1 have higher priority at y than i, and i is admitted to y in round 1 when she ranks y first. Thus

$$(\mathtt{IA} \circ r_0)_i(\hat{\succ}_i', \succ_{-i}) = y.$$

Consequently,

$$(\mathtt{IA} \circ r_0)_i(\hat{\succ}_i', \succ_{-i}) = y \succ_i x \succsim_i (\mathtt{IA} \circ r_0)_i(\succ_i, \succ_{-i}).$$

Hence $\hat{\succ}_i'$ is a profitable manipulation of $\mathtt{IA} \circ r_0$.

Thus, whenever student i can profitably manipulate $\mathtt{IA} \circ r_1$ at a profile, the same student can profitably manipulate $\mathtt{IA} \circ r_0$ at that profile. Therefore $\mathtt{IA} \circ r_0$ is as manipulable as $\mathtt{IA} \circ r_1$.

It remains to show that the containment is strict. Consider the following school choice problem with three students i_1, i_2, i_3 and three schools s_1, s_2, s_3, each with one seat. Preferences and priorities are (manipulations in parentheses):

EXAMPLE 2.

$\succ_{i_1}$	$(\succ'_{i_1})$	$\succ_{i_2}$	$\succ_{i_3}$	$\trianglerighteq_{s_1}$	$\trianglerighteq_{s_2}$	$\trianglerighteq_{s_3}$
s_1	(s_2)	s_1	s_2	i_2	i_1	$\cdot$
s_2	(s_1)	s_3	s_3	i_1	i_3	$\cdot$
s_3	(s_3)	s_2	s_1	i_3	i_2	$\cdot$

Under truthful reporting, the first-stage IA outcome is

$$(i_1 - s_3,\ i_2 - s_1,\ i_3 - s_2).$$

Indeed, in round 1, i_1 and i_2 apply to s_1, and i_2 is accepted; i_1 is rejected. Student i_3 applies to s_2 and is accepted. In round 2, i_1 applies to s_2 but is rejected because s_2 is already full. In round 3, i_1 applies to s_3 and is accepted.

Under $\texttt{IA} \circ r_0$, student i_1 can profitably manipulate by reporting

$$s_2 \succ'_{i_1} s_1 \succ'_{i_1} s_3.$$

Then in round 1, student i_1 applies to s_2 and is accepted, i_2 applies to s_1 and is accepted, and i_3 applies to s_2 but is rejected. Student i_3 then applies to s_3 and is accepted. Hence the outcome is

$$(i_1 - s_2,\ i_2 - s_1,\ i_3 - s_3),$$

which is strictly better for i_1 than her truthful $\texttt{IA} \circ r_0$ assignment s_3.

Now consider $\texttt{IA} \circ r_1$ at the truthful profile. The first-stage outcome is again

$$(i_1 - s_3,\ i_2 - s_1,\ i_3 - s_2).$$

Student i_1 has a successful appeal to s_2, because i_1 ranks s_2 above her first-stage assignment s_3 and

$$i_1 \rhd_{s_2} i_3.$$

Thus, after appeals, i_1 obtains s_2.

No student can profitably manipulate $\texttt{IA} \circ r_1$ at this profile. Students i_2 and i_3 obtain their top choices truthfully, so they cannot improve. Student i_1 obtains s_2 truthfully after appeals. Her only strictly better school is s_1. But i_2 ranks s_1 first and has higher priority than i_1 at s_1, so i_1 cannot obtain s_1 directly by any report and cannot successfully appeal to s_1. Therefore i_1 cannot profitably manipulate $\texttt{IA} \circ r_1$.

Hence this profile is manipulable under $\texttt{IA} \circ r_0$ but not under $\texttt{IA} \circ r_1$. Therefore $\texttt{IA} \circ r_0$ is more manipulable than $\texttt{IA} \circ r_1$.

□

SUPPLEMENTARY APPENDIX B: ANALYZING r_2.

We now consider the probabilistic appeal rule r_2. Since r_2 is random, the outcome of a mechanism followed by r_2 is a lottery over final matchings rather than a deterministic matching. We therefore use stochastic dominance to define incentive and welfare comparisons.

B.1 *Definitions*

Random matchings Let $\mathcal{M}$ denote the set of possibly infeasible matchings. A *randomized matching* is a probability distribution

$$\lambda \in \Delta(\mathcal{M}).$$

For a student i, the *marginal lottery* induced by λ is denoted by λ_i, where

$$\lambda_i(s) = \sum_{\mu \in \mathcal{M}:\, \mu_i = s} \lambda(\mu) \qquad \text{for each } s \in S \cup \{s_\emptyset\}.$$

For two marginal lotteries λ_i and λ_i', we write

$$\lambda_i \succeq_i^{sd} \lambda_i'$$

if λ_i *first-order stochastically dominates* λ_i' with respect to $\succ_i$: for every school $s \in S \cup \{s_\emptyset\}$,

$$\sum_{t:\, t \succeq_i s} \lambda_i(t) \geq \sum_{t:\, t \succeq_i s} \lambda_i'(t).$$

Equivalently, $\lambda_i \succeq_i^{sd} \lambda_i'$ if every von Neumann – Morgenstern utility representation u_i of $\succ_i$ gives weakly higher expected utility under λ_i than under λ_i'.

For two randomized matchings λ and λ', we write

$$\lambda \succeq^{sd} \lambda'$$

if

$$\lambda_i \succeq_i^{sd} \lambda_i'$$

for every student i. We write $\lambda_i \succ_i^{sd} \lambda_i'$ if $\lambda_i \succeq_i^{sd} \lambda_i'$ and the inequality defining stochastic dominance is strict for at least one upper contour set.

Strategy-proofness and equilibrium for randomized mechanisms. A randomized mechanism Φ maps each reported preference profile σ to a lottery

$$\Phi(\sigma) \in \Delta(\mathcal{M}).$$

It is *sd-strategy-proof* if, for every student i, every true preference profile $\succ$, and every report $\hat{\sigma}_i$,

$$\Phi_i(\succ_i, \succ_{-i}) \succeq_i^{sd} \Phi_i(\hat{\sigma}_i, \succ_{-i}).$$

A report profile σ^* is an *sd-Nash equilibrium* of Φ if there is no student i and no unilateral report $\hat{\sigma}_i$ such that

$$\Phi_i(\hat{\sigma}_i, \sigma^*_{-i}) \succ_i^{sd} \Phi_i(\sigma^*).$$

This is the ordinal equilibrium notion appropriate for the present model: it rules out deviations that are unambiguously profitable for every von Neumann–Morgenstern utility representation of $\succ_i$.

The probabilistic appeal rule r_2. Fix $p \in (0,1)$. Given a first-stage matching μ and a reported preference profile σ, the rule r_2 generates a lottery over final matchings as follows. A student i may appeal only to schools that she ranked above her first-stage match μ_i under σ_i and from which she was rejected. For every admissible appeal by student i to school s:

- if (i,s) is a strict blocking pair in μ, that is, if $i \rhd_s j$ for some $j \in \mu_s^{-1}$, the appeal is upheld with probability one;
- otherwise, the appeal is upheld independently with probability p.

Thus r_2 nests r_1: every appeal upheld under r_1 is upheld under r_2 with certainty, while appeals not upheld under r_1 are upheld independently with probability p. For each realization of appeal outcomes, student i is assigned to the highest-ranked school under her true preference $\succ_i$ among the upheld appeals, if any, and otherwise remains at her first-stage match μ_i. The resulting lottery over final matchings is denoted

$$r_2(\mu, \sigma).$$

For any first-stage mechanism M, we write

$$(M \circ r_2)(\sigma) = r_2(M(\sigma), \sigma).$$

B.2 *Welfare Ordering of Appeal Rules*

The probabilistic rule nests the strict rule: it preserves every appeal that succeeds under r_1 and may uphold additional appeals. Since the student chooses her most-preferred upheld claim according to her true preference, the comparison is pointwise and therefore also holds in stochastic-dominance terms.

PROPOSITION 3. *For any mechanism M, any true preference profile $\succ$, and any reported preference profile σ,*

$$(M \circ r_2)(\sigma) \succeq^{sd} (M \circ r_1)(\sigma).$$

PROOF. Fix a student i, and let $\mu = M(\sigma)$ be the first-stage matching. Under r_1, student i keeps μ_i or moves to her most-preferred school under $\succ_i$ among the appeals supported by a strict blocking pair. Call this deterministic outcome a_i.

Under r_2, every appeal upheld under r_1 is upheld with probability one, while additional admissible appeals may also be upheld. Thus, in every realization of the r_2 lottery, the set from which student i's final assignment is selected contains a_i, possibly together

with schools she prefers to a_i. Her realized assignment under r_2 is therefore weakly preferred to a_i according to $\succ_i$. Hence

$$(M \circ r_2)_i(\sigma) \succeq_i^{sd} (M \circ r_1)_i(\sigma).$$

Since i was arbitrary, the result follows. □

Because random appeals may be upheld without a priority claim, the realized matchings under r_2 need not contain fewer blocking pairs.

B.3 *A Stochastic Equilibrium Reversal*

The reversal result also has an equilibrium counterpart under probabilistic appeals. The relevant equilibrium notion is sd-Nash equilibrium, because the model specifies ordinal preferences but not a particular cardinal utility representation.

