False but Harmless?
Misreporting under Strategy-Proof Matching

Yuxing Liang
Concordia University

October 2026

Abstract

Under strategy-proof matching mechanisms, no student can gain by misreporting her ranking of the schools, yet many students submit non-truthful rankings. A misreport is harmless if it gives the student the same school as the truth, whatever the other students report; such a ranking is a dominant report, just as the truth is. This paper asks which misreports are harmless when every student must rank every school. Consider Deferred Acceptance, Top Trading Cycles, and every other strategy-proof mechanism that respects top-priority sets, is top-top consistent, and grants feasible first choices. Under each of them, every student has a guarantee rank, which depends only on her own ranking, the priorities, and the capacities. Reordering the schools below this rank never changes her school. With at least one more seat than students, the harmless misreports are exactly those that reorder only the schools below the guarantee rank, in every priority profile. A dominant report then reveals the student's true ranking down to that rank, and nothing below it. Without a spare seat, some priority profiles allow other harmless misreports. The results sort non-truthful rankings by what they do to the student's assignment, which complements behavioral explanations of why students submit them. In the Chen–Sönmez experiment, 38 of 58 non-truthful rankings kept the realized assignment, but only 6 are harmless whatever the others report; the other 52 could have cost the student a better school.

Keywords: strategy-proofness; deferred acceptance; top trading cycles; serial dictatorship; dominant reports; school choice; option sets; partial identification.
JEL classification: C78; D47.

1 Introduction

Strategy-proof mechanisms such as Deferred Acceptance (DA), Top Trading Cycles (TTC), Serial Dictatorship (SD), and Random Serial Dictatorship give students simple advice: report your true ranking of the schools. Under these mechanisms, no student can ever do better by submitting a different ranking. Yet non-truthful rankings are common, both in laboratory experiments and in high-stakes assignment systems (Chen and Sönmez, 2006; Li, 2017; Hassidim et al., 2017, 2021; Rees-Jones, 2018; Chen et al., 2024; Kloosterman and Troyan, 2026). This paper starts from a simple observation about such rankings. Strategy-proofness guarantees that the true ranking is a best report. It does not guarantee that it is the only best report. Other rankings can be just as good, in the sense that they give the student exactly the same school as the truth, whatever the other students report. A student who submits such a ranking misreports, but she has lost nothing.

Strategy-proofness implies that truth-telling is a weakly dominant strategy, and a weakly dominant strategy is used under two conventions. Under the stronger one, a weakly dominant strategy must also do strictly better than each other strategy against some strategies of the others, so at most one strategy can be weakly dominant. This paper uses the weaker convention. A ranking is a *dominant report* if it does at least as well as every other ranking against every combination of rankings that the other students can submit. Nothing strict is required, so there can exist several dominant reports at once, and the truth need not be unique.

The title of the paper asks whether a false ranking can be harmless, and a misreport can be harmless in two quite different senses. In the robust sense, a ranking gives the student the same school as the truth against every combination of rankings that the other students could submit. Such a ranking is a dominant report, just as the truth is. In the realized sense, a ranking gives the student the same school as the truth only against the rankings that the other students actually submitted. Against other rankings, she might get a worse school. The paper asks a question that comes before any explanation of why students misreport, that is, which misreports are robustly harmless, so that the mechanism itself gives the student no reason to avoid them?

The main contribution is one theorem that answers this question when every student must rank every school. The theorem covers DA, TTC, and every other strategy-proof mechanism with three further properties including respect for top-priority sets, top-top consistency, and feasible first choices. In every market, a ranking is dominant if it lists the student's favorite schools first, in their true order, down to a certain rank. When there is at least one more seat than students, these rankings are exactly her dominant reports in every priority profile. Without a spare seat, dominant reports might not be unique. The theorem also has three other uses beyond the theory. For example, it tells which misreports can cost the student a better school. It shows what a submitted ranking reveals about her preferences. And it shows how an algorithm can ask for less than a complete ranking. A small example shows the main idea.

Example 1.1 (One spare seat)**.** Three students, s_1, s_2, and s_3, apply to three schools, c_1, c_2, and c_3, and each school has one seat. Consider student s_3, whose true ranking is $c_1 \succ_3 c_2 \succ_3 c_3$: she likes c_1 best and c_3 least. Each school ranks the three students by priority. The table lists these priority orders, from highest to lowest; the column headed $\succ_{c_1}$ gives the order of c_1, and so on:

$\succ_{c_1}$	$\succ_{c_2}$	$\succ_{c_3}$
s_1	s_1	s_1
s_2	s_2	s_3
s_3	s_3	s_2

Student s_1 has the highest priority everywhere. Student s_3 has the lowest priority at c_1 and c_2, but at c_3 she comes before s_2.

Which schools can s_3 obtain? Fix any rankings of the other two students and, for a moment, leave s_3 out. Under DA, TTC and their variants, s_1 then receives the school she ranks first, because nobody has higher priority than she does anywhere. Student s_2 receives the school she prefers among the two that remain, and the third school is left with an empty seat. Now bring s_3 back. She can always obtain the school with the empty seat as she only has to rank it first. She can also obtain c_3 whenever s_2 sits there, because c_3 gives her priority over s_2. She can obtain nothing else since at c_1 and c_2, both other students come before her, and at c_3 so does s_1. The set of schools that a student can obtain by changing her own ranking, while the rankings of the others are fixed, is called her *menu*. The table below lists the six ways in which s_1 and s_2 can be seated when s_3 is left out, and the menu that each of them leaves to s_3.

Seat of s_1	Seat of s_2	Empty seat	Menu of s_3
c_1	c_2	c_3	$\{c_3\}$
c_1	c_3	c_2	$\{c_2, c_3\}$
c_2	c_1	c_3	$\{c_3\}$
c_2	c_3	c_1	$\{c_1, c_3\}$
c_3	c_1	c_2	$\{c_2\}$
c_3	c_2	c_1	$\{c_1\}$

Every school appears in some menus, so each of the three schools is one that s_3 can sometimes obtain. But c_1 and c_2 never appear in the same menu. Either of them can be in her menu only through the empty seat, and there is only one empty seat.

This has an important result. Whatever the others submit, s_3 receives the school in her menu that her own submitted ranking puts highest. The reason is explained after the example. Because c_1 and c_2 are never in her menu together, the order in which she ranks them never affects her assignment. In particular, the false ranking $c_2 \succ c_1 \succ c_3$ gives her exactly the same school as the truth, whatever the others report. By contrast, the rank of c_3 does matter. When her menu is $\{c_1, c_3\}$, a ranking that puts c_3 above c_1 gives her c_3 instead of c_1. Against every possible ranking of the others, a dominant report must give the same school as the truth. Student s_3 therefore has exactly two dominant reports: the truth, and the ranking that switches the order of c_1 and c_2. Both rankings put c_3 in the last rank.

Now give c_3 a second seat, so that there is one more seat than students. The other two students can now leave c_1 and c_2 open at the same time. For example, if s_1 and s_2 both rank c_3 first, they both sit at c_3, and s_3 can obtain any of the three schools by ranking it first. Facing this menu, she gets whichever of c_1 and c_2 she ranks higher, so reversing their order is no longer harmless. With the second seat, the truth is her only dominant report.

In general, a spare seat does not make the truth the only dominant report. It gives the dominant reports a simple form, which Theorem 1 describes. The student must list her favorite schools in their true order down to a certain rank, and she may order the remaining schools in any way. Therefore, she can have many dominant reports. In the example, that rank is 3, which is the last rank, so nothing is left to reorder. The first case of Example 4.1, in Section 4.1, has a spare seat and two dominant reports.

The example lies on an idea from social choice. In his direct proof of the Gibbard–Satterthwaite theorem, Barberà (1983) fixed the reports of all other agents and looked at the set of outcomes that one agent can obtain by changing her own report. He called this set her option set. Under a strategy-proof rule, the agent must receive her favorite outcome in her option set, since otherwise some other report can give her something better. The menus above are option sets, so under a strategy-proof mechanism a student receives her favorite school in her menu. The same is true of

any ranking she might submit. Every ranking is a possible true preference, and strategy-proofness must hold for a student whose true preference it is. So, against fixed rankings of the others, every ranking gives her the school in her menu that it ranks highest. Therefore, a ranking does as well as the truth against every ranking of the others exactly when it picks the same school as the truth from every menu she has.

A simple observation follows. Her ranking of two schools can matter only if, for some rankings of the other students, her menu contains both schools and no school that she prefers to both. Only then can the order in which she ranks the two schools change what she gets. Such a pair of schools is called a *decisive pair*. A ranking is dominant exactly when it keeps the true order of every decisive pair (Proposition 2.5). What it does with the other pairs never matters. In Example 1.1, the decisive pairs are c_1 with c_3 and c_2 with c_3, but not c_1 with c_2. Decisive pairs resemble Barberà's pivotal voters, who can switch an outcome between two alternatives by changing their vote. Section 2.4 discusses the connection.

This turns the question of the paper into a question about markets. Under DA and TTC, which pairs of schools can the other students put in a student's menu together, through the rankings they submit, and so which comparisons must every dominant report keep? The answer is not obvious, because the other students shape a student's menu only indirectly. Under DA, their rankings start chains of rejections, and these chains determine which schools fill up. Under TTC, students trade the priorities that they hold at different schools. Example 1.1 shows the kind of subtlety involved: each of c_1 and c_2 can be in the menu of s_3, but never both at once. Knowing which schools a student can possibly obtain is therefore not enough. What matters is which schools she can obtain against the same rankings of the others.

The main result, Theorem 1, answers this question for every market with a spare seat, when every student must rank every school. It is built on a guarantee that every student has. Go down a student's true ranking one school at a time. At each step, ask whether the students who outrank her at those schools could take every seat there. Each of them can take only one seat. The first point at which they no longer can is her *guarantee rank*. Whatever the others submit, at least one of her schools down to this rank stays available to her. Under DA and TTC, the guarantee rank is exactly the rank of the worst school she can receive by reporting truthfully. The guarantee rank depends only on the student's own ranking, the priorities, and the capacities. As it does not depend on what the others submit, it can be computed before anyone submits a ranking. In Example 1.1, the other two students can fill the seats of c_1 and c_2, but they cannot fill all three seats, so the guarantee rank of s_3 is 3.

Theorem 1 has three parts. First, in every market, every ranking that is truthful through the guarantee rank is dominant. Such a ranking lists the student's favorite schools down to that rank first, in their true order, and lists the remaining schools in any order. Second, if there is at least one more seat than students, these are exactly the dominant reports. This holds in every market, that is, whatever priorities the schools use. A student with m schools and guaranteed rank k then has $(m-k)!$ dominant reports, one for each order of the schools below rank k. The truth is her only dominant report exactly when k is m or $m-1$. The dominant reports are also the same under DA, TTC, and every other mechanism with the three properties described below, although these mechanisms can give her different schools and different menus. Among these mechanisms, which comparisons matter is a property of the market, not of the mechanism. Third, without a spare seat, a student can have more dominant reports than these. Under DA and TTC, this happens for every capacity vector with at least three schools and no more seats than students. In some markets, a student's dominant reports are exactly the rankings that put her least favorite school last. She is then free to order all her other schools in any way. Example 1.1 is the smallest market of this kind.

Why does one spare seat make such a difference? For a student's comparison of two schools to

matter, the other students must submit rankings under which three things happen at once. First, students who have higher priority than she does fill every school that she prefers. Second, each of the two schools still has a seat that she can take. Third, every other student still gets a seat somewhere. The third requirement is the one that can fail. Because everyone must rank every school, nobody can be kept out of the market: as long as there are at least as many seats as students, every other student ends up with a seat. With a spare seat, there are always enough seats to meet all three requirements at once. The students with higher priority fill the schools that she prefers to both; one seat is left vacant at each of the two schools, and everyone else takes one of the remaining seats. A count in Section 4.6 shows that S terms cancel, which gives $Q \geq n + 1$. Therefore, the condition depends only on the total number of seats. One spare seat anywhere in the market is enough, even at a school that every student ranks last. Without a spare seat, once every other student has a seat, at most one seat is vacant. Under DA, the second school can then be in her menu only if one of its seats is held by a student whom she outranks there. The market may have no such student to spare. This is what happens in Example 1.1: s_3 outranks nobody at c_1 or c_2, so one of them can be in her menu only through the single empty seat. For balanced markets, Section 5 turns this observation into an exact description of the decisive pairs under DA.

The same reasoning explains why DA and TTC have the same dominant reports when there is a spare seat, even though they treat priorities very differently: DA respects priorities, while TTC lets students trade them. The argument uses only a few properties of the mechanism, so Theorem 1 is stated and proved for every strategy-proof mechanism with three further properties that DA and TTC share (Section 4.2). Two of these properties come from Decerf et al. (2024). Respect for top-priority sets says that if the students with higher priority cannot fill all the seats of the first few schools in a student's ranking, she receives one of these schools. Top-top consistency says that a student who ranks a school at which she has the highest priority in the first rank can be removed together with one seat of that school, and it will not change anyone else's assignment. The third, feasible first choices, requires the mechanism to give every student her first choice whenever these first choices fit into the available seats. In the proof, each property has a specific role. Respect for top-priority sets gives the student her guarantee. With top-top consistency, it also ensures that the rivals who fill the schools that she prefers to both schools of a pair keep those seats, whatever she reports. Feasible first choices then ensure that she can obtain either school of the pair by ranking it first. All three parts of the theorem hold for every mechanism with these properties. In particular, the boundary holds mechanism by mechanism: whichever of these mechanisms is used, some market without a spare seat gives a student dominant reports that no single rank describes.

The requirement that every student rank every school is what creates this boundary. When students may submit shorter lists, Kumano and Watabe (2011) show that every dominant DA report must order truthfully all the schools that the student prefers to her assignment. Combined with the guarantee above, their result shows what happens when short lists are allowed. With or without a spare seat, the dominant reports are then the rankings that are truthful through the guarantee rank in every market. The safe sets of Decerf and Van der Linden (2021), defined for constrained lists, express a related idea. Their argument works by letting the other students leave the market, which they do by listing one school or none. When every student must rank every school, nobody can leave the market. A spare seat then serves the purpose that leaving serves when lists can be short: with it, every other student can have a seat while two schools stay open to the student.

The difference between the two senses of harmless misreporting matters for how data are read. The characterization is easy to compute: it requires only bipartite matching problems built from the student's ranking, the priorities, and the capacities. It can therefore be applied directly to experimental data. Section 6 applies it to the DA treatment of Chen and Sönmez (2006). That market has exactly as many seats as students, so it illustrates the case without a spare seat. Of the

144 submitted rankings, 58 differ from the student's true ranking. Of these 58, 38 gave the student the same school that the truth would have given her, against the rankings that the others actually submitted. Only 6 of these 38 are dominant. The other 32 kept her school only because of what the others happened to submit, and each of them could have cost her a better school against other submissions. In the remaining 20, the misreport changed the school she received. The mechanism gives students no reason to avoid the 6 dominant misreports. Each of them reorders only schools that the student likes less than her district school, where her priority guarantees her a seat. The other 52 misreports reverse a comparison that matters for the student's assignment, so each of them exposed her to the risk of a worse school. The theory thus sorts the observed misreports into three groups: 6 that are robustly harmless, 32 that were harmless only as realized, and 20 that changed the assignment. Section 1.1 explains how this classification complements behavioral explanations of misreporting.