Let

$$\mu = \mathtt{DA}(\succ).$$

For each student i, define the *DA-cutoff report* σ_i^μ as follows. If $\mu_i \in S$, student i lists, in her true order, every school that she weakly prefers to μ_i, and truncates her list immediately after μ_i. If $\mu_i = s_\emptyset$, she lists every school she prefers to being unassigned, again in her true order. Let

$$\sigma^\mu = (\sigma_i^\mu)_{i \in I}.$$

THEOREM 2 (Stochastic reversal under probabilistic appeals). *For every school choice problem with strict priorities and every $p \in (0,1)$, the DA-cutoff profile σ^μ is an sd-Nash equilibrium of $\mathtt{IA} \circ r_2$. Moreover,*

$$(\mathtt{IA} \circ r_2)(\sigma^\mu) \succeq^{sd} \mathtt{DA}(\succ),$$

where the deterministic DA assignment is viewed as a degenerate randomized matching. If the two randomized matchings differ, at least one student is strictly better off in stochastic-dominance terms.

PROOF. Write

$$\nu = \mathtt{IA}(\sigma^\mu) \qquad \text{and} \qquad x = (\mathtt{IA} \circ r_1)(\sigma^\mu).$$

We first show that

$$x_i \succeq_i \mu_i \qquad \text{for every student } i.$$

If $\nu_i \neq s_\emptyset$, this is immediate: the DA-cutoff report contains no school that i ranks below μ_i. Now suppose that $\nu_i = s_\emptyset$ and $\mu_i = s \in S$. Student i applied to s and was rejected. School s must be full under μ. Indeed, if it had an empty DA seat, non-wastefulness of DA would imply that no student assigned elsewhere under μ ranks s above her DA assignment. Hence only students in μ_s^{-1} could apply to s under the DA-cutoff profile, and i could not be rejected.

Since s is full under μ, it has q_s DA assignees, one of whom is i. The q_s first-stage occupants of s under ν exclude i, so at least one of them, say j, is not assigned to s under

μ. Because j applies to s under her DA-cutoff report,

$$s \succ_j \mu_j.$$

Stability of μ and strict priorities imply that every DA assignee at s, including i, has higher priority at s than j. Thus (i,s) is a strict blocking pair of ν, and r_1 upholds i's appeal to s. Hence $x_i \succeq_i \mu_i$.

Next observe that every school $t \succ_i x_i$ is an admissible appeal at σ^μ. The report σ_i^μ follows the true order, and $t \succ_i x_i \succeq_i \nu_i$, so i applied to t before reaching her first-stage assignment and was rejected. Moreover, the appeal to t is not supported by a strict blocking pair; otherwise r_1 would assign i to t or to a still better school, contradicting the definition of x_i. Under r_2, student i therefore has an independent probability-p appeal ticket to every school she strictly prefers to x_i, while x_i is guaranteed.

We now show that σ^μ is an sd-Nash equilibrium. Fix student i and a unilateral deviation $\hat{\sigma}_i$. Let

$$\hat{\nu}_i = \mathtt{IA}_i(\hat{\sigma}_i, \sigma^\mu_{-i}) \qquad \text{and} \qquad \hat{x}_i = (\mathtt{IA} \circ r_1)_i(\hat{\sigma}_i, \sigma^\mu_{-i})$$

denote the first-stage and strict-rule outcomes under the deviation.

First suppose that $\hat{x}_i \preceq_i x_i$. For any upper contour set whose cutoff is weakly below x_i, the equilibrium lottery assigns probability one to that set. For an upper contour set strictly above x_i, the deviation can enter the set only through non-strict probability-p appeals: a first-stage assignment or a successful strict appeal in that set would imply $\hat{x}_i \succ_i x_i$. At σ^μ, by contrast, i holds a probability-p ticket to every school in the upper contour set. Independence of the appeal draws therefore implies that the equilibrium probability of reaching each upper contour set is weakly greater than under the deviation. Thus the equilibrium lottery sd-dominates the deviation lottery.

Now suppose that $\hat{x}_i \succ_i x_i$. We claim that at least one school $t \succ_i \hat{x}_i$ is not an admissible appeal under the deviation. Suppose, to the contrary, that every such school is appealable.

If $\hat{x}_i = \hat{\nu}_i$, then every school preferred to $\hat{x}_i$ must be ranked before $\hat{x}_i$ and must reject i. Reordering applications at which i is rejected does not change any other student's IA history: in each such round the same applicants are accepted as if i had not applied. Student i therefore reaches $\hat{x}_i$ no earlier than under the true-order DA-cutoff report and faces the same or a fuller school. But she is rejected from $\hat{x}_i$ under σ_i^μ, because $\hat{x}_i \succ_i x_i$. Hence she cannot obtain $\hat{x}_i$ directly.

If instead $\hat{x}_i$ is obtained through a strict appeal, student i is rejected from $\hat{x}_i$ in the first stage. Applying to $\hat{x}_i$ later cannot create a new strict claim because IA acceptances are permanent. If she applies earlier, then either the school is already full, in which case any lower-priority occupant would also be present when she applies under σ_i^μ, or its remaining seats are filled in that round by applicants who outrank her. In neither case can moving the application earlier turn the non-strict rejection at σ_i^μ into a strict blocking pair. Thus $\hat{x}_i$ cannot be obtained through a new strict appeal either.

This contradiction proves the claim. Let

$$T_i(\hat{x}_i) = \{t \in S : t \succ_i \hat{x}_i\}.$$

At the equilibrium profile, every school in $T_i(\hat{x}_i)$ is an independent probability-p appeal ticket. Under the deviation, at least one of these tickets is missing, and no school in $T_i(\hat{x}_i)$ can be obtained directly or through a strict appeal by the definition of $\hat{x}_i$. Hence

$$\Pr_{(\mathtt{IA}\circ r_2)_i(\sigma^\mu)}\big(T_i(\hat{x}_i)\big) = 1-(1-p)^{|T_i(\hat{x}_i)|}$$

is strictly larger than the corresponding probability under the deviation, which is at most

$$1-(1-p)^{|T_i(\hat{x}_i)|-1}.$$

The deviation therefore cannot stochastically dominate the equilibrium lottery. Since the deviation was arbitrary, σ^μ is an sd-Nash equilibrium.

Finally, $x_i \succeq_i \mu_i$ for every i, and r_2 preserves every assignment obtained under r_1 while possibly adding better random outcomes. It follows that

$$(\mathtt{IA}\circ r_2)(\sigma^\mu) \succeq^{sd} \mu = \mathtt{DA}(\succ).$$

If the two randomized matchings differ, at least one marginal stochastic-dominance comparison is strict. □

B.4 *DA Incentive Properties Remain*

PROPOSITION 4. *For every $p \in (0,1)$, $\mathtt{DA}\circ r_2$ is sd-strategy-proof.*

PROOF. Fix student i, true preferences $\succ_i$, and reports $\succ_{-i}$ of the other students. Let

$$\mu = \mathtt{DA}(\succ_i,\succ_{-i}) \quad \text{and} \quad \nu = \mathtt{DA}(\hat{\sigma}_i,\succ_{-i})$$

be the truthful and misreported DA outcomes. Since DA is strategy-proof,

$$\mu_i \succeq_i \nu_i.$$

Let

$$\Lambda = (\mathtt{DA}\circ r_2)(\succ_i,\succ_{-i}) \quad \text{and} \quad \widehat{\Lambda} = (\mathtt{DA}\circ r_2)(\hat{\sigma}_i,\succ_{-i})$$

be the induced lotteries. We show

$$\Lambda_i \succeq_i^{sd} \widehat{\Lambda}_i.$$

Fix a school s, and let

$$U_i(s) = \{t \in S \cup \{s_\emptyset\} : t \succeq_i s\}.$$

Case 1: $\mu_i \in U_i(s)$. Truthful reporting gives student i a first-stage match in $U_i(s)$. Since appeals can only move a student to schools ranked above her first-stage match, every realization under truthful reporting places i in $U_i(s)$. Hence

$$\Pr_{\Lambda_i}(U_i(s)) = 1 \geq \Pr_{\widehat{\Lambda}_i}(U_i(s)).$$

Case 2: $\mu_i \notin U_i(s)$. Every school in $U_i(s)$ is strictly preferred by i to μ_i. Under truthful DA, student i applies to every such school before reaching μ_i, and is rejected from each of them. Hence every school in $U_i(s)$ is an admissible appeal school under truthful reporting.

Under the misreport, student i can receive a school in $U_i(s)$ after the appeal stage only if some admissible appeal to a school in $U_i(s)$ is upheld. Every such school is also admissible under truthful reporting, by the previous paragraph. Thus the set of truthful admissible appeals to schools in $U_i(s)$ contains the set of misreport-induced admissible appeals to schools in $U_i(s)$.

Moreover, no such appeal is upheld with probability one under either report. Indeed, DA is stable with respect to the reported preference profile and the priority order used by the mechanism. Thus, whenever student i is rejected from a school t that she ranked above her DA match, (i,t) is not a strict blocking pair of the DA outcome. Consequently, under r_2, any admissible appeal by i to such a school is upheld with probability p, and not with probability one.

Since truthful reporting gives i a superset of the probability-p appeal tickets to schools in $U_i(s)$, the probability that at least one appeal to a school in $U_i(s)$ is upheld is weakly higher under truthful reporting than under the misreport. Therefore

$$\Pr_{\Lambda_i}(U_i(s)) \geq \Pr_{\widehat{\Lambda}_i}(U_i(s)).$$

Since this holds for every upper contour set,

$$\Lambda_i \succeq_i^{sd} \widehat{\Lambda}_i.$$

Because i and $\hat{\sigma}_i$ were arbitrary, $\texttt{DA} \circ r_2$ is sd-strategy-proof. □

The proof reveals two distinct channels through which truthful reporting dominates under $\texttt{DA} \circ r_2$. First, DA's strategy-proofness ensures a weakly better first-stage outcome. Second, truthful reporting generates weakly more appeal opportunities, since listing a school above one's match is necessary for filing an appeal, and truthful reporting lists every truly preferred school. These channels reinforce each other: the truthful report maximizes both the fallback, namely the DA match, and the set of lottery tickets.

B.5 *IA becomes Less Manipulable*

Unlike DA, IA remains manipulable under r_2. However, probabilistic appeals reduce the gain from standard IA manipulations. Under r_2, a rejected application to a desired school is not wasted: even absent a priority claim, it creates a lottery ticket. We formalize this effect in two propositions.

PROPOSITION 5 (Lottery value of truthful applications). *Fix $p \in (0,1)$, a preference profile $\succ$, a student i, and a von Neumann – Morgenstern utility representation u_i of $\succ_i$. Consider two reports for i. Under report σ_i, student i applies to school b, is rejected, and receives first-stage match a, where*

$$b \succ_i a,$$

and the appeal to b is admissible under r_2. Under report $\hat{\sigma}_i$, student i also receives a in the first stage but does not generate an admissible appeal to b. Holding fixed all other appeal opportunities:

1. *if the appeal to b is upheld under r_1, that is, if (i,b) is a strict blocking pair, the expected utility gain from σ_i over $\hat{\sigma}_i$ is*

$$u_i(b) - u_i(a);$$

2. *if the appeal to b is not upheld under r_1, the expected utility gain is*

$$p\big(u_i(b) - u_i(a)\big).$$

PROOF. Under $\hat{\sigma}_i$, student i receives a with no appeal to b. Holding fixed all other appeal opportunities, the comparison-relevant payoff is $u_i(a)$.

Under σ_i, student i receives a in the first stage but has an admissible appeal to b. If (i,b) is a strict blocking pair, r_2 upholds the appeal with probability one, giving payoff $u_i(b)$. The gain is therefore

$$u_i(b) - u_i(a).$$

If (i,b) is not a strict blocking pair, r_2 upholds the appeal with probability p. The expected payoff is

$$pu_i(b) + (1-p)u_i(a),$$

and the gain is

$$pu_i(b) + (1-p)u_i(a) - u_i(a) = p\big(u_i(b) - u_i(a)\big).$$

□

PROPOSITION 6 (Reduced gain from safe-school manipulation). *Fix $p \in (0,1)$, a preference profile $\succ$, a student i, and a von Neumann – Morgenstern utility representation u_i of $\succ_i$. Suppose that under truthful IA, student i applies to school b, is rejected, and obtains a, with*

$$b \succ_i a.$$

Suppose also that a deviation $\hat{\sigma}_i$ obtains b directly in the first stage. Holding fixed appeal opportunities to schools strictly preferred to b:

1. *if the truthful appeal to b is upheld under r_1, the deviation yields no gain from obtaining b;*
2. *if the truthful appeal to b is not upheld under r_1, the expected utility gain from the deviation is at most*

$$(1-p)\big(u_i(b) - u_i(a)\big).$$

PROOF. Under truthful reporting, student i receives a in the first stage and has an admissible appeal to b.

If the appeal is upheld under r_1, then r_2 upholds it with probability one. Truthful reporting already gives i school b, so the deviation yields no additional gain from b.

If the appeal is not upheld under r_1, then r_2 upholds it with probability p. Truthful expected utility is at least

$$pu_i(b) + (1-p)u_i(a),$$

ignoring additional appeal opportunities that only increase truthful expected utility. The deviation gives b in the first stage, so the comparison payoff is at most $u_i(b)$. The gain is at most

$$u_i(b) - \big[pu_i(b) + (1-p)u_i(a)\big] = (1-p)\big(u_i(b) - u_i(a)\big).$$

□

Propositions 5 and 6 capture a local incentive effect: probabilistic appeals create a positive option value for truthful applications to risky schools. For this particular safe-school manipulation channel, the direct gain from obtaining b in the first stage is reduced by a factor of $1-p$, and vanishes as $p \to 1$. This does not make IA sd-strategy-proof, because other manipulations, including preemption and manipulations that create appeal opportunities to still better schools, may remain profitable.

B.6 *Discussion*

The results in this appendix complete the incentive analysis across all three appeal rules. Table B.1 summarizes.

TABLE B.1. Incentive properties across appeal rules

	r_0	r_1	r_2	r_3
DA strategy-proof	yes	yes	yes (sd)	no
IA manipulable	yes	yes, but less	yes, lower expected gains	yes, but less

Among the rules considered here, r_3 is the one under which DA loses strategy-proofness in the presence of weak priorities, because appeals can be upheld within a priority class after a tie-break loss. Under strict priorities, r_3 coincides with r_1, and DA remains strategy-proof. Both r_1 and r_2 preserve DA's strategy-proofness deterministically and in stochastic dominance, respectively, while reducing IA's manipulability.

The probabilistic rule r_2 therefore has a distinctive role. Like r_1, it preserves DA's incentive property, now in stochastic-dominance terms. Under IA, it also supports an sd-Nash equilibrium whose lottery stochastically Pareto-dominates truthful DA. At the same time, r_2 gives students a positive reason to rank desired schools even when they have no priority claim: a rejected application becomes a lottery ticket. The rule does not make IA sd-strategy-proof, but it reduces the expected gain from an important class of safe-school manipulations.

These findings have a direct institutional interpretation. England's School Admission Appeals Code mandates that panels uphold appeals when admission arrangements were misapplied, corresponding to r_1, but panels also exercise discretion in borderline

cases, leading to substantial variation in success rates across local authorities. This discretionary component is captured by r_2. Our results show that this heterogeneity, while potentially concerning from an equity perspective, does not undermine DA's incentive properties and further discourages strategic misreporting under IA.

SUPPLEMENTARY APPENDIX C: ANALYZING r_3.

C.1 r_3, *Pareto improvements, and weak blocking pairs*

Lemma 1 shows that strict appeals produce a Pareto improvement and eliminate every appealable strict blocking pair involving a school that the student truly prefers to her final assignment. The corresponding result under the weak appeal rule is as follows.

LEMMA 4. *For every mechanism M and every report profile $\hat{\succ}$, $(M \circ r_3)(\hat{\succ})$ weakly Pareto-dominates $(M \circ r_0)(\hat{\succ})$ with respect to true preferences and admits no weak blocking pair (i, s) such that*

$$s \in A_i(\hat{\succ}, M) \quad \textit{and} \quad s \succ_i (M \circ r_3)_i(\hat{\succ}).$$

The proof is identical to that of Lemma 1, replacing strict priority comparisons with weak priority comparisons, and is therefore omitted.

C.2 $\texttt{IA} \circ r_0$ *is more manipulable than* $\texttt{IA} \circ r_3$

In the main text, we proved that $\texttt{IA} \circ r_0$ is more manipulable than $\texttt{IA} \circ r_1$. Now we show the same for $\texttt{IA} \circ r_3$.

PROPOSITION 7. $\texttt{IA} \circ r_0$ *is more manipulable than* $\texttt{IA} \circ r_3$.

PROOF FOR r_3. Suppose student i can profitably manipulate $\texttt{IA} \circ r_3$. Let $\succ_i'$ be a profitable report, holding all other reports fixed at $\succ_{-i}$. Let

$$\mu^0 = \texttt{IA}(\succ_i', \succ_{-i})$$

be the first-stage IA outcome under this deviation, and let

$$\mu = (\texttt{IA} \circ r_3)(\succ_i', \succ_{-i})$$

be the final outcome after weak appeals. Write $y = \mu_i$. Since the deviation is profitable,

$$y \succ_i (\texttt{IA} \circ r_3)_i(\succ).$$

Appeals never make a student worse off relative to the first-stage IA outcome, so

$$(\texttt{IA} \circ r_3)_i(\succ) \succeq_i \texttt{IA}_i(\succ).$$

Hence

$$y \succ_i \texttt{IA}_i(\succ).$$

We show that i can also obtain y by manipulating plain IA.

If $y = \mu_i^0$, then i already obtains y in the first-stage IA outcome under the report $\succ_i'$. Therefore the same report is a profitable manipulation of $\texttt{IA} = \texttt{IA} \circ r_0$.

Suppose instead that $y \neq \mu_i^0$. Then i obtains y through a successful weak appeal. Hence

$$y \succ_i' \mu_i^0$$

and there exists a first-stage occupant $j \in (\mu_y^0)^{-1}$ such that

$$i \trianglerighteq_y j.$$

We claim that i obtains y directly under plain IA by reporting y as her first choice. Suppose, toward a contradiction, that she does not. Since i applies to y in the first round, this means that all q_y seats of y are filled in the first round by a set H of students, all different from i, who rank y first and who beat i under the tie-broken priority order used by IA.

Because only i's report has changed, every student in H also ranks y first in the truthful profile. Moreover, the same students in H are accepted by y in the first round of truthful IA: removing i from the first-round applicant pool at y cannot prevent the students in H from filling all q_y seats, and IA acceptances are permanent. Hence every student in H is assigned to y in the truthful first-stage IA outcome.

If some $h \in H$ is tied with i under the original weak priority at y, then

$$i \trianglerighteq_y h.$$

Since h is assigned to y in the truthful first-stage IA outcome and since $y \succ_i \texttt{IA}_i(\succ)$, student i has a truthful weak appeal to y. Under r_3, this appeal is upheld. Therefore the truthful final outcome $(\texttt{IA} \circ r_3)_i(\succ)$ assigns i to y or to a school she prefers to y, contradicting

$$y \succ_i (\texttt{IA} \circ r_3)_i(\succ).$$

Therefore every student in H must strictly outrank i at y. Since the students in H rank y first, they occupy all seats at y from the first round onward under any unilateral report by i. Thus, under the deviation $\succ_i'$, every first-stage occupant of y strictly outranks i. This contradicts the fact that i's appeal to y is successful under r_3, which requires some first-stage occupant $j \in (\mu_y^0)^{-1}$ with

$$i \trianglerighteq_y j.$$

Hence i must be admitted to y directly when she reports y first. Since $y \succ_i \texttt{IA}_i(\succ)$, reporting y first is a profitable manipulation of plain IA. Therefore every profitable manipulation of $\texttt{IA} \circ r_3$ implies a profitable manipulation of $\texttt{IA} \circ r_0 = \texttt{IA}$. □

C.3 $\texttt{DA} \circ r_3$ *is manipulable*

To see that $\texttt{DA} \circ r_3$ is manipulable, consider the school choice problem below with four students and four schools, each with one seat, and a common tie-break that favors students with larger indices.