The results rest on a few assumptions, which are worth stating at the outset. First, every student ranks every school, and every school is acceptable to every student. The theory is about this complete-list domain, and Remark 4.15 discusses what changes when shorter lists are allowed. Second, the exact characterization requires aggregate slack, that is, more seats in total than students. The guarantee in the first part of the theorem holds in every market. Third, priorities are known when students submit their rankings, or they are set by a lottery drawn afterward. In the second case, a ranking is dominant exactly when it is dominant after every possible draw (Proposition 6.1). Section 6 uses this for the Chen–Sönmez data. Fourth, the main theorem covers DA, school-priority TTC, and every other mechanism with the properties described above.

1.1 Contribution and relation to the literature

The main contribution is Theorem 1, which concerns the domain where every student must rank every school. It gives a condition that is sufficient for dominance in every market: the ranking is truthful down to the student's guarantee rank. It also shows that this condition is necessary exactly when there is a spare seat. The result holds under DA and TTC, and under every other strategy-proof mechanism with the three further properties that the two share. Fix the capacities. If there is at least one more seat than students, then for every priority profile a student's dominant reports are exactly the rankings that are truthful through her guarantee rank. That rank is the rank of the worst school that she can receive by reporting truthfully. Every such mechanism then has these same dominant reports, although the mechanisms can give her different schools and different menus. The dominant reports are therefore determined by the student's preferences, the priorities and the capacities, and not by the mechanism. If there is no spare seat, the picture changes. For each such mechanism, some priority profile gives a student more dominant reports than these, and no single rank describes them.

The proof replaces one question with another, and the difference between the two is the core of the paper. A natural first question is, school by school, which schools a student could ever be assigned. Call this the range question. The answer to the range question does not determine the dominant reports. In Example 1.1, each of the three schools can be assigned to s_3. Yet her ranking of c_1 against c_2 never matters, because no rankings of the others give her a menu that contains both schools. What determines the dominant reports is a joint question. For which pairs of schools can the others submit rankings that give her a menu containing both schools and no school that she prefers to both? The proof answers this joint question. The spare seat is what allows the others to do this for every pair that involves one of her schools down to the guarantee rank.

The method comes from social choice. Section 2 builds on the option-set argument of Barberà (1983): under a strategy-proof rule, each agent receives her favorite outcome among those she can

obtain by changing her own report. Decisive pairs resemble Barberà's pivotal voters. Arribillaga and Massó (2016) also study dominance type by type. For generalized median voter schemes, they collect the true preferences at which the truthful report is not a dominant report and compare mechanisms by inclusion of these sets. In the present paper, the mechanisms are strategy-proof, so truthful reporting is always dominant. The object is instead the set of all dominant reports of a given type.

The closest antecedents in matching are Kumano and Watabe (2011, 2012), Decerf and Van der Linden (2021), and Decerf et al. (2024). Kumano and Watabe show that, when students may truncate their lists, every dominant DA report must be truthful above the student's assignment, and that dominant play reproduces the truthful DA matching. Decerf and Van der Linden characterize related safe sets under constraints on list length. Decerf, Haeringer, and Van der Linden show that priorities and capacities determine the set of assignments that a student can possibly receive under incontestable and top-top-consistent mechanisms. This answers a range question. As Example 1.1 shows, the range does not determine the dominant reports. The present paper uses their two properties to define its class of mechanisms. The guarantee in Theorem 1(i) is the counterpart of their result for complete rankings. Three ingredients of the present analysis are known. Reporting truthfully through a guaranteed school is sufficient for dominance, non-truthful dominant reports can exist, and a student's menu under DA has a simple description (Kumano and Watabe, 2011; Gonczarowski et al., 2026; Katuščák and Kittsteiner, 2025). The contribution here is to determine when, on the domain of complete rankings, dominance requires an entire top segment of the ranking to be truthful. In the other markets, the paper shows how joint menus leave more of the ranking free.

Sirguiado and Torres-Martínez (2024) study one-to-one matching markets in which agents cannot declare partners unacceptable, the analogue of the complete-list domain studied here. They show that the uniqueness of a stable mechanism that is strategy-proof for one side survives only for the short side of an unbalanced market. Suppose that each school in the present setting has one seat. Students are then the short side exactly when there are more schools than students, that is, exactly when there is a spare seat, as in Theorem 1. The two questions are very different: they study how many stable strategy-proof mechanisms there are, while this paper studies the structure of dominant reports within a fixed mechanism. In both cases, however, an incentive result that can fail in balanced markets with complete lists holds once there is a spare seat.

Beyond the theory, the results have three uses, which the paper develops as consequences of Theorem 1. First, in experimental data they separate the misreports that cannot hurt the student from those that can (Section 6). Second, they show what a submitted ranking reveals about preferences. Under the weaker assumption that a submitted ranking is a dominant report, a ranking submitted in a market with a spare seat identifies the student's preferences down to the guarantee rank, and nothing below it (Corollary 4.6). This complements empirical work that estimates preferences from submitted rankings under assumptions weaker than truth-telling, such as Fack et al. (2019). Third, they show how a mechanism can ask for less than a complete ranking: DA can ask for one school at a time and stop at the guarantee rank without changing the matching (Section 7).

Finally, the paper complements behavioral explanations of misreporting. The classification sorts observed misreports into three groups by their consequences for the student's assignment: those that are robustly harmless, which are dominant reports; those that were harmless only as realized; and those that changed the assignment. Within the first group, assignment incentives do not favor any report over another, so which dominant report a student submits is a question of selection. Reporting costs, defaults, norms, or behavioral factors can answer it, and Section S5 of the supplement models one such channel, a small cost of departing from a displayed default

order. A misreport outside the first group exposes the student to a worse school against some rankings of the others, so an explanation must say why she accepts that risk. Misunderstanding the mechanism is the usual answer. Others are beliefs about the rankings of the other students, expectations-based loss aversion (Dreyfuss et al., 2022), disappointment aversion, which Chen et al. (2024) test experimentally, rational inattention (Bucher and Caplin, 2021), report-dependent or rankings-dependent preferences (Meisner, 2023; Kloosterman and Troyan, 2026), and mistakes of the kind that Artemov et al. (2023) model. In the Chen–Sönmez data, most observed misreports fall outside the first group (Section 6), so they do call for such explanations.

Li (2017) asks a related question: when is a dominant strategy easy to recognize? A strategy is obviously dominant when recognizing it requires no contingent reasoning. His laboratory experiment includes two implementations of random serial dictatorship: an obviously strategy-proof one, in which subjects choose in turn, and the direct one, in which they submit rankings. Deviations from the dominant strategy are much rarer under the first. The classification of the present paper applies before any such explanation. In Li's direct treatment, subjects rank four prizes before the queue is drawn, so only the truth is dominant, and every misreport calls for an explanation (Section S7.1 of the supplement). In the Chen–Sönmez data, by contrast, some misreports are dominant (Section 6).

1.2 Organization

Section 2 sets up the model and recalls Barberà's option-set argument, which reduces dominance to menus and leads to decisive pairs. Section 3 treats serial dictatorship, where a student's place in the queue determines the answer. Section 4 introduces the guarantee rank, which takes the place of the queue position under DA and TTC, and describes the class of mechanisms that the main theorem covers. It then states and proves the theorem and shows what a dominant report reveals about preferences. Section 5 describes the menus of DA and the dominant reports in balanced markets, and Section 6 applies the classification to the Chen–Sönmez experiment. Section 7 concludes and discusses how a mechanism can ask for less than a complete ranking. The appendix contains the remaining proofs. A separate supplement collects further results. It covers excess demand, comparative statics, and more on balanced markets. It also treats serial dictatorship with a random queue, how a student might choose among her dominant reports, preference elicitation, two more experiments, and large markets.

2 Model

This section sets up the model and introduces the tool on which the paper builds: the option sets of Barberà (1983), which the paper calls menus. Section 2.1 describes students, schools, and rankings, and Section 2.2 describes the mechanisms. Section 2.3 defines dominant reports and menus and recalls Barberà's argument, which links the two. Section 2.4 uses this link to describe the dominant reports through decisive pairs of schools. Example 1.1 serves as a running illustration.

2.1 Students, schools, and rankings

Let $N = \{s_1, \dots, s_n\}$ be a finite set of students and $C = \{c_1, \dots, c_m\}$ a finite set of schools. School c has $q_c \geq 1$ seats, and c_0 stands for the outside option of being unassigned. Write $Q = \sum_{c \in C} q_c$ for the total number of seats. A matching μ assigns each student to one school or to c_0, without giving any school more students than it has seats. Write μ_i for the assignment of s_i. Each school c ranks the students by a strict priority order $\succ_c$. A school prefers any student to an empty seat. When more students apply than it has seats, it admits those with the highest priority. That is, its choice

is responsive to its priority order. For fixed schools, capacities, and students, a priority profile is called a *market*.

The comparison between the number of seats and the number of students plays a central role. There are three cases:

$$Q > n \ \text{ (aggregate slack)}, \qquad Q = n \ \text{ (balanced market)}, \qquad Q < n \ \text{ (excess demand)}. \tag{1}$$

Under aggregate slack, there is at least one more seat than students, and the market is also said to have a *spare seat*.

Student s_i has a strict true preference P_i over $C \cup \{c_0\}$, written $\succ_i$, with weak version $\succeq_i$. Throughout, every school is acceptable, and every ranking is complete, so the analysis uses the domain

$$\mathcal{P}_i^C = \{P_i : c \succ_i c_0 \text{ for every } c \in C\}.$$

Let $\mathcal{P}^C = \prod_{i \in N} \mathcal{P}_i^C$ and $\mathcal{P}_{-i}^C = \prod_{j \neq i} \mathcal{P}_j^C$. When no other report domain is under discussion, these sets are abbreviated to $\mathcal{P}_i$, $\mathcal{P}$, and $\mathcal{P}_{-i}$. The schools in the student's true order are written

$$P_i : \quad c_{i,1} \succ_i c_{i,2} \succ_i \cdots \succ_i c_{i,m} \succ_i c_0,$$

so that $c_{i,r}$ is the school she ranks r-th.

A ranking that a student submits is written with a hat, as $\hat{P}_i$, to keep it apart from her true preference P_i. The submitted ranking may or may not be the true one. When two submitted rankings are compared, the second is written $\tilde{P}_i$. Most of the analysis looks at one student, s_i, and asks which rankings are safe for her. She is called the *focal student*, and the other students are her *rivals*. The rankings submitted by her rivals are written $\hat{P}_{-i} = (\hat{P}_j)_{j \neq i}$, and a profile of submitted rankings of all students is written $\hat{P}$. For any complete ranking $\hat{P}_i$, let $\hat{P}_i(r)$ be the school it puts at rank r, and let $\text{Top}_h(\hat{P}_i) = (\hat{P}_i(1), \ldots, \hat{P}_i(h))$ be the list of its first h schools, in order. The same notation applies to the true ranking, so that $P_i(r) = c_{i,r}$. By convention, $\text{Top}_0(\hat{P}_i)$ is the empty list, and $\text{Top}_h(\hat{P}_i) = \text{Top}_m(\hat{P}_i)$ when $h > m$. A ranking $\hat{P}_i$ is *truthful through rank h* if $\text{Top}_h(\hat{P}_i) = \text{Top}_h(P_i)$: it lists the student's h favorite schools first, in their true order, and lists the remaining schools in any order.

2.2 Mechanisms

A deterministic direct mechanism φ maps each profile of complete rankings to a matching. When s_i submits $\hat{P}_i$ and the others submit $\hat{P}_{-i}$, her assignment is written $\varphi_i(\hat{P}_i, \hat{P}_{-i})$. Some mechanisms use a random queue or random priorities. The realized draw is then treated as an exogenous state ω, and φ^ω denotes the deterministic mechanism that applies once ω is known. Unless stated otherwise, the state is held fixed, and the rankings of the other students range over the whole domain.

The mechanism is *strategy-proof* for students if no student can ever gain by misreporting: for every student s_i, true preference P_i, alternative ranking $\hat{P}_i$, and rankings $\hat{P}_{-i}$ of the others,

$$\varphi_i(P_i, \hat{P}_{-i}) \succeq_i \varphi_i(\hat{P}_i, \hat{P}_{-i}). \tag{2}$$

Strategy-proofness makes the true ranking a best report. Whether other rankings can be equally good is the question this paper studies.

Two mechanisms are central. Student-proposing Deferred Acceptance, denoted DA, returns the student-optimal stable matching and is strategy-proof for students (Gale and Shapley, 1962; Dubins and Freedman, 1981; Roth, 1982). In school-priority top trading cycles, denoted TTC (Abdulkadiroğlu and Sönmez, 2003), every remaining student points to her favorite school with a

free seat, and every such school points to its highest-priority remaining student. Each student on a cycle receives the school she points to, and the procedure repeats. TTC is also strategy-proof. Pápai (2000) introduced hierarchical exchange rules, a general class of strategy-proof top trading cycle mechanisms in a one-to-one model. The school-priority TTC studied here is the many-to-one school-choice formulation of these mechanisms by Abdulkadiroğlu and Sönmez (2003). Leshno and Lo (2021) characterize the TTC assignment by a cutoff for each pair of schools, in a model with a continuum of students.

2.3 Dominant reports and menus

The paper studies the rankings that give a student exactly what the truth gives her, whatever the other students submit. It uses the weaker convention described in Section 1: a dominant report must do at least as well as every other report against every profile of the others' reports, and nothing strict is required.

Definition 2.1 (Dominant reports)**.** *Fix φ, a student s_i, and a strict preference P_i.*

(a) A report $\hat{P}_i \in \mathcal{P}_i^C$ is a dominant report *for P_i if*

$$\varphi_i(\hat{P}_i, \hat{P}_{-i}) \succeq_i \varphi_i(\tilde{P}_i, \hat{P}_{-i}) \qquad \text{for every } \tilde{P}_i \in \mathcal{P}_i^C \text{ and every } \hat{P}_{-i} \in \mathcal{P}_{-i}^C.$$

The set of these reports is the dominant set $\mathcal{D}_i^{\varphi}(P_i)$*. When the mechanism is clear, it is written $\mathcal{D}_i(P_i)$.*

(b) Two reports $\hat{P}_i, \tilde{P}_i \in \mathcal{P}_i^C$ are own-assignment equivalent*, written $\hat{P}_i \sim_i^{\varphi} \tilde{P}_i$, if*

$$\varphi_i(\hat{P}_i, \hat{P}_{-i}) = \varphi_i(\tilde{P}_i, \hat{P}_{-i}) \qquad \text{for every } \hat{P}_{-i} \in \mathcal{P}_{-i}^C.$$

Strategy-proofness says exactly that the truth is a dominant report. Under strategy-proofness, a report is dominant if and only if it gives the student the same school as the truth against every profile of the others' reports Lemma 2.2.

Lemma 2.2 (Dominant reports under strategy-proofness)**.** *Suppose that φ is deterministic and strategy-proof and that P_i is strict. A report is a dominant report for P_i if and only if it is own-assignment equivalent to P_i:*

$$\mathcal{D}_i^{\varphi}(P_i) = \{\hat{P}_i \in \mathcal{P}_i^C : \hat{P}_i \sim_i^{\varphi} P_i\}.$$

Proof. If $\hat{P}_i$ gives the truthful assignment against every $\hat{P}_{-i}$, then $\hat{P}_i$ is a best report against every $\hat{P}_{-i}$, because the truth is. Conversely, suppose that $\hat{P}_i$ is a dominant report. Against every $\hat{P}_{-i}$, it is then at least as good as the truth, and the truth is at least as good as it, by strategy-proofness. So the two assignments are equally good under P_i. Because P_i is strict, they are the same school. □

Own-assignment equivalence is symmetric. So, under a strategy-proof mechanism, $\hat{P}_i$ is a dominant report for P_i exactly when P_i is a dominant report for $\hat{P}_i$. Section 4.3 uses this fact to read a student's preferences from her report (Corollary 4.6).

Lemma 2.2 gives every non-truthful ranking one of two meanings. A ranking inside $\mathcal{D}_i^{\varphi}(P_i)$ changes what the student says without changing what she gets, whatever the others submit. A ranking outside it is never better than the truth, and it is strictly worse against some rankings of the others. Two remarks help to interpret the dominant set. First, it concerns the student's own assignment only. Under a bossy mechanism, the assignments of other students may change.

Second, membership in the dominant set describes what a ranking does for the student's assignment, whatever her reason for choosing it. A ranking may be equivalent to the truth because the student deliberately uses the freedom that she has, follows a default, copies an earlier list, or makes a mistake that happens to be harmless. Assignment incentives do not favor any one of these rankings, so which of them a student submits is a separate question; Section S5 of the supplement models one answer.

To describe the dominant set, the paper uses the option sets of Barberà (1983). Fix the rankings $\hat{P}_{-i}$ of the other students. The set of schools that s_i can obtain by changing her own ranking,

$$M_i^{\varphi}(\hat{P}_{-i}) = \{\varphi_i(\hat{P}_i, \hat{P}_{-i}) : \hat{P}_i \in \mathcal{P}_i^C\},$$

is her option set in Barberà's terminology. This paper calls it her *menu* at $\hat{P}_{-i}$. In Example 1.1, the table lists the menus of s_3. Barberà (1983)

Lemma 2.3 (Menu choice; Barberà, 1983)**.** *If φ is deterministic and strategy-proof, then*

$$\varphi_i(P_i, \hat{P}_{-i}) = \operatorname{top}_{P_i} M_i^{\varphi}(\hat{P}_{-i}),$$

where $\operatorname{top}_{P_i} M$ denotes the best school in M according to P_i.

Proof. The truthful assignment is in the menu. If another school in the menu were strictly better under P_i, the ranking that obtains it would be a profitable deviation from the truth. Because P_i is strict, the best school in the menu is unique. □

The lemma applies to every ranking a student might submit, not only to her true one. Every complete ranking $\hat{P}_i$ is itself a possible true preference, and the menu $M_i^{\varphi}(\hat{P}_{-i})$ does not depend on her preference. Applying the lemma to a student whose true preference is $\hat{P}_i$ shows that, against fixed $\hat{P}_{-i}$, the ranking $\hat{P}_i$ gives her the school in her menu that $\hat{P}_i$ ranks highest. Let $\mathfrak{M}_i^{\varphi}$ be the family of all menus of s_i as $\hat{P}_{-i}$ varies. Two rankings are therefore equivalent exactly when they pick the same school from every menu in this family. This turns the paper's question into a question about menus: for which pairs of schools can the order in her ranking change the school that she gets from some menu? Section 2.4 answers it with decisive pairs.

The words "every menu" matter. A ranking that picks the same school as the truth from the one menu that the student actually faced has preserved her realized assignment, which is much weaker than being dominant. Likewise, a comparison that an algorithm never consults when it runs on the submitted rankings may still matter against other rankings of the others. Dominance, harmlessness at the realized rankings, and what an algorithm happens to consult are therefore three different things.

2.4 Decisive pairs

By Lemma 2.3, a ranking matters only through the school it picks from each menu. Taking one menu, the truth picks the school she likes best in it. Any other ranking picks the same school as long as it keeps that school above the other schools in the menu. How it orders those other schools makes no difference, since she receives none of them.

Therefore, her ranking of two schools matters when its pair is decisive. From that menu, the truth gives her the better of the two. If she reverses their order in her ranking, she receives the worse one instead. Against those rankings of the others, she is therefore pivotal between the two schools: the order in which she ranks them decides which of them she gets. Such a pair is called a decisive pair.

The idea resembles the pivotal voters of Barberà (1983), who can switch a social outcome between two alternatives by changing their votes. The difference is that a decisive pair concerns the student's own assignment, and the question here is which pairs of schools the other students can make available to her together, through their rankings.

To state the definition, let

$$\Delta(P_i, \hat{P}_i) = \big\{\{c, c'\} : P_i \text{ and } \hat{P}_i \text{ order } c \text{ and } c' \text{ differently}\big\}$$

be the set of pairs of schools on which two complete rankings disagree. Throughout, a pair of schools $\{c_{i,r}, c_{i,t}\}$ is also written as its pair of true ranks (r, t), where $r < t$, so that $\Delta(P_i, \hat{P}_i)$ can be read as the set of rank pairs that $\hat{P}_i$ reverses.

Definition 2.4 (Decisive pairs)**.** *Fix φ, s_i, and P_i. True ranks $1 \le r < t \le m$ form a* decisive pair *if some menu in $\mathfrak{M}_i^{\varphi}$ contains $c_{i,r}$ and $c_{i,t}$ but none of $c_{i,1}, \dots, c_{i,r-1}$. Write $E_i^{\varphi}(P_i)$ for the set of decisive pairs, and $\lhd_i^{\varphi}$ for the partial order it generates, that is, its transitive closure.*

Equivalently, a menu whose best school under the true preference has rank r_M contributes the pairs (r_M, t) for every other school $c_{i,t}$ in it. The set $E_i^{\varphi}(P_i)$ collects these pairs over all menus. In Example 1.1, the menus $\{c_1, c_3\}$ and $\{c_2, c_3\}$ contribute the pairs $(1, 3)$ and $(2, 3)$. No menu contributes $(1, 2)$, because no menu contains c_1 and c_2 together.

A ranking is dominant if and only if it keeps the true order of every decisive pair. How it orders the other pairs never affects the student's assignment.

Proposition 2.5 (Joint-menu characterization)**.** *Let φ be deterministic and strategy-proof, and let P_i be strict. A complete report is dominant for type P_i if and only if it reverses no decisive pair:*

$$\hat{P}_i \in \mathcal{D}_i^{\varphi}(P_i) \iff \Delta(P_i, \hat{P}_i) \cap E_i^{\varphi}(P_i) = \varnothing. \tag{3}$$

Consequently, $\mathcal{D}_i^{\varphi}(P_i)$ is the set of linear extensions of $\lhd_i^{\varphi}$. The order $\lhd_i^{\varphi}$ can be recovered from the dominant set, and $|\mathcal{D}_i^{\varphi}(P_i)|$ is its number of linear extensions.

Proof. Fix $\hat{P}_{-i}$ and let $M = M_i^{\varphi}(\hat{P}_{-i})$. Every complete ranking is itself a strict type in the domain, so Lemma 2.3 gives $\varphi_i(\hat{P}_i, \hat{P}_{-i}) = \operatorname{top}_{\hat{P}_i} M$ for every $\hat{P}_i$. If $M = \{c_0\}$, every ranking gives c_0. Otherwise let $c_{i,r} = \operatorname{top}_{P_i} M$. Every other school $c_{i,t} \in M$ has $t > r$, and the menu M shows that (r, t) is decisive. A ranking that reverses no decisive pair therefore ranks $c_{i,r}$ above every other school in M, so it picks $c_{i,r}$, as the truth does. Because $\hat{P}_{-i}$ was arbitrary, the ranking is equivalent to the truth.

Conversely, suppose that $\hat{P}_i$ reverses a decisive pair (r, t), and let $M = M_i^{\varphi}(\hat{P}_{-i})$ be a menu that contains $c_{i,r}$ and $c_{i,t}$ and no school ranked above $c_{i,r}$. The truth picks $c_{i,r}$ from M, whereas $\hat{P}_i$ ranks $c_{i,t}$ above $c_{i,r}$ and so picks a different school. Lemma 2.2 completes the equivalence. Finally, a linear order keeps every pair of a set in its true order if and only if it keeps every pair of the transitive closure of that set. A partial order is the intersection of its linear extensions (Szpilrajn, 1930). □

The simplest form the answer can take is the one that serial dictatorship gives (Section 3): the student must report her top d schools truthfully and in order, and she may order the rest freely. In terms of pairs, this is the order

$$O_d = \{(r, t) : 1 \le r < t \le m,\ r \le d\},$$

which requires each of her top d schools to stay above every school she ranks below it. Call O_d the *top-segment order* of length d. It is empty for $d = 0$, and for $d \ge m - 1$ it requires the whole true ranking.

Corollary 2.6 (Top-segment form)**.** *Under the assumptions of Proposition 2.5,*

$$\mathcal{D}_i^{\varphi}(P_i) = \{\hat{P}_i : \mathrm{Top}_d(\hat{P}_i) = \mathrm{Top}_d(P_i)\} \iff \lhd_i^{\varphi} = O_d.$$

Proof. The rankings that are truthful through rank d are exactly the linear extensions of O_d; apply the last part of Proposition 2.5. □

Truthful top segments and free tails. Suppose the dominant set has the form of Corollary 2.6 for some $d_i \in \{0, \ldots, m\}$, that is, it consists of the rankings that are truthful through rank d_i. Such a set is called a *top-segment set*. The student's top d_i schools, in their true order, then form her *truthful top segment*: every dominant report must list them first, in this order. The remaining schools,

$$\mathcal{U}_i^{\varphi}(P_i) = \{c_{i,d_i+1}, \ldots, c_{i,m}\},$$

form her *free tail*: dominance places no restriction on how she orders them. These terms apply whenever the dominant set has this form. Under DA and TTC with aggregate slack, d_i is the guarantee rank of Section 4.1 (Theorem 1). Under serial dictatorship with a known queue position k, $d_i = k$. The one exception is the case $k = m$, in which only one school is left by the time she chooses. Then $d_i = 0$ (Section 3).

Proposition 2.5 shows that, in general, what dominance requires is a partial order of comparisons rather than a single number of ranks. A truthful top segment is the special case in which one position in the ranking separates the comparisons that matter from those that do not. A free tail exists only in that case. Example 1.1 shows that the special case can fail. There, c_3 must be ranked below both c_1 and c_2, but the order of c_1 and c_2 is free. No single position separates the comparisons that matter from the one that does not. Whether the answer takes the simple form is a property of the menus that a market can generate. Theorem 1 shows that aggregate slack is the condition under which DA, TTC, and the other mechanisms with the four properties of Section 4.2 always give the simple form.

3 Serial Dictatorship: The Benchmark

Serial dictatorship shows the paper's question in its simplest form, without the complications that school-specific priorities create. There are m schools with one seat each and $n \leq m$ students, and the students choose one at a time in a fixed order, the queue. Suppose s_i chooses in position $k \leq n$. Before her turn, the $k-1$ students ahead of her can have taken any $k-1$ schools. Her possible menus are therefore all the sets of $m-k+1$ schools.

Proposition 3.1 (SD benchmark)**.** *If $k < m$, the dominant reports of type P_i are exactly the rankings that are truthful through rank k:*

$$\mathcal{D}_i^{\mathrm{SD}}(P_i) = \{\hat{P}_i : \mathrm{Top}_k(\hat{P}_i) = \mathrm{Top}_k(P_i)\}.$$

If $k = m$, every complete ranking is dominant.

Proof. Suppose $k < m$. Every menu has $m-k+1$ schools, while only $m-k$ schools lie below her top k, so every menu contains at least one of her top k schools. A ranking that is truthful through rank k therefore picks the same school as the truth from every menu. Conversely, take any true ranks $r < t$ with $r \leq k$. The students ahead of her can take every school she likes more than $c_{i,r}$, and they can fill their remaining turns with schools other than $c_{i,r}$ and $c_{i,t}$. This is possible because $k < m$. The resulting menu contains $c_{i,r}$ and $c_{i,t}$ and none of the schools she likes more than $c_{i,r}$, so

the pair (r, t) is decisive. By Proposition 2.5, dominance therefore requires every comparison that involves one of her top k schools and no comparison between two schools below rank k. If $k = m$, only one school is left when her turn comes, so her ranking never affects her assignment. □

The result shows the two sides of the question. As long as more than one school is left when her turn comes, the students ahead of her can remove any $k - 1$ schools. Every comparison that involves one of her top k schools can then matter, while comparisons among the schools below can never matter. When only one school is left for her, her ranking does not matter at all.

The central question is what replaces the queue position when schools have their own priorities. A single number no longer describes what the other students can do. One rival may have higher priority than the student at several schools, and DA or TTC may turn these priorities into different menus. Section 4 introduces the guarantee rank, which plays the role of the queue position, and shows that it determines all the dominant reports exactly when the market has a spare seat.

4 The Capacity Boundary

This section answers the paper's matching question. Section 4.1 introduces the guarantee rank, the object that replaces the queue position of serial dictatorship. Section 4.2 describes the mechanisms that the answer covers: DA, TTC, and every other strategy-proof mechanism with three further properties that DA and TTC share. Section 4.3 states the main result, Theorem 1, explains what it says, and shows what a dominant report reveals about preferences.

The proof works with menus. Section 4.4 begins by asking which menus the other students must be able to create, through their rankings, for a student's dominant reports to be exactly the rankings that are truthful through a given rank. Two conditions are needed. First, every menu that gives her a real choice must contain one of her schools down to that rank. This is the *guarantee condition*. Second, for every comparison that involves one of these schools, some menu must contain both schools and no school that she prefers to both. This is the *exposure condition*. The proof then proceeds in three steps, one for each part of the theorem, and it helps to see them before the details.

1. The guarantee condition holds in every market. Whatever the others submit, one of the student's schools down to the guarantee rank stays available to her, so the rankings that are truthful through the guarantee rank are always dominant (Section 4.5).