EXAMPLE 3. Preferences and priorities are given as follows (manipulations in parentheses):

$\succ_{i_1}$	$\succ_{i_2}$	$\succ_{i_3}$	$(\succ'_{i_3})$	$\succ_{i_4}$	$\trianglerighteq_{s_1}$	$\trianglerighteq_{s_2}$	$\trianglerighteq_{s_3}$	$\trianglerighteq_{s_4}$
s_4	s_4	s_2	(s_2)	s_1	i_1	i_2	$\cdot$	i_4
s_1	s_1	s_4	(s_3)	s_2	i_2	i_3, i_4	$\cdot$	i_3
$\cdot$	s_2	s_3	$\cdot$	s_4	i_4	$\cdot$	$\cdot$	i_2
$\cdot$	$\cdot$	$\cdot$	$\cdot$	$\cdot$	$\cdot$	$\cdot$	$\cdot$	i_1

With truthful preference reporting, DA yields the matching $(i_1 - s_1, i_2 - s_2, i_3 - s_3, i_4 - s_4)$. There are no successful appeals under r_3, and thus the DA matching remains unaltered. Now let student i_3 misreport her preferences by submitting $s_2 \succ'_{i_3} s_3$. Now DA yields the first-stage matching $(i_1 - s_1, i_2 - s_4, i_3 - s_3, i_4 - s_2)$. Student i_3 then appeals her rejection at s_2. Since student i_3 and student i_4 have the same priority at s_2, the pair (i_3, s_2) is a weak blocking pair. Hence the appeal succeeds under r_3, and student i_3 is reassigned from s_3 to s_2. Since student i_3 truly prefers s_2 to s_3, this misreport is profitable. Therefore $\texttt{DA} \circ r_3$ is manipulable.

C.4 $\texttt{DA} \circ r_3$ *can generate more successful appeals than* $\texttt{IA} \circ r_3$

The preceding example shows that weak appeals can undermine the incentive properties of DA. We now show that weak appeals can also overturn the natural intuition that stable mechanisms should generate fewer successful appeals. Under weak appeals, DA may create appeal opportunities for students who lose a tie-break against otherwise equal-priority applicants. As a result, truthful DA may generate more successful appeals than truthful IA.

PROPOSITION 8. *There exists a school choice problem with weak priorities such that, under truthful reporting,* $\texttt{IA} \circ r_3$ *generates fewer successful appeals than* $\texttt{DA} \circ r_3$.

PROOF. Consider a school choice problem with seven students, $i_1, \ldots, i_7$, and three schools, s_1, s_2, s_3, each with two seats. Priorities are weak, and ties within a priority class are broken in favour of students with smaller indices. Preferences and weak priorities are as follows.

$\succ_{i_1}$	$\succ_{i_2}$	$\succ_{i_3}$	$\succ_{i_4}$	$\succ_{i_5}$	$\succ_{i_6}$	$\succ_{i_7}$	$\trianglerighteq_{s_1}$	$\trianglerighteq_{s_2}$	$\trianglerighteq_{s_3}$
s_3	s_3	s_2	s_1	s_1	s_2	s_3	i_6, i_7	i_4	i_4, i_5
s_2	s_2	s_1	s_3	s_2	s_3	s_2	i_5	i_1, i_2, i_3, i_5	i_3, i_6
s_1	s_1	s_3	s_2	s_3	s_1	s_1	i_2, i_3, i_4	i_6, i_7	i_2, i_7
$s_\emptyset$	$s_\emptyset$	$s_\emptyset$	$s_\emptyset$	$s_\emptyset$	$s_\emptyset$	$s_\emptyset$	i_1	$\cdot$	i_1

Under truthful IA, the first-stage matching is

$$(i_1 - s_\emptyset,\ i_2 - s_3,\ i_3 - s_2,\ i_4 - s_1,\ i_5 - s_1,\ i_6 - s_2,\ i_7 - s_3).$$

The only successful weak appeal is by student i_1 to school s_2. Indeed, i_1 was rejected from s_2, prefers s_2 to being unassigned, and belongs to the same weak priority class at s_2 as the incumbent i_3. Hence i_1's appeal to s_2 is upheld under r_3. No other student

who fails to receive a more preferred school has a successful weak appeal. Thus $\texttt{IA} \circ r_3$ generates exactly one successful appeal.

Now consider truthful DA, computed using the same tie-breaking rule. The first-stage DA matching is

$$(i_1 - s_2,\ i_2 - s_2,\ i_3 - s_\emptyset,\ i_4 - s_3,\ i_5 - s_3,\ i_6 - s_1,\ i_7 - s_1).$$

Under this outcome, there are two successful weak appeals. Student i_3 appeals to s_2, where she belongs to the same weak priority class as the incumbents i_1 and i_2. Hence her appeal is upheld. Student i_5 also appeals to s_2, where she too belongs to the same weak priority class as i_1 and i_2. Hence her appeal is also upheld. No further appeal is needed for the comparison.

Therefore, under truthful reporting, $\texttt{IA} \circ r_3$ generates one successful appeal, whereas $\texttt{DA} \circ r_3$ generates two successful appeals. This proves the claim. □

In the previous example, truthful reporting is a Nash equilibrium of $\texttt{IA} \circ r_3$. Indeed, under truthful reporting and weak appeals, every student except i_1 receives her first-choice school, while i_1 obtains s_2 through a weak appeal. The only school i_1 prefers to s_2 is s_3, but both seats at s_3 are occupied by students with strictly higher weak priority than i_1. Hence i_1 cannot obtain s_3, either directly or by appeal, and no student has a profitable deviation. Thus, the lower number of successful appeals under IA arises at an equilibrium outcome. The reversal theorem below shows that this is not an isolated phenomenon: under weak appeals, IA admits equilibrium outcomes with strong welfare properties relative to DA.

C.5 *The Reversal Theorem*

Theorem 1 shows that, under strict priorities, $\texttt{IA} \circ r_1$ admits a Nash equilibrium whose outcome weakly Pareto-dominates truthful DA. We now extend this result to weak priorities under the weak appeal rule r_3.

THEOREM 3. *Fix weak priorities $(\trianglerighteq_s)_{s\in S}$, and fix a tie-breaking rule τ that induces strict priorities $(\rhd_s^\tau)_{s\in S}$ used by* IA *and* DA*. Under the weak appeal rule r_3, there exists a Nash equilibrium σ^* of $\texttt{IA} \circ r_3$ such that*

$$(\texttt{IA} \circ r_3)(\sigma^*) \succsim \texttt{DA}(\succ).$$

Furthermore, if $(\texttt{IA} \circ r_3)(\sigma^*) \neq \texttt{DA}(\succ)$, then at least one student is strictly better off.

PROOF. Throughout the proof, $\texttt{DA}(\succ)$ denotes the student-proposing DA outcome computed using the tie-broken strict priorities $(\rhd_s^\tau)_{s\in S}$. Under r_3, a student may appeal only to schools that she ranked above her first-stage IA assignment and from which she was rejected in the first stage. An appeal to such a school s is successful if and only if the student has weakly higher priority at s than some first-stage occupant of s, that is, if $i \trianglerighteq_s j$ for some first-stage occupant j of s.

We prove the result using an auxiliary mechanism, denoted $\texttt{MIA}^3$.

Modified IA for weak appeals. The mechanism $\texttt{MIA}^3$ proceeds in synchronous rounds. In round k, every currently unassigned student applies to the k-th school on her true preference list. Admissions are permanent. At each school s, regular seats are filled only while fewer than q_s students have been regularly admitted to s. If s has remaining regular capacity, it admits the highest-priority current applicants according to the tie-broken order $\rhd_s^\tau$, up to the remaining regular capacity. Once q_s students have been regularly admitted to s, no further regular admissions occur. Any further current applicant i is admitted by overbooking if and only if

$$i \trianglerighteq_s j$$

for some regularly admitted student j at s. In that case, s's effective capacity increases by one. Students admitted within the regular capacity are called *regularly admitted*; students admitted through the capacity-expansion rule are called *overbooked.* All other applicants are rejected and proceed to the next school on their list in the following round.

Let $\texttt{MIA}^3(\succ)$ denote the outcome.

Rejection lemma.

LEMMA 5. *If student i is rejected from school s under $\texttt{MIA}^3(\succ)$, then every regularly admitted student at s under $\texttt{MIA}^3(\succ)$ has strictly higher weak priority at s than i. That is, for every regularly admitted student j at s,*

$$j \trianglerighteq_s i \quad \textit{and not} \quad i \trianglerighteq_s j.$$

PROOF. Student i is rejected from s only if all regular seats at s have already been filled and i is not weakly higher priority than any regularly admitted student at s. If there existed a regularly admitted student j at s such that $i \trianglerighteq_s j$, then the overbooking rule would admit i. Therefore, for every regularly admitted student j at s, it is not the case that $i \trianglerighteq_s j$. Since $\trianglerighteq_s$ is complete, this means that j has strictly higher weak priority than i. Regular admittees are never removed, and no further regular admittees can be added after the q_s regular seats are filled. Hence the claim holds for every regular admittee of s under $\texttt{MIA}^3(\succ)$. □

By the definition of overbooking, if student i is overbooked at school s under $\texttt{MIA}^3(\succ)$, then there exists a regularly admitted student j at s such that

$$i \trianglerighteq_s j.$$

$\texttt{MIA}^3$ *weakly Pareto-dominates DA.*

LEMMA 6. *For every school choice problem with weak priorities and fixed tie-breaking,*

$$\texttt{MIA}^3(\succ) \succsim \texttt{DA}(\succ).$$

PROOF. Suppose, toward a contradiction, that some student is strictly better off under $\mathtt{DA}(\succ)$ than under $\mathtt{MIA}^3(\succ)$. Let

$$B = \{i \in I : \mathtt{DA}_i(\succ) \succ_i \mathtt{MIA}^3_i(\succ)\}.$$

By assumption, $B \neq \varnothing$.