2. With a spare seat, the exposure condition holds as well, so these rankings are the only dominant ones (Section 4.6). Rivals with higher priority fill every school that she prefers to both schools of the pair. One seat is left vacant at each school of the pair, and all remaining rivals take seats elsewhere.

3. Without a spare seat, the exposure condition can fail. A market can be built in which the student's comparison of her two favorite schools never matters, although some comparison below them does. Under DA and TTC, the only comparisons that matter in this market are those with her least favorite school (Section 4.7).

Section 4.8 explains why the requirement that students rank every school is essential.

4.1 The guarantee rank

The guarantee rank describes the worst that the other students can do to a student. Go down her true ranking one school at a time. At each step, ask whether the rivals who have higher priority than she does at these schools could fill every seat at all the schools seen so far. Each rival can take

only one seat. As long as they can, the rivals could in principle shut her out of every one of these schools. The guarantee rank is the first point at which they cannot: however the rivals are seated, some seat at one of her schools down to that point is out of their reach. It is determined by her own ranking, the priorities, and the capacities alone, so it is known before anyone submits a ranking.

To make this precise, fix s_i and the market. For every school c, the *blockers* of s_i at c are the rivals with higher priority than s_i at c, and her *lower-priority rivals* at c are the rivals with lower priority:

$$B_i(c) = \{s_j \in N \setminus \{s_i\} : s_j \succ_c s_i\}, \qquad L_i(c) = \{s_j \in N \setminus \{s_i\} : s_i \succ_c s_j\}.$$

Write $H_i(h) = \{c_{i,1}, \ldots, c_{i,h}\}$ for the set of her h favorite schools, with $H_i(0) = \varnothing$. In Example 1.1, $B_3(c_1) = B_3(c_2) = \{s_1, s_2\}$ and $B_3(c_3) = \{s_1\}$. The two rivals can fill the seats of c_1 and c_2, one each. They cannot fill the three seats of c_1, c_2, and c_3, so the guarantee rank of s_3 is 3.

Counting blockers school by school is not enough, because the same rival may be a blocker at several schools and can still take only one seat. The next example shows why this matters.

Example 4.1 (Overlapping blockers)**.** Consider three students and four schools with one seat each, so that there are $Q = 4$ seats for $n = 3$ students. Student s_3 has the true ranking

$$P_3 : c_1 \succ_3 c_2 \succ_3 c_3 \succ_3 c_4 \succ_3 c_0,$$

and suppose that $B_3(c_1) = B_3(c_2) = \{s_1\}$: only s_1 has higher priority than s_3 at her two favorite schools. Student s_1 can fill the seat of c_1, so c_1 alone is not guaranteed to s_3. But s_1 can take only one seat, so she cannot fill the seats of c_1 and c_2 at the same time. One of these two schools is therefore always available to s_3, and her guarantee rank is 2. By Theorem 1 below, her dominant reports are the two rankings that list c_1 and then c_2 first; she may order c_3 and c_4 either way. Now suppose instead that $B_3(c_1) = \{s_1\}$ and $B_3(c_2) = \{s_2\}$. Each of her two favorite schools still has exactly one blocker, but the blockers are now two different students, who can fill both seats at once. Two rivals can never fill three seats, so her guarantee rank is then 3, and by Theorem 1 the truth becomes her only dominant report, because only one school is left below rank 3.

In general, whether the rivals can fill all the seats at once is a matching problem, and Hall's marriage theorem answers it. Build a bipartite graph for the set $H_i(h)$. On one side, put q_c copies of each school $c \in H_i(h)$, one for each of its seats. On the other side, put the $n - 1$ rivals. Connect every seat of c to the students in $B_i(c)$, the rivals who are allowed to take it.

Definition 4.2 (Guarantee rank)**.** *The set $H_i(h)$ is* blockable *if the graph has a matching that covers all its seats, that is, if distinct blockers can fill every seat of the schools in $H_i(h)$. Such a matching is called a* covering *of $H_i(h)$. The empty set $H_i(0)$ is blockable. The* guarantee rank *of s_i is*

$$\kappa_i = \min\big(\{h \in \{1, \ldots, m\} : H_i(h) \text{ is not blockable}\} \cup \{m + 1\}\big).$$

When $\kappa_i \le m$, the set $H_i(\kappa_i)$ is her guarantee segment*.*

When $Q \ge n$, the schools have more seats in total than there are rivals, so the rivals cannot fill all of them and $\kappa_i \le m$. The value $\kappa_i = m + 1$ can occur only under excess demand, when distinct rivals who outrank s_i can fill every seat of every school. She then has no guarantee at all. Blockability can only fail as h grows: if the rivals cannot fill the seats of her top h schools, they cannot fill the seats of her top $h + 1$ schools either.

Hall's theorem says when the seats can be filled. Because all the seats of one school are connected to the same rivals, its condition needs to be checked only for sets of schools, not for every set of seats. The lemma is what makes the guarantee rank easy to compute. It also allows a set X of rivals to be left out, which Section 5 uses.

Lemma 4.3 (Capacitated Hall test)**.** *For a set X of rivals, the seats of the schools in $H_i(h)$ can be filled by distinct blockers outside X if and only if*

$$\left|\bigcup_{c\in A} B_i(c)\setminus X\right| \;\geq\; \sum_{c\in A} q_c \qquad \text{for every } A\subseteq H_i(h).$$

With $X=\varnothing$, this is the test for blockability, and κ_i is the smallest h for which the displayed inequality fails for some $A\subseteq H_i(h)$.

Proof. Remove the rivals in X from the graph. By Hall's theorem, the seats can be matched to distinct rivals if and only if every set S of seats is connected to at least $|S|$ rivals. Let A be the set of schools that have a seat in S. All the seats of a school c are connected to the same rivals, namely those in $B_i(c)\setminus X$, so the rivals connected to S are those in $\bigcup_{c\in A} B_i(c)\setminus X$, whichever seats of these schools S contains. Among all sets of seats with the same A, the largest consists of all the seats of the schools in A, which number $\sum_{c\in A} q_c$. Hall's condition therefore holds for every S if and only if it holds for these largest sets, which is the displayed inequality for every $A\subseteq H_i(h)$. □

In the first case of Example 4.1, the inequality fails for $A=\{c_1,c_2\}$, because $|B_3(c_1)\cup B_3(c_2)| = |\{s_1\}| = 1$, while the two schools have 2 seats.

The graph only describes what the rivals could do in principle, and what this means for menus needs some care. If her top h schools are blockable, the graph shows that higher-priority rivals could fill all their seats. Whether some rankings of the rivals actually shut her out of all these schools is a separate question about the mechanism. Under DA they do. Suppose that every rival in such a covering ranks first the school that she covers. Each of these schools then receives, from the very start, as many applications from students who outrank s_i as it has seats, so it never admits her. Under TTC, an arbitrary covering can fail, because the student may trade her priority at a school she likes less for a seat that the rivals were meant to fill. Example 4.9 shows how this happens. Lemma 4.10 shows how to choose the covering so that it works under DA, TTC, and every other mechanism of Section 4.2. In the other direction, once her top κ_i schools are not blockable, at least one of them is in every menu she can face (Lemma 4.11). A school below the guarantee rank can then never be the best school in her menu.

Remark 4.4 (Individual and set-valued guarantees)**.** A school c is individually guaranteed to s_i when $|B_i(c)| < q_c$: too few rivals have higher priority to fill all its seats. If such a school has true rank r, then $\kappa_i \leq r$. The converse fails. A set of schools can be guaranteed even though each school in it could be filled on its own, because the same rivals cannot fill all of their seats at once, as in Example 4.1. The guarantee rank therefore extends the familiar idea of a guaranteed school to a guarantee for a set of schools, created by overlapping priorities. A set that is not blockable is the graph form of a safe set in Decerf and Van der Linden (2021), and it is closely related to potential overdemand in Kesten and Kurino (2019). The guarantee rank is easy to compute: add schools in the order of the student's true ranking, and update a maximum matching in the graph at each step.

4.2 The mechanisms

DA and TTC are the leading examples of the main result, but its proof uses only a few properties that the two mechanisms share. This subsection describes these properties, so that the result can be stated, and proved, once for every mechanism that has them. Two of the properties, respect for top-priority sets and top-top consistency, come from Decerf et al. (2024).

Top-top consistency compares markets with different students and schools. In this section, a mechanism is therefore a single rule that applies to every market. It applies to every set of students

and schools, to every capacity vector, and to the priorities and rankings restricted to them. DA and TTC are rules of this kind. They are run in the same way whatever the size of the market. In particular, when a smaller market arises from a larger one because some students and seats have been removed, the rule treats it exactly as it would treat the same market on its own. Each property below is required in every market.

Two definitions are needed. First, blockability is extended from s_i to any student. A set of schools A is *blockable for* s_j if distinct students other than s_j can fill every seat of the schools in A, each at a school where she has higher priority than s_j. For s_i and $A = H_i(h)$, this is the test of Definition 4.2. Second, at a profile of rankings, a student s_j and a school c form a *top-top pair* if s_j ranks c first and has the highest priority at c among all students. The four properties are the following.

(SP) *Strategy-proofness*, as in (2).

(TP) *Respect for top-priority sets.* For every profile of rankings $\hat{P}$, student s_j, and $h \geq 1$: if the set of the first h schools of $\hat{P}_j$ is not blockable for s_j, then $\varphi_j(\hat{P})$ is one of those schools.

(TT) *Top-top consistency.* Let (s_j, c) be a top-top pair at $\hat{P}$. Remove s_j and one seat of c, and keep the rankings of the other students, restricted to the remaining schools if c had a single seat. Then every student other than s_j receives under φ the same school in this smaller market as at $\hat{P}$.

(FF) *Feasible first choices.* If no school is ranked first by more students than it has seats, every student receives her first choice.

Each property has a simple meaning. Respect for top-priority sets is a minimal respect for priorities. If the students with higher priority than s_j cannot fill all the seats of the first few schools on her list, she receives one of these schools. In particular, a student in a top-top pair receives her first choice, because nobody has higher priority than she does there, so that school on its own is not blockable for her. Top-top consistency then says that she and her seat can be removed from the market without changing anyone else's assignment. Feasible first choices is a weak form of unanimity: when the students' first choices fit into the seats, everyone receives her first choice. It follows from Pareto efficiency, because when the first choices fit, the assignment that gives every student her first choice Pareto dominates every other assignment. It also follows from stability, because in that assignment nobody wants to move. Decerf et al. (2024) call an assignment incontestable if no student can successfully appeal it. They show that incontestability is equivalent to individual rationality, non-wastefulness, and respect for top-priority sets. They also show that DA and TTC are incontestable and top-top consistent. In what follows, *the class* always means the set of mechanisms with all four properties.

DA and TTC have all four properties (Lemma 4.5), so that everything proved for the class applies to them. Each property is checked directly, using the stability of the DA matching and the cycles of TTC; the proof is in Appendix A.1.

Lemma 4.5 (DA and TTC)**.** *DA and TTC satisfy (SP), (TP), (TT), and (FF).*

The class contains other mechanisms as well. Morrill (2015) introduces two variations of TTC, Clinch-and-Trade and First-Clinch-and-Trade, and shows that both are strategy-proof and Pareto efficient. Decerf et al. (2024, Proposition 2) show that both are incontestable and top-top consistent, so both have the four properties.[1] Within the class, DA and TTC are the leading examples, and the

[1]The outcome of Clinch-and-Trade can depend on the order in which cycles are processed (Morrill, 2015), so it is a mechanism once that order is fixed. The outcome of First-Clinch-and-Trade does not depend on this order.

four properties describe what the two have in common that matters for dominant reports. Respect for top-priority sets and top-top consistency give the student her guarantee and allow the rivals to keep her out of chosen schools. With feasible first choices, two chosen schools are both in her menu whenever the rivals' first choices fit into the seats and leave a vacant seat at each of the two.

4.3 The main result

For a mechanism φ, write $\mathcal{D}_i^{\varphi}(P_i)$ for the dominant set of type P_i, as in Section 2.3. A student's *worst truthful assignment* is the worst school that she can receive by reporting truthfully, as the rankings of the others vary. Its rank is

$$\max_{\hat{P}_{-i} \in \mathcal{P}_{-i}} \operatorname{rank}_{P_i}\big(\varphi_i(P_i, \hat{P}_{-i})\big),$$

where $\operatorname{rank}_{P_i}(c)$ is the position of c in P_i, and rank $m+1$ means that she is unassigned. The main result can now be stated.

Theorem 1 (Capacity Boundary)**.** *Let φ be DA or TTC, or more generally any mechanism that satisfies (SP), (TP), and (TT). Fix $m \geq 3$ schools, capacities $(q_c)_{c \in C}$, and n students, and let every ranking be complete.*

(i) Universal guarantee. *In every market, each student's worst truthful assignment under φ has rank exactly κ_i, and every report that is truthful through rank $\min\{\kappa_i, m\}$ is dominant.*

(ii) Exactness under slack. *If $Q > n$ and φ also satisfies (FF), as DA and TTC do, then for every priority profile, student, and preference,*

$$\mathcal{D}_i^{\varphi}(P_i) = \{\hat{P}_i : \operatorname{Top}_{\kappa_i}(\hat{P}_i) = \operatorname{Top}_{\kappa_i}(P_i)\}.$$

(iii) Failure without slack. *If $Q \leq n$, there exist a priority profile, a student, and a preference for which $\mathcal{D}_i^{\varphi}(P_i)$ is not the set of reports that are truthful through rank d, for any d. For DA and TTC, a single market serves both, and in it the dominant reports are exactly the rankings that put the student's least preferred school last.*

Hence, fix the capacity vector. Every mechanism with the four properties, DA and TTC among them, has the following property if and only if $Q > n$: in every priority profile, the dominant reports of every student are the reports that are truthful through some rank. When the property holds, that rank is the guarantee rank, which is the rank of the student's worst truthful assignment. All mechanisms with the four properties then have the same dominant reports, although they can assign different schools and give the student different menus.[2]

Parts (ii) and (iii) differ in their quantifiers, and the difference matters. Part (ii) is a statement about every market: once the capacities provide a spare seat, no choice of priorities can break the characterization. Part (iii) is a statement of existence: for every capacity vector without a spare seat, at least one market breaks it. Many other markets without a spare seat still have dominant

[2]For example, take three students and four schools with one seat each, so that $Q = 4 > n = 3$. Student s_3 has the true ranking $c_1 \succ_3 c_2 \succ_3 c_3 \succ_3 c_4$, and the priority orders are $s_1 \succ s_3 \succ s_2$ at c_1, $s_3 \succ s_1 \succ s_2$ at c_2, $s_1 \succ s_3 \succ s_2$ at c_3, and $s_1 \succ s_2 \succ s_3$ at c_4. Her guarantee rank is 2, so under DA and under TTC her dominant reports are the two rankings that list c_1 and then c_2 first. Suppose that s_1 ranks $c_2 \succ c_1 \succ c_3 \succ c_4$ and that s_2 ranks $c_2 \succ c_4 \succ c_1 \succ c_3$. Under DA, the menu of s_3 is $\{c_1, c_2, c_3\}$: if she ranks c_4 first, s_2 is rejected at c_2 and takes the seat at c_4. Under TTC, her menu contains all four schools: if she ranks c_4 first, she trades her top priority at c_2 for the top priority of s_1 at c_4.

reports of the simple form. Markets in which all schools use the same priority order are among them, because there DA works like serial dictatorship (Section 3).