For each $i \in B$, let

$$s_i = \mathtt{DA}_i(\succ).$$

Since $s_i \succ_i \mathtt{MIA}^3_i(\succ)$, student i applied to s_i under $\mathtt{MIA}^3$ before receiving her $\mathtt{MIA}^3$-assignment, and was rejected. Let r_i be the round in which i applied to s_i under $\mathtt{MIA}^3$. Let ρ_i be the round in which i receives her $\mathtt{MIA}^3$-assignment, and set $\rho_i = \infty$ if i is unassigned under $\mathtt{MIA}^3$. Then

$$r_i < \rho_i.$$

By Lemma 5, every regularly admitted student at s_i under $\mathtt{MIA}^3$ has strictly higher weak priority at s_i than i. Since i was rejected from s_i, the q_{s_i} regular seats at s_i had already been filled. Let $D(s_i)$ denote the set of these q_{s_i} regular admittees.

Since $i \notin D(s_i)$, $\mathtt{DA}_i(\succ) = s_i$, and school s_i has only q_{s_i} seats under DA, at most $q_{s_i} - 1$ students from $D(s_i)$ can be assigned to s_i under $\mathtt{DA}(\succ)$. Hence there exists some student

$$h \in D(s_i)$$

who is not assigned to s_i under $\mathtt{DA}(\succ)$.

By Lemma 5, h has strictly higher weak priority than i at s_i. Therefore h also has higher tie-broken priority than i at s_i:

$$h \rhd^{\tau}_{s_i} i.$$

Since i is assigned to s_i under DA and DA is stable with respect to the tie-broken priorities, stability implies

$$\mathtt{DA}_h(\succ) \succ_h s_i.$$

Otherwise, (h, s_i) would be a blocking pair of the DA outcome under the tie-broken priorities. But h's $\mathtt{MIA}^3$-assignment is s_i, because h is regularly admitted to s_i under $\mathtt{MIA}^3$. Therefore

$$\mathtt{DA}_h(\succ) \succ_h \mathtt{MIA}^3_h(\succ),$$

so $h \in B$.

Moreover, h was regularly admitted to s_i before or at the round in which i was rejected from s_i. Hence

$$\rho_h \leq r_i < \rho_i.$$

Thus, from any $i \in B$, we can construct another student $h \in B$ with

$$\rho_h < \rho_i.$$

If all students in B had $\rho_i = \infty$, this argument would produce a student in B with finite ρ_h, a contradiction. Hence B contains at least one student with finite ρ_i. Choose $i^* \in B$ with minimal finite ρ_{i^*}. Applying the argument above to i^* produces $h \in B$ with finite $\rho_h < \rho_{i^*}$, contradicting the minimality of ρ_{i^*}. Therefore $B = \varnothing$, and

$$\texttt{MIA}^3(\succ) \succsim \texttt{DA}(\succ).$$

□

Constructing the report profile. Partition students according to their status under $\texttt{MIA}^3(\succ)$:

- D: regularly admitted students;
- O: overbooked students;
- U: unassigned students.

Define σ^* as follows:

- If $i \in D$ and $\texttt{MIA}^3_i(\succ) = s$, then i ranks s first, followed by the remaining schools in her true order.
- If $i \in O$ and $\texttt{MIA}^3_i(\succ) = s$, then i reports truthfully down to s, truncating her list after s.
- If $i \in U$, then i reports truthfully.

Implementation. We show that

$$(\texttt{IA} \circ r_3)(\sigma^*) = \texttt{MIA}^3(\succ).$$

Fix a school s, and let D_s denote the set of students regularly admitted to s under $\texttt{MIA}^3$. Every student in D_s ranks s first under σ^*, so every student in D_s applies to s in round 1 of IA.

We first show that the round-1 admits at s under $\texttt{IA}(\sigma^*)$ are exactly the students in D_s.

Consider any additional student $i \notin D_s$ who also applies to s in round 1 under σ^*. There are three cases.

First, suppose $i \in O$ and $\texttt{MIA}^3_i(\succ) = s$. Then i is assigned to s by overbooking under $\texttt{MIA}^3$. If i ranks s first under σ^*_i, then s is also i's first choice under $\succ_i$. Thus i applied to s in round 1 under $\texttt{MIA}^3$. Since i was not regularly admitted at s, the students in D_s selected for the regular seats at s in round 1 have higher tie-broken priority than i. Hence i is not selected for a regular seat at s in round 1 of IA under σ^*.

Second, suppose $i \in O$, $\texttt{MIA}^3_i(\succ) = t \neq s$, and $s \succ_i t$. Then i was rejected from s under $\texttt{MIA}^3$. By Lemma 5, every student in D_s has strictly higher weak priority at s than i, and therefore higher tie-broken priority than i.

Third, suppose $i \in U$ and s is her first choice. Then i was rejected from s under $\texttt{MIA}^3$. Again, by Lemma 5, every student in D_s has strictly higher weak priority at s than i, and therefore higher tie-broken priority than i.

Thus every student in D_s has higher tie-broken priority at s than every additional round-1 applicant to s under σ^*. Since $|D_s| \leq q_s$, the students in D_s are exactly the round-1 admits at s under $\texttt{IA}(\sigma^*)$.

Now consider $i \in O$, and let $\texttt{MIA}_i^3(\succ) = s$. Under σ_i^*, student i applies to all schools she truly prefers to s, and then to s. At every school $t \succ_i s$, student i was rejected under $\texttt{MIA}^3$. By Lemma 5, every regular admittee at t has strictly higher weak priority than i, and therefore higher tie-broken priority than i. Since the round-1 IA admits at t are exactly the students in D_t, student i is rejected from every such t in the IA first stage.

At s, the regular seats are held by D_s, so i is rejected in the first stage. Her list is truncated at s, so she applies to no further school. Since i is overbooked at s under $\texttt{MIA}^3$, there exists $j \in D_s$ such that

$$i \trianglerighteq_s j.$$

This student j occupies s in the IA first stage. Therefore i's weak appeal to s is upheld under r_3, and i obtains s, exactly as under $\texttt{MIA}^3$.

Finally, consider $i \in U$. Student i was rejected from every school she prefers to being unassigned under $\texttt{MIA}^3$. By Lemma 5, every regular admittee at every such school has strictly higher weak priority than i, and therefore higher tie-broken priority than i. Since the first-stage IA seats at each such school are filled by the students in the corresponding D_t, student i is rejected everywhere in the first stage and has no successful weak appeal. Hence i remains unassigned, as under $\texttt{MIA}^3$.

Therefore

$$(\texttt{IA} \circ r_3)(\sigma^*) = \texttt{MIA}^3(\succ).$$

σ^* is a Nash equilibrium. Fix a student i, and consider any school t such that

$$t \succ_i \texttt{MIA}_i^3(\succ).$$

Student i was rejected from t under $\texttt{MIA}^3$. By Lemma 5, every regular admittee of t has strictly higher weak priority at t than i, and therefore higher tie-broken priority than i. Under σ^*_{-i}, these regular admittees rank t first and fill the first-stage IA seats at t, independently of i's report. Since IA admissions are permanent, i cannot obtain t directly, regardless of when she applies to t.

Nor can i obtain t through a weak appeal. The first-stage occupants of t are precisely the students in D_t, all of whom have strictly higher weak priority at t than i. Thus i is not weakly higher priority than any first-stage occupant of t, so no weak appeal to t succeeds.

Therefore no deviation allows i to obtain a school she strictly prefers to $\texttt{MIA}_i^3(\succ)$. Since i was arbitrary, σ^* is a Nash equilibrium of $\texttt{IA} \circ r_3$.

Combining the implementation result with Lemma 6, we obtain

$$(\texttt{IA} \circ r_3)(\sigma^*) = \texttt{MIA}^3(\succ) \precsim \texttt{DA}(\succ).$$

If $(\texttt{IA} \circ r_3)(\sigma^*) \neq \texttt{DA}(\succ)$, then Pareto dominance together with strict preferences implies that at least one student is strictly better off. This proves the theorem. □

SUPPLEMENTARY APPENDIX D: ADDITIONAL LABORATORY RESULTS.

This appendix reports a small set of robustness and supplementary outcomes in the same order as the main text. Participant-level analyses use the paper sample of 10,195 participant-round observations and 59 independent matching groups. The blocking-pair analysis uses the 1,435 complete seven-person market-rounds constructed from the same sample.

D.1 *First-choice truth-telling*

The main analysis defines truth-telling as submitting the complete induced preference ranking. A less demanding measure asks only whether the participant ranks her induced first-choice school first. Figure D.1 and Table D.1 repeat the main truth-telling analysis using this measure.