Part (iii) is also a statement about each mechanism on its own. Because DA belongs to the class, the DA case alone would show that the characterization fails somewhere in the class. Part (iii) says more: whichever mechanism in the class is used, some market without a spare seat gives a student dominant reports of another form. This part uses only (SP), (TP), and (TT). Feasible first choices is used only in part (ii), to place the two reserved schools in the student's menu.

The restriction to at least three schools excludes only a trivial case: with two schools, every set of complete rankings that contains the truth consists of the rankings that are truthful through some rank.

What the spare seat does. With a spare seat, the answer is exact and simple in every priority profile: the dominant reports are the rankings that are truthful through the guarantee rank. The student must keep her top κ_i schools in their true positions, and she can list the schools below them in any order without ever changing her assignment, whatever complete rankings the others submit. She therefore has $(m - \kappa_i)!$ dominant reports, and the truth is her only dominant report exactly when $\kappa_i \geq m - 1$. That is the case in Example 1.1 with the second seat, where $\kappa_3 = 3 = m$. In the first case of Example 4.1, by contrast, the market has a spare seat, $\kappa_3 = 2 < m - 1 = 3$, and the student has two dominant reports. A spare seat therefore does not single out the truth. It is the condition under which the dominant reports take the same simple form in every priority profile. More seats do not single out the truth either. Once there is a spare seat, more seats keep the guarantee rank or lower it (Section S2 of the supplement), so they keep or enlarge the set of dominant reports.

The theorem can also be read from the side of the data, as a statement about what a submitted ranking reveals. Own-assignment equivalence is symmetric (Section 2.3), and this gives the following corollary of part (ii).

Corollary 4.6 (What a dominant report reveals)**.** *Let φ satisfy (SP), (TP), (TT), and (FF), and suppose that $Q > n$. Fix a priority profile and a student s_i. For a complete ranking $\hat{P}_i$, let $\hat{\kappa}_i$ be the guarantee rank of a student whose true ranking is $\hat{P}_i$, computed with the same priorities and capacities. Then, for every preference $P_i \in \mathcal{P}_i^C$,*

$$\hat{P}_i \in \mathcal{D}_i^{\varphi}(P_i) \iff \mathrm{Top}_{\hat{\kappa}_i}(P_i) = \mathrm{Top}_{\hat{\kappa}_i}(\hat{P}_i).$$

So the preferences for which $\hat{P}_i$ is a dominant report are exactly the $(m - \hat{\kappa}_i)!$ rankings that list the same $\hat{\kappa}_i$ schools first, in the same order.

Proof. By Lemma 2.2, $\hat{P}_i \in \mathcal{D}_i^{\varphi}(P_i)$ if and only if $\hat{P}_i \sim_i^{\varphi} P_i$. This relation is symmetric, so the preferences for which $\hat{P}_i$ is a dominant report are those in $\mathcal{D}_i^{\varphi}(\hat{P}_i)$, the dominant set of a student whose true ranking is $\hat{P}_i$. By part (ii) of Theorem 1, applied to that student, this set consists of the rankings that agree with $\hat{P}_i$ through rank $\hat{\kappa}_i$. □

Suppose that an analyst observes a student's submitted ranking, the priorities, and the capacities, and assumes only that the ranking is a dominant report, which is weaker than assuming that it is truthful. The analyst can compute $\hat{\kappa}_i$ from the submitted ranking itself. Corollary 4.6 then says exactly what the assumption identifies. The student's $\hat{\kappa}_i$ favorite schools and their order are identified: they are the first $\hat{\kappa}_i$ schools of the submitted ranking. Her order of the remaining schools is not identified at all, because every order of them is consistent with the assumption. In the language of partial identification, the identified set is the set of $(m - \hat{\kappa}_i)!$ rankings in the corollary.

In large markets with one spare school and priorities drawn independently at random, the guarantee rank covers on average about 63 percent of the ranking (Proposition S9.2 of the supplement). A dominant report then identifies a large part of the ranking and leaves the rest free.

The 'if' direction of the corollary holds in every market, by part (i), and the symmetry holds under every strategy-proof mechanism. Without a spare seat, the preferences for which $\hat{P}_i$ is dominant therefore still form the dominant set of a student whose true ranking is $\hat{P}_i$, which Proposition 2.5 describes through decisive pairs. In the Chen–Sönmez design of Section 6, for example, a ranking that is dominant before the lottery identifies the student's ranking down to her district school, and nothing below it. The corollary is in the spirit of empirical work that replaces the assumption of truthful reporting with weaker ones, such as the stability-based approach of Fack et al. (2019).

Remark 4.7 (The part of a ranking that DA and TTC read)**.** Under DA and TTC, a further fact holds in every market, with or without a spare seat. Take a student s_i with $\kappa_i \leq m$, and suppose that her ranking lists her top κ_i schools first, in their true order. By respect for top-priority sets, she then receives one of these schools, whatever the others report. Under DA, she proposes in the order of her ranking and never proposes beyond the school that she finally receives, so she never proposes to a school below position κ_i. Under TTC, she always points to her favorite school with a free seat. The school that she finally receives has a free seat at every earlier step, so she never points to a school below it, and so never to a school below position κ_i. Neither algorithm therefore ever reads her ranking below position κ_i, and reordering the schools there changes no student's assignment, not only her own. Applied to one student at a time, this also shows that if every student's ranking is truthful through her guarantee rank, DA and TTC produce the same matching as with complete truthful rankings.

The theorem connects two ideas that are usually kept apart: how secure a student's assignment is, and how much of her ranking must be reported truthfully. With a spare seat, the connection is exact. A student who is guaranteed one of her top k schools must report exactly those k schools truthfully, and the order below them is free. Greater security therefore means that less of the ranking matters. Without a spare seat, the two ideas can come apart in both directions. Dominance can leave the top of a ranking free while fixing its bottom, as in Example 1.1. A better priority can even reduce the number of dominant reports (Section 5.2).

Part (ii) also says something about the choice of mechanism. DA and TTC can assign different schools and give a student different menus, and they treat priorities differently: DA respects them, while TTC lets students trade them. With a spare seat, they nevertheless have exactly the same dominant reports, and so does every other mechanism with the four properties. These reports are determined by the student's preferences, the priorities and the capacities, through the guarantee rank. Within the class, which comparisons a student must report truthfully is then a property of the market rather than of the mechanism. Without a spare seat, two mechanisms in the class can have different dominant reports (Section 4.7).

The intuition behind parts (ii) and (iii) is the one given in the introduction. For the comparison between two schools to matter, the rivals must rank the schools so that three things happen. First, rivals who outrank the student fill every school that she prefers to both. Second, each of the two schools has a seat that she can take. Third, every remaining rival has a seat. The third requirement comes from the complete lists: no rival can be kept out of the market. A spare seat always leaves enough seats for this. Without it, under DA, a rival whom the student outranks at one of the two schools has to take the place of the second vacant seat. The market may not have such a rival to spare. In Example 1.1, s_3 outranks nobody at c_1 or c_2, so these two schools are never in her menu together. With a second seat at c_3, both rivals can sit at c_3 and leave c_1 and c_2 vacant together.

4.4 Two conditions on the menus

The three parts of Theorem 1 answer one question: when does a single rank separate the comparisons that matter from those that do not? Lemma 4.8 links the guarantee rank, which describes what a student is sure to receive, to dominance, which depends on which comparisons matter, through two conditions on her menus. Its central idea is the description of dominant reports by decisive pairs in Proposition 2.5. A comparison matters exactly when some menu contains both schools and no school that she prefers to both.

Lemma 4.8 (Two conditions on the menus)**.** *Let φ be strategy-proof, and fix s_i, P_i, and a rank $d \in \{0, \dots, m\}$. Consider two conditions on the menus of s_i.*

> Guarantee condition: *every menu that contains at least two schools contains one of her top d schools.*
>
> Exposure condition: *for every pair of true ranks $r < t$ with $r \leq d$, some menu contains $c_{i,r}$ and $c_{i,t}$ and none of the schools that she prefers to $c_{i,r}$.*

If the guarantee condition holds, every report that is truthful through rank d is dominant. If both conditions hold, these are exactly her dominant reports.

Proof. A pair (r, t) is decisive when some menu contains $c_{i,r}$ and $c_{i,t}$ and none of the schools that she prefers to $c_{i,r}$. Such a menu contains at least two schools, and its best school is $c_{i,r}$. Under the guarantee condition, it contains one of her top d schools, so $r \leq d$. The guarantee condition therefore says that every decisive pair belongs to the top-segment order O_d. Because O_d is transitively closed, the order $\lhd_i^{\varphi}$ generated by the decisive pairs is then contained in O_d. Every ranking that is truthful through rank d keeps O_d, so it keeps $\lhd_i^{\varphi}$, and it is dominant by Proposition 2.5. The exposure condition says that every pair in O_d is decisive. Under both conditions, the decisive pairs are exactly the pairs in O_d, and Corollary 2.6 gives the dominant set. □

The guarantee condition is the student's guarantee, seen through her menus: she is never offered a real choice that lies entirely below rank d. The exposure condition requires that every comparison down to rank d can matter, one comparison at a time. For each such comparison, some rankings of the rivals must give her a menu that contains both schools and no school that she prefers to both. Together, the two conditions are slightly stronger than the conclusion requires, because a comparison can also follow from two others through the transitivity of $\lhd_i^{\varphi}$. They are exactly what the proof of the theorem establishes. Part (i) of Theorem 1 is the guarantee condition for $d = \kappa_i$, and part (ii) is the exposure condition under slack. Part (iii) shows that without slack, the exposure condition can fail in a way that no other rank repairs.

4.5 Step 1: the guarantee

The first step shows that under every mechanism in the class, the truthful assignment has rank at most κ_i. One half is immediate. Because her top κ_i schools are not blockable, respect for top-priority sets, applied to her truthful ranking, gives her one of them whatever the others submit. The other half requires the rivals to keep her out of every school that she prefers to c_{i,κ_i}. This needs more care than it may seem. By the definition of κ_i, distinct blockers can fill every seat of these schools. Under DA, it is enough to let each of these blockers rank first the school she is meant to fill. Under TTC, this can fail, as the next example shows.

Example 4.9 (Choosing the blockers under TTC)**.** There are three students and three schools. Schools c_1 and c_2 have one seat each and c_3 has two, so there are $Q = 4$ seats for $n = 3$ students. Student s_3 has the true ranking $c_1 \succ_3 c_2 \succ_3 c_3$, and the priorities are as follows:

$\succ_{c_1}$	$\succ_{c_2}$	$\succ_{c_3}$
s_1	s_3	s_3
s_2	s_1	s_1
s_3	s_2	s_2

Student s_3 has the highest priority at c_2 and c_3, so the set $\{c_1, c_2\}$ is not blockable, and her guarantee rank is 2. Her favorite school c_1 on its own is blockable: both s_1 and s_2 have higher priority than she does there, and either of them could take its seat.

Suppose that s_2 is the rival chosen to fill c_1, so that she ranks c_1 first, and that s_1 ranks c_3 first. Under DA, this works: s_2 outranks s_3 at c_1, so s_3 can never take the seat from her. Under TTC, it fails. If s_3 ranks c_1 first, she points to c_1, which points to its highest-priority student s_1, who points to c_3, which points back to s_3. This cycle gives c_1 to s_3: she trades her top priority at c_3, which s_1 wants, for the top priority of s_1 at c_1. The seat of c_1 was not safe, because the rival chosen to fill it did not hold the strongest claim to it.

Suppose instead that s_1 is chosen to fill c_1. She has the highest priority there, so if she ranks c_1 first, she receives it under DA and TTC alike, whatever the others report. The menu of s_3 is then $\{c_2, c_3\}$ against every ranking of s_2.

The example suggests the remedy: choose the blockers so that each of them, when her turn comes, has the highest priority at her school among the students who are left. Such a rival can never be displaced or traded out. By respect for top-priority sets, she receives the school that she ranks first. By top-top consistency, she can then be removed from the market together with her seat without changing anyone else's assignment, and the next rival is then in the same position. Lemma 4.10 shows that such a choice is always possible. With it, the rivals can keep the student out of any blockable set of schools, under every mechanism in the class at once. The proof of the theorem uses the lemma twice. Here, it shows that the guarantee rank is the worst case. In Section 4.6, it ensures that no school that she prefers to both schools of a pair is in her menu.

Lemma 4.10 (Protected covering)**.** *Let φ satisfy (TP) and (TT), and let V be a set of schools that is blockable for s_i. There is an assignment of $\sum_{c \in V} q_c$ distinct rivals to the seats of V, each to a school at which she outranks s_i, with the following property: if each of these rivals ranks her assigned school first, she receives it under φ, whatever the other students report. Against every such profile of rival rankings, no school of V is in the menu of s_i.*

The assignment is built in two steps. The first step selects the right rivals. It starts from any covering of V, lets each rival in it rank her school first, and runs school-proposing DA on the schools of V alone. The result is a *closed* covering: whenever a selected rival sits at a school c, every rival with higher priority than her at c is also selected. In Example 4.9, starting from the covering that uses s_2, this step selects s_1. The second step puts the selected rivals in the right order. It runs a top-trading-cycles procedure among them, in which each school points to its highest-priority selected rival. The procedure reassigns the selected rivals to the seats of V and puts them in an order. In this order, every rival, when her turn comes, has the highest priority at her school among the students still present, s_i included. Appendix A.2 gives the proof and walks through both steps.

With the protected covering, the guarantee can be stated precisely, in the form that Lemma 4.8 requires. Lemma 4.11 describes the guarantee through the student's menus: it gives the guarantee

condition with $d = \kappa_i$, and it describes the worst case. The two parts use different ideas. Part (i) applies respect for top-priority sets to the student's own ranking. Part (ii) uses the protected covering of her schools above rank κ_i.

Lemma 4.11 (Guarantee)**.** *Let φ satisfy (TP) and (TT), and fix s_i and P_i.*

(i) If $\kappa_i \leq m$, then against every profile of rival rankings, the truthful assignment of s_i is one of her top κ_i schools. If φ is also strategy-proof, every menu of s_i contains one of these schools.

(ii) Some profile of rival rankings gives s_i a menu that contains none of her top $\kappa_i - 1$ schools. Against it, she receives c_{i,κ_i} by reporting truthfully if $\kappa_i \leq m$, and she is unassigned whatever she reports if $\kappa_i = m + 1$.

Proof. (i) Because $H_i(\kappa_i)$ is not blockable for s_i, respect for top-priority sets, applied to her truthful ranking with $h = \kappa_i$, gives her one of these schools. If φ is strategy-proof, her truthful assignment is the best school in her menu (Lemma 2.3), so every menu contains one of her top κ_i schools.

(ii) The set $V = H_i(\kappa_i - 1)$ is blockable by the definition of κ_i. Let the rivals of Lemma 4.10 for V rank their assigned schools first, and let the other rivals report anything. These rivals then take every seat of V, so no school of V is in her menu. If $\kappa_i = m + 1$, the set V contains every school, so she is unassigned whatever she reports. If $\kappa_i \leq m$, part (i) gives her one of her top κ_i schools when she reports truthfully. As she receives none of the first $\kappa_i - 1$, she receives c_{i,κ_i}. □

Proof of Theorem 1(i). By Lemma 4.11(i), the truthful assignment has rank at most κ_i against every profile of rival rankings, and by Lemma 4.11(ii), some profile gives it rank exactly κ_i. When $\kappa_i \leq m$, Lemma 4.11(i) gives the guarantee condition of Lemma 4.8 with $d = \kappa_i$, so every report that is truthful through rank κ_i is dominant. When $\kappa_i = m + 1$, the only report that is truthful through rank m is the truth, which is dominant by strategy-proofness. □

Part (i) gives the guarantee rank a direct economic meaning in every market, with or without slack. It is the rank of the worst school that the student can receive by reporting truthfully, under DA, TTC, and every other mechanism in the class. It is the counterpart, for complete rankings, of a result of Decerf et al. (2024), discussed at the end of Section 4.7. Reporting truthfully through the guarantee rank is therefore always enough for dominance. The next step shows that, with a spare seat, it is also necessary.

4.6 Step 2: why one spare seat suffices

Fix two of the student's schools, $c_{i,r}$ and $c_{i,t}$ with $r < t$. The better one, $c_{i,r}$, lies at or above her guarantee rank. To make her comparison of these two schools matter, the rivals must rank the schools so that both of them are in her menu and no school that she prefers to $c_{i,r}$ is. Three moves achieve this.

1. *Block the better schools.* Protected blockers, chosen as in Lemma 4.10, fill every seat of the schools that she prefers to $c_{i,r}$. This is possible because $r \leq \kappa_i$, so these schools are blockable.

2. *Leave two seats vacant.* One seat at $c_{i,r}$ and one seat at $c_{i,t}$ are kept out of every rival's first choice.

3. *Give everyone else a seat.* Every remaining rival ranks first a school that still has a seat for her.

No school is then ranked first by more students than it has seats, even after the student herself ranks $c_{i,r}$ or $c_{i,t}$ first. By feasible first choices, she obtains $c_{i,r}$ by ranking it first and $c_{i,t}$ by ranking it first, while the blockers keep the better schools. Her menu therefore contains both schools and no school that she prefers to $c_{i,r}$, so the pair is decisive. The only question is whether there are enough seats for the third move, and this is where the spare seat enters. Whether the third move is possible is a matter of counting seats, which the proof does.

Proof of Theorem 1(ii). Suppose that $Q > n$ and that φ satisfies (FF). Then $\kappa_i \leq m$, because $n-1$ rivals cannot fill more than $n-1$ seats. By Lemma 4.11(i), the guarantee condition of Lemma 4.8 holds with $d = \kappa_i$, so it remains to check the exposure condition. Fix true ranks $r < t$ with $r \leq \kappa_i$. Because $r - 1 < \kappa_i$, the set $V = H_i(r-1)$ of schools that she prefers to $c_{i,r}$ is blockable. Let $S = \sum_{c \in V} q_c$ be its number of seats. Let the rivals of Lemma 4.10 for V rank their assigned schools first. Reserve one seat of $c_{i,r}$ and one seat of $c_{i,t}$. Give each of the other $n - 1 - S$ rivals her own seat among the remaining seats outside V, and let her rank the school of that seat first. This is possible because $Q \geq n + 1$ gives

$$Q - S - 2 \;\geq\; n - 1 - S.$$

Complete all rankings arbitrarily. If s_i ranks $c_{i,r}$ first, no school is ranked first by more students than it has seats, so feasible first choices gives her $c_{i,r}$. In the same way, she obtains $c_{i,t}$ by ranking it first. By Lemma 4.10, she obtains no school of V, whatever she reports. Her menu therefore contains $c_{i,r}$ and $c_{i,t}$ and none of the schools that she prefers to $c_{i,r}$, which is the exposure condition for the pair (r, t). Lemma 4.8 completes the proof. □

The proof shows exactly where each assumption is used. Respect for top-priority sets and top-top consistency are used to keep the better schools out of her menu, feasible first choices is used to put the two reserved schools in it, and the spare seat is used to give every other rival a seat. From the menu that the construction produces, the truth picks $c_{i,r}$. A ranking that puts $c_{i,t}$ above $c_{i,r}$ picks a different school, because $c_{i,t}$ is also in the menu. With one seat fewer, the third move can fail. Two schools can then each be in some menu without ever being in the same one, as in Example 1.1.

4.7 Step 3: why the spare seat is necessary

The construction of Step 2 needs a spare seat. That alone does not show that the characterization needs one. Without a spare seat, a student can still have two schools in her menu, because her priority may allow her to take a seat that a rival holds. The last step shows that the spare seat is needed for the characterization to hold in every priority profile. It builds markets without a spare seat in which the exposure condition fails: some comparison that the characterization requires never matters, so a ranking that reverses it is dominant too. These markets have a simple shape: most rivals take their seats first, as under serial dictatorship, and the student then shares the last two seats with a single rival. The construction works for every capacity vector with $Q \leq n$, for every mechanism in the class, and for balanced markets as well as markets with excess demand.

Write $c_1, \ldots, c_m$ for the schools in the true order of s_i. Because every school has at least one seat, $Q \geq m \geq 3$. Choose one rival, w, and $Q - 2$ further rivals, called the *leading rivals*. The remaining $n - Q$ students, if there are any, are the *trailing rivals*. At every school, the leading rivals come first, in the same order at every school. Then come s_i and w, in an order that depends on the school, and then the trailing rivals. Let F be the set of schools at which s_i comes before w. Such a market is called a *core market*, and the set F is the only thing that varies in the proof.

The first lemma says that, in a core market, the leading rivals behave exactly as under serial dictatorship. It predicts what a general mechanism does from its properties alone, and it reduces

the student's menu to a competition for two seats. The idea is that the first leading rival has the highest priority at every school, so she and the school that she ranks first form a top-top pair. She receives that school and can be removed from the market together with her seat, and the next leading rival is then in the same position.

Lemma 4.12 (Leading rivals)**.** *Let φ satisfy (TP) and (TT), and consider a core market. Whatever the students report, the leading rivals receive the schools that serial dictatorship gives them: one after another, in their priority order, each takes her favorite school among those with a seat left. Every other student receives the school she receives in the market that remains when the leading rivals and their seats are removed. That market has two seats, and in it s_i and w have higher priority than every trailing rival at every school.*

Proof. The first leading rival has the highest priority at every school, so she and the school she ranks first form a top-top pair. By respect for top-priority sets, she receives that school. By top-top consistency, every other student receives what she receives in the market without her and one seat of that school. In that market, the second leading rival has the highest priority at every school, and the argument repeats. After $Q - 2$ rounds, two seats remain. □

Lemma 4.13 describes the competition for the last two seats. It determines which pairs of schools the student can ever find in one menu, and so which comparisons she ever faces. Its idea is the same as before: a student who has the highest priority at the school she ranks first receives it. The one case that this leaves open is the one in which s_i and w each rank first the school at which the other has the higher priority. Each of them could then gain by exchanging their claims. Say that φ *lets them exchange* if it then gives each of them her first choice.

Lemma 4.13 (The last two seats)**.** *Let φ satisfy (TP) and (TT), and consider a market with two seats in which s_i and w have higher priority than every other student at every school. Then s_i and w receive the two seats, whatever the others report, and the menu of s_i is as follows.*

(a) If both seats are at one school c, her menu is $\{c\}$.

(b) If the seats are at two schools a and b, and w comes before s_i at both, her menu consists of the one of a and b that w ranks lower.

(c) If the seats are at two schools a and b, and s_i comes before w at both, her menu is $\{a, b\}$.

(d) If the seats are at two schools a and b, and s_i comes before w at a but after w at b, her menu is $\{a\}$ when w ranks b above a. When w ranks a above b, her menu is $\{a, b\}$ if φ lets s_i and w exchange when s_i ranks b first, and $\{a\}$ otherwise.

DA and TTC always let s_i and w exchange.

Proof. The proof uses one argument throughout. A student who ranks first a school at which she has the highest priority among the students present receives it, by respect for top-priority sets. She and her seat can then be removed from the market, by top-top consistency. In case (a), whichever of s_i and w comes first at c receives a seat there, and the other then receives the second seat. In case (b), w receives the school she ranks higher, and s_i, who then comes first at the other school, receives it. In case (c), s_i receives the school she ranks first, and w receives the other. In case (d), if w ranks b above a, she receives b, and s_i receives a. If w ranks a above b and s_i ranks a first, s_i receives a, and w receives b.

The remaining case is the one in which w ranks a first and s_i ranks b first. The set $\{a, b\}$ is not blockable for s_i, because nobody has higher priority than she does at a, so respect for top-priority

sets gives her a or b. In the same way, $\{a, b\}$ is not blockable for w, because nobody has higher priority than she does at b, so w also receives a or b. The two students therefore take the two seats, either as each ranked them or the other way around. Under DA, each of them proposes to her first choice and is accepted, so they exchange. Under TTC, s_i points to b, which points to w, who points to a, which points back to s_i. This cycle gives each of them her first choice. In every case, s_i and w take both seats, so the other students receive none. □

Together, the two lemmas describe every menu of s_i in a core market. By ranking the other seats first, the leading rivals can leave the last two seats at any two schools, or at any one school with at least two seats. Every menu of s_i therefore contains at most two schools. A two-school menu $\{a, b\}$ arises exactly when s_i comes before w at both schools. It also arises when she comes before w at only one of them and φ lets the two students exchange there, for some rankings of the trailing rivals. Because a pair is decisive exactly when it forms a two-school menu, two schools outside F never form a decisive pair.

Proof of Theorem 1(iii). Suppose that $Q \leq n$, so that a core market can be formed. Recall from Corollary 2.6 that the dominant reports are the rankings that are truthful through some rank d exactly when the decisive pairs generate the top-segment order O_d. A top-segment order that restricts the ranking at all keeps her favorite school above her second favorite.

DA and TTC. Take a core market with $F = \{c_m\}$: the student comes before w only at her least preferred school. By Lemma 4.13, both mechanisms let s_i and w exchange, so the decisive pairs are exactly the pairs (r, m) with $r < m$. Her dominant reports are therefore the rankings that put c_m last, with the other schools in any order. This set restricts the ranking but leaves c_1 and c_2 free, because $m \geq 3$, so it is not a top-segment set.

Any mechanism in the class, with $m \geq 4$. Take a core market with $F = \{c_{m-1}, c_m\}$. By Lemma 4.13(c), the pair $(m-1, m)$ is decisive, so the dominant reports restrict the ranking. The pair $(1, 2)$ is not decisive, because c_1 and c_2 both lie outside F. It does not follow from other decisive pairs either, because no school lies between c_1 and c_2 in her ranking. So the dominant set again restricts the ranking but leaves c_1 and c_2 free, and it is not a top-segment set.

Any mechanism in the class, with $m = 3$. Here two markets are needed, and they differ only in the priorities at c_2. In market A, $F = \{c_3\}$, and the student's true ranking is $c_1 \succ c_2 \succ c_3$. In market B, $F = \{c_2, c_3\}$, and her true ranking is $c_1 \succ c_3 \succ c_2$. The school names are the same in both markets, so in market B they no longer follow her ranking. When the leading rivals leave the last two seats at c_1 and c_3, they fill c_2. The market that remains, with one seat at each of c_1 and c_3, is then the same in A and in B. So φ treats it in the same way in both markets. Either φ lets s_i and w exchange there for some rankings of the trailing rivals, or it never does. If it does, then in market A the pair $\{c_1, c_3\}$ is decisive, while $\{c_1, c_2\}$, the pair of her two favorite schools, is not. As in the case $m \geq 4$, the dominant set in market A is then not a top-segment set. If it never does, then in market B the pair $\{c_2, c_3\}$ is decisive, by Lemma 4.13(c), while $\{c_1, c_3\}$, the pair of her two favorite schools in market B, is not. The dominant set in market B is then not a top-segment set. □

The construction shows why the spare seat is necessary. In a core market, the leading rivals take all but two seats, and the student can have two schools in her menu only by sharing the last two seats with w. Which pairs she can face then depends on who comes first at the two schools, so the comparisons that matter are determined by the set F rather than by her own ranking. With a spare seat, she would have a seat that nobody else needs, and every comparison down to her guarantee rank could then matter. Under excess demand, the trailing rivals change nothing: they come after s_i and w everywhere and never receive one of the last two seats. Example 1.1 is the smallest core market.

Example 4.14 (Example 1.1, continued)**.** Example 1.1 has $Q = n = 3$ and no trailing rivals. Student s_1 is the only leading rival, w is s_2, and $F = \{c_3\}$, because s_3 comes before s_2 only at c_3. By Lemma 4.12, s_1 receives the school she ranks first, as the example observed directly. If she takes c_3, the last two seats are at c_1 and c_2, where s_2 comes first, so by Lemma 4.13(b) the menu of s_3 is whichever of the two schools s_2 ranks lower. If s_1 takes c_1 or c_2, the last two seats are at c_3 and at the other of these two schools. By Lemma 4.13(d), the menu of s_3 is then $\{c_3\}$ when s_2 prefers the other school to c_3, and it contains both schools when s_2 prefers c_3. These are the six menus in the table of Example 1.1. The guarantee rank of s_3 is 3, so truthful reporting can leave her at c_3, her last choice. The truth is dominant, as Theorem 1(i) says. Because c_1 and c_2 never share a menu, the ranking that puts c_2 above c_1 is dominant too, so dominance fixes only the position of c_3. A second seat at c_3 leaves $\kappa_3 = 3$ unchanged but creates a spare seat. By Theorem 1(ii), her dominant reports are then the rankings that are truthful through rank 3. With only three schools, that is the truth alone, as the example showed directly.

Without slack, two mechanisms in the class can also have different dominant reports. To see this, change Example 1.1 by moving s_3 to the top of the priority order at c_3. Under DA, c_1 and c_2 still never share a menu, so both rankings that put c_3 last remain dominant. Under TTC, the change makes a trade possible: if s_1 ranks c_3 first, s_3 can trade her top priority at c_3 for the top priority of s_1 at c_1 or at c_2. All three schools are then in one menu, and the truth is her only dominant report. With a spare seat, Theorem 1(ii) shows that such differences disappear.