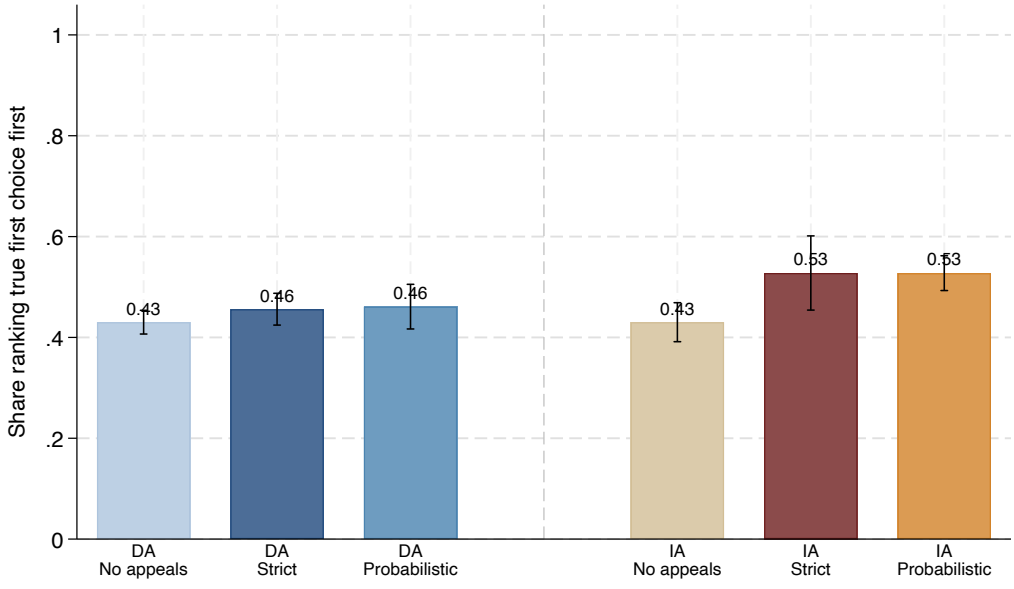


FIGURE D.1. First-choice truth-telling by treatment.

Bars show the share of participants who rank their induced first-choice school first. Whiskers are 95% confidence intervals clustered at the matching-group level.

The pattern is the same as for complete truth-telling. Under IA, first-choice truth-telling rises from 0.43 without appeals to 0.53 under both strict and probabilistic appeals. Under DA, it rises only to 0.46 under either rule. In the fully controlled specification, strict and probabilistic appeals increase first-choice truth-telling under IA by approximately 10 percentage points. The corresponding DA effects are approximately 3 percentage points and are not statistically significant.

Evolution over rounds. Figure D.2 reports role-balanced first-choice truth-telling in each round. From round 2 onward, both IA appeal treatments remain above IA without

TABLE D.1. Impact of appeals on first-choice truth-telling

	Specifications			
	(1.00)	(2.00)	(3.00)	(4.00)
Panel A. IA treatments				
IA strict	0.10**	0.10**	0.10**	0.10**
	(0.04)	(0.04)	(0.04)	(0.04)
IA probabilistic	0.10***	0.10***	0.10***	0.10***
	(0.03)	(0.03)	(0.03)	(0.03)
Panel B. DA treatments				
DA strict	0.03	0.03	0.03	0.03
	(0.02)	(0.02)	(0.02)	(0.02)
DA probabilistic	0.03	0.03	0.03	0.03
	(0.02)	(0.02)	(0.02)	(0.03)
Round effects	No	Yes	Yes	Yes
Participant-type effects	No	No	Yes	Yes
Demographics + risk att.	No	No	No	Yes
Observations	10,195	10,195	10,195	10,185
Matching groups	59	59	59	59

Notes. $^{***}p < 0.01$, $^{**}p < 0.05$, $^{*}p < 0.100$. Entries are average marginal effects from logit regressions. Standard errors are clustered at the matching-group level. First-choice truth-telling equals one when the participant ranks her induced first-choice school first. The baseline in each panel is the corresponding mechanism without appeals.

appeals in every round. The DA series overlap and display no persistent ordering. The treatment pattern therefore does not arise from a short-lived response in the opening rounds.

D.2 *Efficiency in experimental points*

The main text measures efficiency using assigned rank. Because experimental points are an affine transformation of true-preference rank, the points measure yields the same descriptive ordering. Figure D.3 reports mean points by treatment.

Under IA, mean points rise from 8.36 without appeals to 9.70 under strict appeals and 9.80 under probabilistic appeals. Strict appeals leave DA essentially unchanged, at 6.29 without appeals and 6.31 with strict appeals, while probabilistic appeals raise mean points to 6.97.

Unassigned participants. Figure D.4 shows that part of IA's efficiency gain occurs on the extensive margin. The unassigned share falls from 0.14 without appeals to 0.07 under strict appeals and 0.06 under probabilistic appeals. Under DA, strict appeals leave the unassigned share at 0.14, while probabilistic appeals reduce it to 0.12.

Table D.2 repeats the regression specifications from the rank analysis using experimental points. All specifications use the 10,195-observation paper sample; the specification with personal controls uses the 10,185 observations for which those controls are

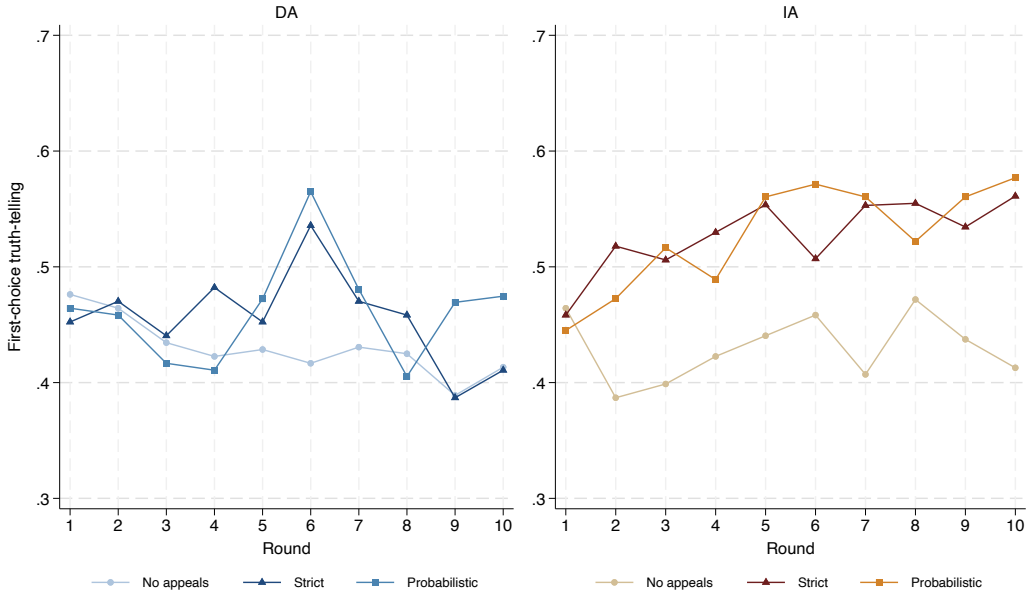


FIGURE D.2. First-choice truth-telling over rounds.
Each observation is the mean first-choice truth-telling rate in a treatment-round, giving equal weight to each of the seven participant types.

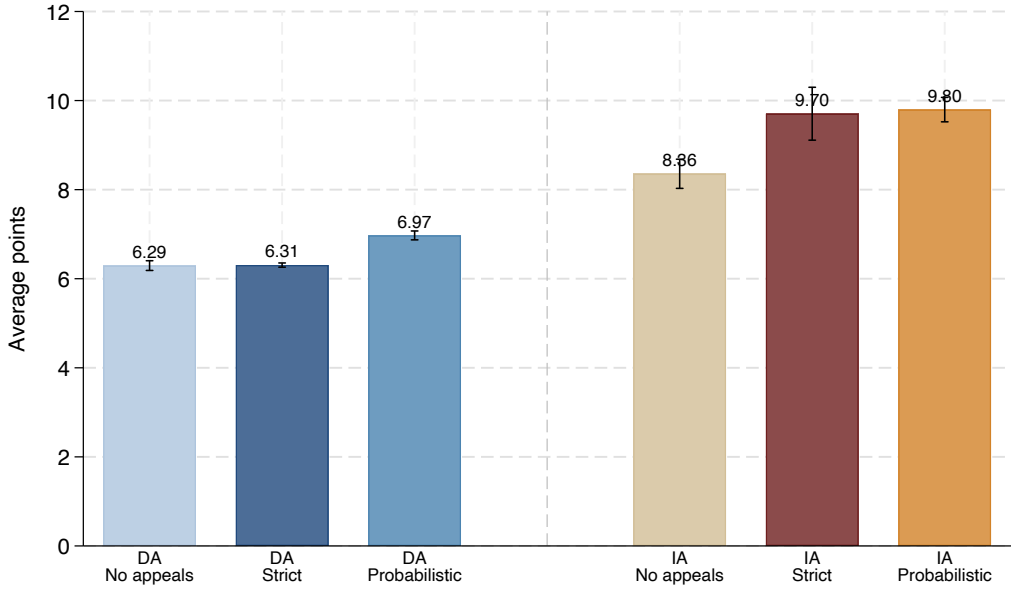


FIGURE D.3. Experimental points by treatment.
Bars report mean points from the final assignment. Higher values indicate better assignments. Whiskers are 95% confidence intervals clustered at the matching-group level.