These examples also show how dominance differs from the question of which schools a student can receive. Decerf et al. (2024) study the set of schools that a student can receive as the other students' preferences vary, on a domain in which students may find schools unacceptable. For every incontestable and top-top-consistent mechanism, this set is her smallest top-priority set, which for a complete ranking is $H_i(\kappa_i)$ when $\kappa_i \leq m$. Theorem 1(i) is the counterpart, for complete rankings, of the bottom end of this set: the worst school that she can receive by reporting truthfully is exactly her school at rank κ_i. The set of possible assignments looks at one school at a time. Dominance depends instead on *joint* menus: the student's ranking of two schools matters only if she can obtain both against the same rival rankings. In Example 1.1, each of the three schools is a possible truthful assignment of s_3, yet c_1 and c_2 never share a menu. In the variant above, DA and TTC give s_3 the same possible truthful assignments, all three schools, but different dominant reports.

4.8 Why complete rankings matter

Remark 4.15 (The truncated domain)**.** On a domain that allows truncated lists, the capacity boundary disappears. Proposition 1 of Kumano and Watabe (2011) shows that every dominant DA report must rank truthfully the schools that a student prefers to her assignment. Combined with the worst-case guarantee of Lemma 4.11, this shows that the dominant reports are exactly those that are truthful through the guarantee rank, whether or not there is a spare seat. Compare the safe-set formulation of Decerf and Van der Linden (2021). The necessity part of the Kumano–Watabe argument uses rivals who list one school or none. With complete rankings, every rival stays in the market. The rivals then have fewer ways to shape the student's menu, so fewer comparisons can be made to matter, and her dominant set can be larger. Theorem 1 shows exactly when the conclusion survives this restriction: a spare seat lets every rival be seated while two schools remain available. Kumano and Watabe allow substitutable, quota-filling choice functions for schools, which are more general than the responsive choice induced by strict priorities here. The distinction that this paper isolates is the domain of student reports. The restricted-length environment of Marutani (2018) poses a different question: there DA can lose strategy-proofness, and the solution concept is Nash implementation.

5 Dominant Reports without Slack

When there are no more seats than students, reporting truthfully through the guarantee rank remains sufficient for dominance (Theorem 1(i)), but other rankings can be dominant as well. The missing seat changes which menus can occur: two schools can each be available on their own without ever being available together. For balanced markets under DA, this section describes exactly which comparisons matter. Section 5.1 recalls a known description of DA menus and shows which of these menus can arise when every student ranks every school. Section 5.2 uses it to describe the decisive pairs in balanced markets. The answer turns on a rival who can take the place of the second vacant seat used in the slack proof. The supplement treats excess demand (Section S1) and comparative statics (Section S2), and it adds more on balanced markets (Section S3).

5.1 Menus under DA

Throughout this section, the mechanism is DA. Write $\mathcal{D}_i(P_i) = \mathcal{D}_i^{\mathrm{DA}}(P_i)$, and write $E_i(P_i)$ and $\lhd_i$ for the decisive pairs under DA and the order they generate.

The starting point is a known description of a student's menu under DA. Fix the reports $\hat{P}_{-i}$ of the rivals of s_i. Remove s_i from the market, and call what remains the *reduced market*. Run school-proposing DA in the reduced market, among her rivals alone, and call the resulting matching η^{-i}. Then

$$c \in M_i^{\mathrm{DA}}(\hat{P}_{-i}) \iff |\eta^{-i}(c)| < q_c \quad \text{or} \quad \exists s_j \in \eta^{-i}(c) \text{ with } s_i \succ_c s_j \tag{4}$$

(Gonczarowski et al., 2026), and her menu is $\{c_0\}$ when no school satisfies the condition. In words, a school is in her menu when, in the reduced market, it has an empty seat or it has admitted a rival whom she outranks there. This formula is usually stated for a student who may submit a partial list, and it remains true when her reports must be complete rankings. Indeed, if some partial list obtains c, consider the complete ranking that puts c first. If this ranking were her true preference, strategy-proofness would give her a school at least as good as c, which can only be c itself. So the complete ranking obtains c as well. The formula is therefore taken as given. What the analysis adds is the answer to a different question: which of these menus can actually arise when every *rival* must also submit a complete ranking?

Consider any matching η of the rivals that respects capacities; some rivals may be unassigned. Write $A_i(\eta)$ for the set of schools that satisfy the condition in (4) at η, that is, the schools that have an empty seat at η or hold a rival whom s_i outranks there. Call η *realizable* if it is the matching η^{-i} produced by some complete reports $\hat{P}_{-i}$ of the rivals, and let

$$\mathfrak{R}_i = \{\eta^{-i}(\hat{P}_{-i}) : \hat{P}_{-i} \in \mathcal{P}_{-i}^C\}$$

be the set of realizable matchings. The menus of s_i are exactly the sets $A_i(\eta)$ with $\eta \in \mathfrak{R}_i$. So the question of which comparisons can matter becomes a question about which matchings of the rivals their complete reports can produce. When there are at least as many seats as students, every way of seating all the rivals can be produced (Lemma 5.1). It is what makes DA tractable in balanced markets, because it turns a question about rankings into a question about seatings. The idea of the proof is to let each rival rank her own school first and a school with an empty seat second. This matching is then stable, and the rural hospitals theorem shows that it is the only stable one.

Lemma 5.1 (Realizable matchings)**.** *Suppose $Q \geq n$, and let η be a matching that seats every rival of s_i at a school without exceeding any capacity. Then there are complete reports of the rivals for which η is the unique stable matching of the reduced market; in particular, η is the outcome of school-proposing DA in that market. Conversely, when $Q \geq n$, every realizable matching seats every rival.*

The proof is in Appendix A.3. So when $Q \geq n$, the realizable matchings are exactly the ways of seating all the rivals without exceeding any capacity. Under excess demand, fewer matchings are realizable. The supplement describes them (Section S1).

Combining the menu formula with Definition 2.4 gives a test for decisive pairs under DA. The description of balanced markets is built on this test. A pair is decisive when some realizable matching has two features. Every school that the student prefers to both schools of the pair is full, and only her blockers sit there. And each school of the pair is open to her.

Lemma 5.2 (DA decisive pairs)**.** *True ranks $r < t$ form a decisive pair if and only if some $\eta \in \mathfrak{R}_i$ satisfies both of the following:*

(W1) for every $\ell < r$, school $c_{i,\ell}$ is full at η, and every rival assigned to it belongs to $B_i(c_{i,\ell})$;

(W2) each of $c_{i,r}$ and $c_{i,t}$ has an empty seat at η or holds a rival whom s_i outranks there, that is, a rival in $L_i(c_{i,r})$ or in $L_i(c_{i,t})$, respectively.

Consequently, $\mathcal{D}_i(P_i)$ is the set of linear extensions of the order generated by these pairs.

Proof. By (4), a school c is missing from $A_i(\eta)$ exactly when it is full and every rival assigned to it outranks s_i there. Conditions (W1) and (W2) therefore say that the menu produced by η contains $c_{i,r}$ and $c_{i,t}$ and no school ranked above $c_{i,r}$. Apply Definition 2.4 and Proposition 2.5. □

5.2 Balanced markets

When $Q = n$, every realizable matching seats all $n - 1$ rivals (Lemma 5.1) and leaves exactly one seat empty. Only one of two schools can have that empty seat, so two schools can be in the menu of s_i together only if at least one of them holds a rival whom s_i outranks there. Whether such a rival is available depends on which rivals are needed to block the schools that she prefers to both. For $r \leq \kappa_i$, call a rival *essential* for rank r if she appears in every way of filling the seats of $H_i(r-1)$ with distinct blockers. Let $\mathrm{Ess}_i(r)$ be the set of these rivals, with $\mathrm{Ess}_i(1) = \varnothing$. By Lemma 4.3 with $X = \{w\}$, a rival w is not essential for rank r if and only if the seats of $H_i(r-1)$ can be filled without w.

Proposition 5.3 (Balanced markets)**.** *Suppose $Q = n$. True ranks $r < t$ form a decisive pair if and only if the seats of $H_i(r-1)$ can be filled by distinct blockers while leaving unused at least one rival in $L_i(c_{i,r}) \cup L_i(c_{i,t})$. Equivalently,*

$$(r,t) \in E_i(P_i) \iff r \leq \kappa_i \quad \text{and} \quad L_i(c_{i,r}) \cup L_i(c_{i,t}) \not\supseteq \mathrm{Ess}_i(r). \tag{5}$$

In words, a rival takes the place of the second empty seat that a spare seat would supply. She must be a rival whom s_i outranks at one of the two schools. She must also not be needed to block the schools that s_i prefers to both. Computing $E_i(P_i)$ requires at most n bipartite matching problems for each rank. In Example 1.1, s_3 outranks nobody at c_1 or c_2, so $L_3(c_1) = L_3(c_2) = \varnothing$, and the pair $(1,2)$ fails (5). At c_3 she outranks s_2, so $L_3(c_3) = \{s_2\}$, and this makes $(1,3)$ and $(2,3)$ decisive. Take the pair $(2,3)$, for instance. The only school that s_3 likes more than c_2 is c_1, and s_1 alone can fill its seat. This leaves s_2 free to sit at c_3, where s_3 outranks her. The empty seat is then at c_2, and the menu of s_3 is $\{c_2, c_3\}$. The proof is in Appendix B.

Corollary 5.4 (Consequences for balanced markets)**.** *Suppose $Q = n$ and $m \geq 2$.*

(i) The reports that are truthful through the guarantee rank, $\{\hat{P}_i : \mathrm{Top}_{\kappa_i}(\hat{P}_i) = \mathrm{Top}_{\kappa_i}(P_i)\}$, are exactly the dominant reports whenever $L_i(c_{i,r}) \not\supseteq \mathrm{Ess}_i(r)$ for every $r \leq \kappa_i$.

(ii) Report irrelevance. *Every complete report is dominant if and only if s_i has the lowest priority at every school.*

(iii) Common priorities. *Suppose every school uses one priority order in which s_i has position p. If $p < n$, then $\mathcal{D}_i(P_i) = \{\hat{P}_i : \text{Top}_{\kappa_i}(\hat{P}_i) = \text{Top}_{\kappa_i}(P_i)\}$ with $\kappa_i = \min\{h : \sum_{\ell \le h} q_{c_i,\ell} \ge p\}$. If $p = n$, every report is dominant.*

Part (iii) and Theorem 1(ii) contain the serial dictatorship benchmark of Section 3 as the case of common priorities. This includes its jump at the last position: the SD student who finds only one school left corresponds to the balanced-market student with the lowest priority at every school. With equal capacities q, the formula gives $\kappa_i = \lceil p/q \rceil$.

Part (ii) also shows that a better priority can leave a student with fewer dominant reports. Suppose that s_3 in Example 1.1 had the lowest priority at c_3 as well, and so at every school. By part (ii), all six of her rankings would then be dominant: she would always receive the school that the others leave. Her actual priority at c_3, above s_2, sometimes lets her choose between c_3 and another school. These choices fix the position of c_3 in her ranking, so only the two rankings that put c_3 last remain dominant. With a spare seat, a better priority always keeps or enlarges the set of dominant reports, because it keeps the guarantee rank or lowers it. Section S2 of the supplement treats these comparative statics in general. Section S3 reports how often, in random balanced markets, the dominant reports differ from the rankings that are truthful through the guarantee rank. It also gives a second example in which DA and TTC have different dominant reports.

6 Classifying non-truthful Reports

Three properties of a report can be determined from experimental data, and they need to be kept apart. A report may differ from the student's true ranking. It may also give her the same assignment as the truth against the reports that the other students actually submitted, that is, the same realized assignment. Or it may give her the same assignment as the truth against every possible profile of the others' reports. Only the third makes the report dominant. Counting the reports that kept the realized assignment can therefore greatly overstate the freedom that strategy-proofness gives students. Koutout et al. (2021) also classify sub-optimal DA strategies in a school choice experiment, using a typology of strategies. The classification here sorts reports by what they do to the student's assignment.

The DA treatment of Chen and Sönmez (2006) provides a direct application. It has 36 students and seven schools with capacities $(3, 3, 6, 6, 6, 6, 6)$, and therefore exactly 36 seats. Every student lives in the district of one school, and each school's district has exactly as many residents as the school has seats. Each school gives priority to its own residents over everyone else. Ties within each of these two groups are broken by a single lottery, which was drawn after the rankings had been submitted. The analysis uses the subject-level file `chen_sonmez_sc_exp_data.txt` from openICPSR project 119341. Rerunning DA on these reports reproduces all 144 recorded assignments.

6.1 Rankings submitted before the lottery

Because the lottery was drawn after the rankings were submitted, each student had to choose her ranking without knowing the draw. As in Section 2.2, the draw is an exogenous state ω. Once it is known, a deterministic mechanism φ^ω applies, here DA with the priorities that the draw produces. A ranking is *dominant before the draw* if, against every profile of the others' rankings, its expected utility is at least that of every other ranking. Proposition 6.1 extends Sections 2 to 5, where priorities

are known, to rankings submitted before the draw. Its idea is simple. The truth is the best report in every state. So a ranking that does worse than the truth against some rankings of the others, in a state with positive probability, does strictly worse on average against those same rankings.

Proposition 6.1 (Dominance before a lottery)**.** *Suppose that, after the reports are submitted, a state ω is drawn from a finite set Ω, that every state has positive probability, and that the mechanism φ^ω used in state ω is strategy-proof. Fix a utility function that represents P_i. A report is dominant before the draw if and only if it is dominant in every state:*

$$\mathcal{D}_i(P_i) = \bigcap_{\omega\in\Omega} \mathcal{D}_i^{\varphi^\omega}(P_i).$$

Equivalently, the dominant reports are the linear extensions of the order generated by the decisive pairs of all states together, $\bigcup_{\omega\in\Omega} E_i^{\varphi^\omega}(P_i)$. In particular, they do not depend on the probabilities of the states or on the utility function.

The proof is in Appendix C. A ranking that is dominant once the draw is known need not be dominant before it, because another draw can make more comparisons matter. In the Chen–Sönmez design, the decisive pairs of all draws together have a simple form, which Lemma 6.2 describes. Its idea is the following. A student's district school is guaranteed to her whatever the draw, because the other residents of her district are fewer than its seats. And in the draw that puts her first in the lottery, every comparison down to her district school matters.

Lemma 6.2 (District schools)**.** *Suppose that every student lives in the district of one school, and that each school's district has exactly as many residents as the school has seats. Suppose also that each school gives priority to its residents and breaks ties within each group by one lottery, drawn after the reports are submitted. Then, under DA, the reports of a student that are dominant before the draw are exactly the rankings that are truthful through her district school.*

The proof, also in Appendix C, applies Proposition 5.3 to the market that each draw produces. These markets are balanced, because the districts together contain all the students and all the seats. In the Chen–Sönmez data, a ranking is therefore dominant exactly when it agrees with the truth down to the student's district school. By the symmetry of Section 2.3, the lemma also says what a dominant report reveals in this design. A ranking that is dominant before the draw shows the student's true ranking down to her district school, and nothing below it.