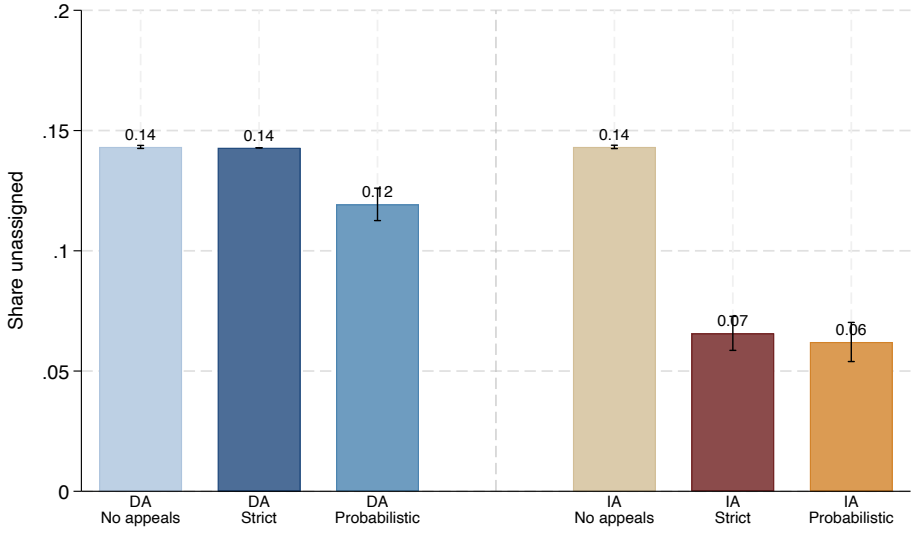


FIGURE D.4. Share of participants unassigned by treatment.
Whiskers are 95% confidence intervals clustered at the matching-group level.

complete. The estimates confirm the descriptive pattern. In the fully controlled specification, strict appeals raise IA welfare by 1.36 points and leave DA unchanged. Probabilistic appeals raise welfare by 1.46 points under IA and 0.67 points under DA. Strict and probabilistic appeals therefore widen IA's points advantage by 1.36 and 0.79 points, respectively.

D.3 *Strict blocking pairs*

The main text reports the share of stable market-rounds and the corresponding regression estimates. Figure D.5 provides the continuous counterpart by reporting the mean number of strict blocking pairs. No additional regression analysis is needed because Table 6 already reports treatment effects on this outcome.

Strict appeals reduce IA's mean number of strict blocking pairs by more than half under both preference benchmarks, while leaving DA essentially unchanged. Probabilistic appeals also reduce blocking pairs under IA, but create priority violations after DA. This is the same qualitative pattern documented by the stable-share measure in the main text.

TABLE D.2. Impact of appeals on experimental points

	Specifications			
	(1.00)	(2.00)	(3.00)	(4.00)
Panel A. IA treatments				
IA strict	1.35***	1.35***	1.34***	1.36***
	(0.33)	(0.33)	(0.33)	(0.32)
IA probabilistic	1.44***	1.44***	1.44***	1.46***
	(0.21)	(0.21)	(0.21)	(0.21)
Panel B. DA treatments				
DA strict	0.01	0.01	0.01	0.00
	(0.06)	(0.06)	(0.06)	(0.07)
DA probabilistic	0.68***	0.68***	0.68***	0.67***
	(0.07)	(0.07)	(0.07)	(0.08)
Panel C. Increase in IA's points advantage relative to no appeals				
Strict appeals	1.34***	1.34***	1.33***	1.36***
	(0.33)	(0.33)	(0.33)	(0.33)
Probabilistic appeals	0.76***	0.76***	0.76***	0.79***
	(0.22)	(0.22)	(0.22)	(0.22)
Round effects	No	Yes	Yes	Yes
Participant-type effects	No	No	Yes	Yes
Demographics + risk att.	No	No	No	Yes
Observations	10,195	10,195	10,195	10,185
Matching groups	59	59	59	59

Notes. $^{***}p < 0.01$, $^{**}p < 0.05$, $^{*}p < 0.100$. Entries are adjusted differences in experimental points from linear regressions. Standard errors are clustered at the matching-group level. The baselines in Panels A and B are the corresponding mechanisms without appeals. Panel C reports (IA appeal − DA appeal) − (IA no appeals − DA no appeals), so positive values indicate that appeals widen IA's points advantage. Demographics + risk att. are age, gender, and risk aversion.

SUPPLEMENTARY APPENDIX E: FULL APPEAL TAKE-UP.

The theoretical analysis assumes that appeals are costless and therefore that each student follows the weakly dominant appeal strategy

$$a_i^{\mathrm{FT}} = \{s \in A_i(\hat{\succ}, M) : s \succ_i M_i(\hat{\succ})\}.$$

We assess this assumption by replacing observed appeal choices with a_i^{FT}, while holding submitted rank-order lists and first-stage assignments fixed. Under the strict rule, every strict-priority claim in this set succeeds. Under the probabilistic rule, strict-priority claims succeed with certainty and every remaining claim succeeds independently with probability 0.1. If several appeals succeed, the student receives her most-preferred successful school according to her induced true preferences, as in the model. We integrate

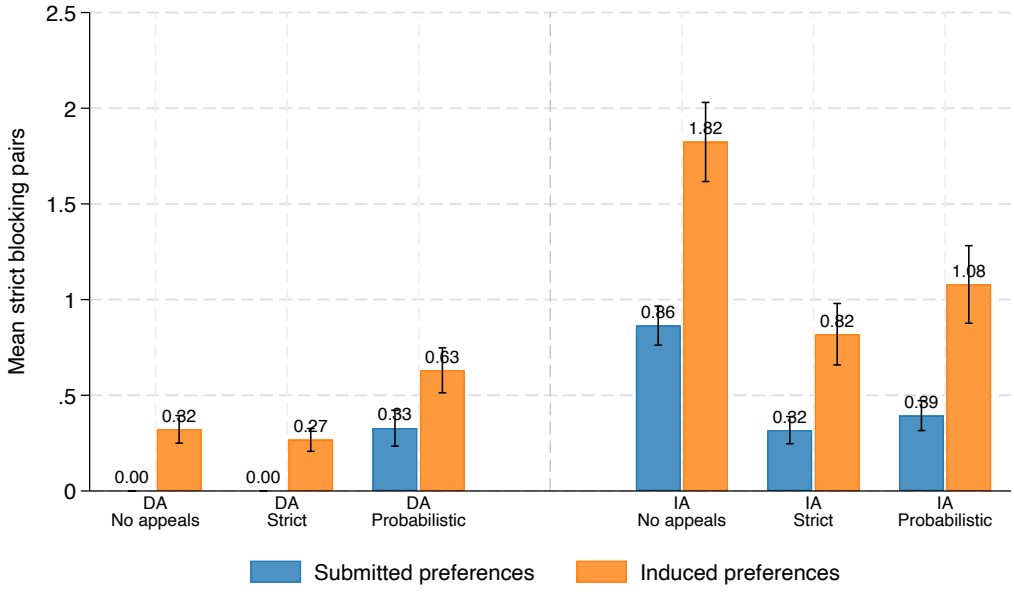


FIGURE D.5. Strict blocking pairs by treatment.

Bars report the mean number of strict blocking pairs per market-round, evaluated against submitted and induced true preferences. Equal-priority tie-break claims are excluded. Treatment means give equal weight to each independent matching group. Whiskers are 95% confidence intervals across matching groups.

exactly over all probabilistic appeal outcomes. This exercise is a conditional theoretical benchmark rather than an additional experimental treatment: in particular, it does not allow initial reports to respond to automatic appeal filing.

Full take-up sharpens the asymmetry under the strict rule. DA remains exactly unchanged. Under IA, average assigned rank improves from 2.059 to 2.010, strict blocking pairs fall from 0.328 to 0.043 per market-round, and procedural stability rises from 0.719 to 0.935. The share of market-rounds that strictly Pareto-dominate truthful DA rises from 0.477 to 0.600. By construction, no appealable strict blocking pair involving a school that the student genuinely prefers remains, as predicted by Lemma 1. The outcome need not be fully procedurally stable because that empirical measure is evaluated against submitted preferences and also requires non-wastefulness.

The probabilistic benchmark reinforces the mechanism-specific comparison. Full take-up further improves IA's average rank and procedural stability. Under DA, it modestly improves average rank but increases strict blocking pairs and reduces procedural stability because additional random admissions can disturb an initially stable assignment and leave first-stage seats vacant. Full-take-up counterparts of every main-text table and figure are produced by the replication code attached.

TABLE E.1. Observed outcomes and the theoretical full-take-up benchmark

	DA		IA	
	Observed	Full take-up	Observed	Full take-up
Panel A. Strict appeals				
Average assigned rank	2.739	2.739	2.059	2.010
Strict blocking pairs	0.000	0.000	0.328	0.043
Procedurally stable share	1.000	1.000	0.719	0.935
Strictly dominates truthful DA	0.000	0.000	0.477	0.600
Panel B. Probabilistic appeals				
Average assigned rank	2.606	2.551	2.041	1.940
Strict blocking pairs	0.307	0.386	0.396	0.142
Procedurally stable share	0.648	0.613	0.606	0.843
Strictly dominates truthful DA	0.346	0.414	0.458	0.630

Notes. Lower average rank is better. Reports and first-stage assignments are held fixed. Full take-up replaces observed appeal choices with every appealable claim to a school genuinely preferred to the first-stage assignment and assigns the most-preferred successful school when several appeals are upheld. Average rank is measured at the participant-round level. Strict blocking pairs and Pareto dominance are averaged across complete seven-person market-rounds. Procedural stability is evaluated against submitted preferences and averaged first within matching group and then across matching groups, as in Figure 4. Probabilistic outcomes are integrated exactly. Treatments without appeals are omitted because the counterfactual leaves them unchanged.

SUPPLEMENTARY APPENDIX F: SURVEY OF APPEAL PANEL MEMBERS.

This appendix reports a survey conducted with members of school admission appeal panels in England. The purpose of the survey was to understand how experienced panel members evaluate different types of appeal claims, and we conducted it to understand how to best model appeal rules in our theoretical model that were consistent not only with the regulations, but with how appeals are upheld in practice.

The survey was implemented online using Qualtrics. We contacted all local authorities in England and asked them to circulate the survey to individuals involved in school admission appeals. Only a limited number of local authorities agreed to circulate the survey or replied positively to our request. The local authorities that circulated the survey were Merton, Barnsley, Barnet, Darlington, Hartlepool, Kirklees, and Wokingham. Bradford expressed possible interest, but had not circulated the survey at the time of writing. Participation was voluntary and anonymous, and the survey was approved through the ethics process at City, University of London.

The survey asked respondents to evaluate a set of hypothetical appeal cases. In each case, respondents were presented with a short description of a child who had been refused admission to a school and whose parents had appealed the decision. The cases varied the reason for the appeal. One case involved a clear priority or admissions-rule violation. A second case involved a medical reason, where the appealed school was close to a hospital and the child had a serious health condition. A third case involved a comparatively weak reason, namely that the school offered an after-school Italian course. A

fourth case involved equal priority and tie-breaking. A fifth case asked respondents how many additional students they would be comfortable admitting when many rejected students appeal to a school already at its published admissions number.

After each likelihood scenario, respondents were asked whether the appeal was likely to be successful. The response categories were *Very unlikely*, *Unlikely*, *Hard to say*, *Likely*, and *Very likely*. These questions appeared in random order. Respondents were also invited to explain their answer in an open text box. The survey also included two broader questions. First, respondents were asked how often appeal panel decisions are unanimous. Second, they were asked an open-ended question about the possible reasons why an appeal may be successful. These questions were included to understand whether panel members view appeal decisions as deterministic or as involving substantial discretion and case-by-case judgement.

Table F.1 reports the distribution of responses across the four main hypothetical cases. Missing responses are reported separately.

TABLE F.1. Perceived likelihood that an appeal succeeds

	Priority violation	Medical reason	Trivial reason	Tie/equal-priority case
Very unlikely	2	1	28	9
Unlikely	3	7	32	16
Hard to say	37	46	20	42
Likely	16	20	2	11
Very likely	17	2	1	4
No answer	19	18	11	12
Total	94	94	94	94

Several patterns are worth noting. First, respondents distinguished sharply between different reasons for appeal. The weak or idiosyncratic reason was generally regarded as unlikely to succeed: 60 out of 94 respondents chose either *Very unlikely* or *Unlikely*. By contrast, the priority-violation case attracted substantially more favourable responses. Only 5 respondents chose *Very unlikely* or *Unlikely*, while 33 chose *Likely* or *Very likely*. Equivalently, among respondents who gave a directional answer rather than selecting *Hard to say*, 33 out of 38 classified the priority-violation appeal as likely or very likely to succeed. This supports our focus on the strict appeal rule r_1 in the main text: experienced panel members broadly agree that appeals are especially compelling when the initial rejection reflects a priority or admissions-rule violation.

Second, the medical-reason case generated considerable uncertainty. Although 22 respondents classified the appeal as *Likely* or *Very likely*, 46 selected *Hard to say*. This is consistent with the idea that some appeals are not resolved by a simple priority comparison. Instead, panel members may need to weigh the strength of the family's circumstances against the prejudice to the school from admitting an additional pupil. This motivates the probabilistic appeal rule in the model: some claims may be institutionally admissible but difficult to predict ex ante.

Third, the tie or equal-priority case also generated substantial uncertainty. A large fraction of respondents selected *Hard to say*, while the remaining responses were spread across both positive and negative categories. This suggests that cases involving equal priority, tie-breaking, or borderline admissibility may be treated differently across panels or local contexts. This evidence motivates our distinction between strict appeals, weak appeals, and probabilistic appeals.

The survey also asked respondents how often panel decisions are unanimous. The answers are reported in Table F.2. Most respondents reported that unanimous decisions occur often, but not always. This is further evidence that appeal decisions involve judgement and that similar cases need not be assessed identically by all panel members.

TABLE F.2. Frequency of unanimous panel decisions

Response	Number of respondents
Sometimes	17
Often	60
Always	4
No answer	13
Total	94

The capacity question provides similar evidence of discretion. Respondents were asked to consider a secondary school with a published admissions number of 100 that had already accepted 100 students, while another 100 rejected students were appealing. They were then asked for the maximum number of additional students they would be comfortable admitting. Responses varied substantially: many respondents answered zero, while others were willing to admit a positive number of additional students. This reinforces the view that appeal panels may differ in how they weigh the prejudice to the school against the claims of appealing families.

The open-ended responses provide additional context. Respondents listed several reasons why appeals may be successful, including mistakes in applying oversubscription criteria, admissions arrangements that were not lawful or not correctly applied, unreasonable decisions, medical or social needs, special educational needs, sibling or twin considerations, recent changes in family circumstances, the lack of a reasonable alternative school, and cases where the prejudice to the child or family was judged to outweigh the prejudice to the school. These answers are consistent with the institutional description in the main text: some appeals operate as corrections of procedural or priority errors, while others involve a balancing exercise.

The survey has limitations. Although we contacted all local authorities in England, only a small subset agreed to circulate the survey or replied positively. The resulting coverage is therefore limited, and the respondents should not be interpreted as representative of all admission appeal panel members in England. Respondents also self-selected into participation, and we do not observe the full population of panel members who received the survey invitation. Finally, the hypothetical cases are necessarily stylized and

cannot capture all legal and factual details that would be available in an actual appeal hearing. For these reasons, we use the survey only as qualitative and descriptive support for the modeling assumptions, not as a basis for estimating appeal success probabilities.

Despite these limitations, the survey is useful for the purposes of this paper. It confirms that experienced panel members distinguish sharply between different grounds for appeal; that priority or procedural violations are viewed as especially important; and that medical, welfare, tie-breaking, and balancing cases generate substantial uncertainty. This supports the modelling choice to study more than one appeal rule. The strict rule captures appeals that correct priority or procedural violations. The weak rule captures cases in which equal-priority or tie-breaking claims may be recognised. The probabilistic rule captures the discretionary and heterogeneous component of appeal adjudication.

F.1 *Survey Verbatim Text*

Consent screen. You are invited to take part in a research survey being conducted by Dr Claudia Cerrone, Dr Yoan Hermstrüwer and Dr Josuè Ortega.

The purpose of this survey is to understand school admissions appeals in England. The survey has received ethics approval from City, University of London. Your participation in the survey is voluntary and the information you provide is anonymous.

The data collected from this survey may be used for scientific publication. By continuing to the questions, you consent to participate in this survey.

I continue to the questions

Question 1. Livia and Matty have applied for the same primary school (non-infant class) and are both being evaluated on the basis of the admission criterion of distance to the school. Livia lives 975 meters from the school and is rejected, while Matty lives 950 meters from the school and is admitted.

Livia's parents have decided to appeal the school's decision to reject Livia. They argue that Livia should be admitted to that school as it offers an after school Italian course and they are very keen on Livia learning Italian. Note that this appeal is for a place in Year 5 and therefore not covered by infant class size legislation.

Do you think that this appeal is likely to be successful?

Very unlikely; Unlikely; Hard to say; Likely; Very likely

If you want to explain your decision, please use the box below.

Question 2. Charlie and Dora have applied for the same primary school (non-infant class) and are both being evaluated on the basis of the admission criterion of distance to the school. Charlie lives 975 meters from the school and is rejected, while Dora lives 950 meters from the school and is admitted.

Charlie's parents have decided to appeal the school's decision to reject Charlie. They argue that Charlie should be admitted to that school as it is close to a hospital and Charlie has a life-threatening heart defect. Note that this appeal is for a place in Year 5 and therefore not covered by infant class size legislation.

Do you think that this appeal is likely to be successful?

Very unlikely; Unlikely; Hard to say; Likely; Very likely

If you want to explain your decision, please use the box below.

Question 3. Appeals are usually successful if school admission arrangements are unlawful or have not been applied correctly. Apart from this, what are possible reasons why an appeal might be successful? Give as many reasons as you want, and please separate them by a comma.

Question 4. Based on your experience, how often are panels' decisions unanimous?
Never; Rarely; Sometimes; Often; Always

Question 5. A selective secondary school admits students based on their grades. The school admits all students with an average grade above 80 and reserves two spots for students with an average grade of exactly 80. Three applicants had an average grade of 80, so a tie-breaker was used to allocate the two available places.

The parents of the student who was not selected are appealing the school's decision to reject their child, arguing that it is unfair for their child to be rejected purely due to random chance.

Is their appeal likely to succeed?
Very unlikely; Unlikely; Hard to say; Likely; Very likely
If you want to explain your decision, please use the box below.

Question 6. Ana and Bob have applied for the same primary school and are both being evaluated on the basis of the admission criterion of distance to the school. Ana lives 925 meters from the school and is rejected, while Bob lives 1,050 meters from the school and is admitted.

Ana's parents have decided to appeal the school's decision to reject Ana. They argue that Ana should be admitted to the school as the admissions criteria have not been applied correctly. Note that this appeal is for a Reception place under infant class size legislation.

Do you think that this appeal is likely to be successful?
Very unlikely; Unlikely; Hard to say; Likely; Very likely
If you want to explain your decision, please use the box below.

Question 7. A secondary school has a Published Admissions Number of 100, which was calculated in compliance with admissions law. While 100 students have been accepted, another 100 were rejected. All the parents of the rejected children are appealing, arguing that their children would benefit from the school's unique mathematics program. The school argues that it has no extra resources to accommodate more students, and that each additional student admitted would lead to larger classes, and consequently less resources per pupil.

You must decide how many appeals will be successful.

What is the maximum number of additional students you would be comfortable admitting to the school?

End of survey.