6.2 The 144 rankings

Table 1 sorts all 144 rankings of the DA treatment into four groups. The groups do not overlap, and together they contain every ranking.

Of the 58 rankings that differ from the truth, 38 gave the student the same school as the truth against the rankings that the others actually submitted, but only 6 of these 38 are dominant. The other 32 were harmless only because of what the other students happened to submit. Against some other rankings of the others, each of them would have given the student a worse school than the truth. The remaining 20 misreports changed the student's assignment.

The published study counts 52 misrepresentations under DA, because it treats a report as truthful when it agrees with the true ranking down to the district school. The six reports that differ from the truth only below the district school are the six non-truthful dominant reports. In this experiment, the published convention for partial truth therefore coincides with the comparisons that matter for incentives. The other 52 misreports switch at least one comparison that matters for the student's assignment.

Table 1: The 144 rankings in the Chen–Sönmez DA treatment

Group	Rankings	Share of all	Share of misreports
Truthful	86	59.7%	
non-truthful, dominant	6	4.2%	10.3%
Not dominant, same realized assignment	32	22.2%	55.2%
Not dominant, different realized assignment	20	13.9%	34.5%
Total	144	100.0%	100.0%

Notes. "Dominant" means dominant before the lottery is drawn (Lemma 6.2). A ranking has the same realized assignment if it gives the student the same school as the truth when the other students' rankings and the lottery are fixed at their realized values. The 6 dominant misreports also have the same realized assignment, so 38 of the 58 misreports keep it.

The 6 dominant misreports are robustly harmless. Assignment incentives give the student no reason to prefer the truth to them, so which of her dominant reports she submits is a question of selection. Section 1.1 discusses the reason behind of such behavior. The other 32 misreports that kept the realized assignment were harmless only as realized. Each of them exposed the student to a worse school against other rankings of the others. The 20 misreports that changed the assignment are those for which the risk did materialize. They cost the students who submitted them 4.65 dollars on average, in the induced payoffs of the experiment (Section S7.3 of the supplement). For the last two groups, an explanation must say why a student accepts the risk of a worse school, and models of misunderstanding, beliefs, or nonstandard preferences can provide one (Section 1.1).

The Chen–Sönmez market is balanced, so it illustrates the case without a spare seat. The capacity boundary itself calls for a direct test, which would compare otherwise identical markets with $Q = n$ and with $Q = n + 1$. Such a test would elicit preferences independently, reveal the relevant priorities before the rankings are submitted, and require complete rankings. In the market with a spare seat, Theorem 1 classifies each ranking by the guarantee rank. In the balanced market, Proposition 5.3 classifies it by decisive pairs. The comparison would show whether students use the freedom that each market leaves them, and whether their non-truthful comparisons fall where the theory says they are free.

The same classification applies to other strategy-proof mechanisms. The supplement applies it to two experiments with serial dictatorship (Section S7). In that of Li (2017), the queue is drawn after the reports are submitted. In that of Hakimov et al. (2023), students must pay to learn their own preferences.

7 Conclusion

Strategy-proof mechanisms make the truth a best report, but not necessarily the only one. This paper has asked which misreports are harmless whatever the others submit, so that they are dominant reports just as the truth is. The answer separates two senses in which a misreport can be harmless: robustly, against every combination of rankings that the others could submit, or only as realized, against the rankings that they actually submitted.

The analysis rests on the option sets of Barberà (1983). Because a strategy-proof mechanism gives a student the school that her own ranking puts highest in her menu, a ranking is dominant exactly when it keeps the true order of every decisive pair. Therefore, a comparison between two schools matters exactly when, for some rankings of the rivals, her menu contains both schools and no school that she prefers to both. The rivals' rankings determine which schools are in her menu;

the student herself still ranks every school. The dominant reports keep exactly the comparisons that matter, the decisive pairs.

Theorem 1 describes these comparisons when every student must rank every school. It covers DA, TTC, and every other strategy-proof mechanism that respects top-priority sets and is top-top consistent. Each student is guaranteed one of her schools down to her guarantee rank, which her own ranking, the priorities and the capacities determine. Every ranking that is truthful through this rank is dominant. Suppose there is at least one more seat than students, and the mechanism also gives every student her first choice whenever these first choices fit. Then these are exactly the dominant reports, in every priority profile, and they are the same under every such mechanism, although the mechanisms can give her different schools and different menus. A student with guarantee rank κ_i then has $(m - \kappa_i)!$ dominant reports, so the spare seat makes the answer exact and simple rather than unique. Without a spare seat, each of these mechanisms has markets in which the dominant reports take another form. The spare seat matters because every rival must be seated somewhere. With it, the rivals can leave two schools vacant for the student while they take every school that she prefers to both. In this way, with a spare seat, the guarantee rank measures both how secure her assignment is and how much of her ranking she must report truthfully.

The requirement to rank every school is what creates this boundary. When lists can be truncated, rivals can leave the market. Under DA, the dominant reports are the rankings that are truthful through the guarantee rank, in every market (Remark 4.15). When every school must be ranked, nobody can leave the market, and a spare seat has the effect that leaving has when lists can be short. The form of the reports is therefore part of the design, and it is a part that a mechanism controls more directly than the number of seats.

The same logic shows which features of a market determine how much of a ranking matters. A comparison between two schools can matter when the rivals can leave both schools vacant for the student while they take every school that she prefers to both. A spare seat makes this possible, and so does the freedom to leave the market. Priorities can too. Without a spare seat, under DA, a rival whom the student outranks at one of the two schools can take the place of the second vacant seat. This is why priorities determine which comparisons matter in balanced markets (Proposition 5.3). Under TTC, she can also reach a school by trading a priority, which is one reason why DA and TTC can differ without a spare seat (Section 4.7). Timing is another such feature. When priorities are drawn only after the rankings are submitted, a comparison that matters after some draw already matters before it (Proposition 6.1). Capacity has two effects. The first spare seat leaves enough seats for the rivals wherever it is, even at a school that every student ranks last. But an additional seat at a school that she likes also makes that school harder to fill. So once there is a spare seat, more seats and better priorities keep the guarantee rank or lower it, and they leave the student at least as many dominant reports (Section S2 of the supplement). A spare seat is therefore the condition under which the answer is exact.

These results also suggest a way to ask students for less. DA mechanism can ask each student for her favorite school, then for her next favorite, and so on. It can stop as soon as the schools named so far are not blockable, that is, at her guarantee rank. The clearinghouse computes this rank from her answers, the priorities, and the capacities, so the student needs neither to compute it nor to know the priorities. It then completes her ranking with a default order and runs DA on the completed rankings. If the answers are kept private until all of them are in, truthful answers are dominant. When every student answers truthfully, DA produces the same matching as it would with complete truthful rankings (Proposition S6.2 of the supplement). Remark 4.7 gives the reason, and it shows that the same holds for TTC. With a spare seat, no procedure of this kind can stop earlier and still be sure that the completed ranking is a dominant report for the student (Proposition S6.1). In the Chen–Sönmez design, where the lottery is drawn after the rankings are submitted, a safe procedure

can stop once the student has named her district school, and in general not before (Lemma 6.2). Asking for one school at a time resembles sequential and obviously strategy-proof implementations (Li, 2017). The procedure keeps dominance, however, rather than adding obviousness: for general priorities, no mechanism that implements stable matchings is obviously strategy-proof (Ashlagi and Gonczarowski, 2018).

The Chen–Sönmez experiment shows why the distinction between the two senses of harmlessness matters in data. Of 144 rankings, 58 differ from the truth. Of these 58, 38 kept the student's assignment at the realized rankings, but only 6 are dominant. In this experiment, counting realized harmlessness would therefore overstate the freedom that strategy-proofness leaves students about sixfold. Empirical work on strategy-proof mechanisms should keep literal misreporting, harmlessness at the realized rankings, and dominance apart, because each requires a different explanation. The results also tell such work what a submitted ranking reveals. In a market with a spare seat, a ranking that is a dominant report reveals the student's preferences down to her guarantee rank, and nothing below it (Corollary 4.6).

References

Abdulkadiroğlu, A. and Sönmez, T. (2003). School choice: A mechanism design approach. *American Economic Review*, 93(3):729–747.

Arribillaga, R. P. and Massó, J. (2016). Comparing generalized median voter schemes according to their manipulability. *Theoretical Economics*, 11(2):547–586.

Artemov, G., Che, Y.-K., and He, Y. (2023). Stable matching with mistaken agents. *Journal of Political Economy Microeconomics*, 1(2):270–320.

Ashlagi, I. and Gonczarowski, Y. A. (2018). Stable matching mechanisms are not obviously strategy-proof. *Journal of Economic Theory*, 177:405–425.

Barberà, S. (1983). Strategy-proofness and pivotal voters: A direct proof of the Gibbard–Satterthwaite theorem. *International Economic Review*, 24(2):413–417.

Bucher, S. F. and Caplin, A. (2021). Inattention and inequity in school matching. Working Paper 29586, National Bureau of Economic Research.

Chen, R., Katuščák, P., Kittsteiner, T., and Kütter, K. (2024). Does disappointment aversion explain non-truthful reporting in strategy-proof mechanisms? *Experimental Economics*, 27(5):1184–1210.

Chen, Y. and Sönmez, T. (2006). School choice: An experimental study. *Journal of Economic Theory*, 127(1):202–231.

Decerf, B., Haeringer, G., and Van der Linden, M. (2024). Incontestable assignments. Working paper, version of June 17, 2024.

Decerf, B. and Van der Linden, M. (2021). Manipulability in school choice. *Journal of Economic Theory*, 197:105313.

Dreyfuss, B., Heffetz, O., and Rabin, M. (2022). Expectations-based loss aversion may help explain seemingly dominated choices in strategy-proof mechanisms. *American Economic Journal: Microeconomics*, 14(4):515–555.

Dubins, L. E. and Freedman, D. A. (1981). Machiavelli and the Gale–Shapley algorithm. *American Mathematical Monthly*, 88(7):485–494.

Fack, G., Grenet, J., and He, Y. (2019). Beyond truth-telling: Preference estimation with centralized school choice and college admissions. *American Economic Review*, 109(4):1486–1529.

Gale, D. and Shapley, L. S. (1962). College admissions and the stability of marriage. *American Mathematical Monthly*, 69(1):9–15.

Gonczarowski, Y. A., Heffetz, O., and Thomas, C. (2026). Strategyproofness-exposing descriptions of matching mechanisms. arXiv:2209.13148, version 4.

Hakimov, R., Kübler, D., and Pan, S. (2023). Costly information acquisition in centralized matching markets. *Quantitative Economics*, 14(4):1447–1490.

Hassidim, A., Marciano, D., Romm, A., and Shorrer, R. I. (2017). The mechanism is truthful, why aren't you? *American Economic Review*, 107(5):220–224.

Hassidim, A., Romm, A., and Shorrer, R. I. (2021). The limits of incentives in economic matching procedures. *Management Science*, 67(2):951–963.

Katuščák, P. and Kittsteiner, T. (2025). Strategy-proofness made simpler. *Management Science*, 71(9):7560–7578.

Kesten, O. and Kurino, M. (2019). Strategy-proof improvements upon deferred acceptance: A maximal domain for possibility. *Games and Economic Behavior*, 117:120–143.

Kloosterman, A. and Troyan, P. (2026). Rankings-dependent preferences: A real goods matching experiment. *Games and Economic Behavior*, 155:167–191.

Koutout, K., Dustan, A., Van der Linden, M., and Wooders, M. (2021). Mechanism performance under strategy advice and sub-optimal play: A school choice experiment. *Journal of Behavioral and Experimental Economics*, 94:101755.

Kumano, T. and Watabe, M. (2011). Untruthful dominant strategies for the deferred acceptance algorithm. *Economics Letters*, 112(2):135–137.

Kumano, T. and Watabe, M. (2012). Dominant strategy implementation of stable rules. *Games and Economic Behavior*, 75(1):428–434.

Leshno, J. D. and Lo, I. (2021). The cutoff structure of top trading cycles in school choice. *Review of Economic Studies*, 88(4):1582–1623.

Li, S. (2017). Obviously strategy-proof mechanisms. *American Economic Review*, 107(11):3257–3287.

Marutani, K. (2018). Gaming the deferred acceptance when message spaces are restricted. *Mathematical Social Sciences*, 93:153–158.

Meisner, V. (2023). Report-dependent utility and strategy-proofness. *Management Science*, 69(5):2733–2745.

Morrill, T. (2015). Two simple variations of top trading cycles. *Economic Theory*, 60(1):123–140.

Pápai, S. (2000). Strategyproof assignment by hierarchical exchange. *Econometrica*, 68(6):1403–1433.

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

Roth, A. E. (1982). The economics of matching: Stability and incentives. *Mathematics of Operations Research*, 7(4):617–628.

Roth, A. E. (1986). On the allocation of residents to rural hospitals: A general property of two-sided matching markets. *Econometrica*, 54(2):425–427.

Sirguiado, C. J. and Torres-Martínez, J. P. (2024). Strategic behavior in one-to-one matching markets without outside options. *Games and Economic Behavior*, 148:385–397.

Szpilrajn, E. (1930). Sur l'extension de l'ordre partiel. *Fundamenta Mathematicae*, 16:386–389.

Appendix

This appendix proves the results of the main text whose proofs are deferred. Appendix A proves Lemmas 4.5, 4.10, and 5.1, Appendix B proves Proposition 5.3 and Corollary 5.4, and Appendix C proves Proposition 6.1 and Lemma 6.2. Further results are collected in a separate supplement.

A Proofs for the Capacity Boundary

A.1 DA and TTC have the four properties

This appendix proves Lemma 4.5. Strategy-proofness is known for both mechanisms. The other three properties are checked one at a time, first for DA and then for TTC. Throughout, fix a profile of rankings $\hat{P}$ and a student s_j, and let A be the set of the first h schools of $\hat{P}_j$.

Strategy-proofness. DA is strategy-proof for students by Dubins and Freedman (1981) and Roth (1982), and TTC is strategy-proof by Pápai (2000), for one seat per school, and by Abdulkadiroğlu and Sönmez (2003), for schools with several seats.

DA, respect for top-priority sets. Suppose that s_j does not receive a school of A. She proposes in the order of $\hat{P}_j$, so every school of A has rejected her. A school rejects a student only when it already holds as many students it ranks above her as it has seats, and the students it holds only improve afterward. At the end, every school $c \in A$ therefore holds q_c students who outrank s_j at c. These students are distinct and fill every seat of A, so A is blockable for s_j.

DA, top-top consistency. Let (s_j, c) be a top-top pair. Every stable matching assigns s_j to c: otherwise s_j and c would form a blocking pair, because s_j ranks c first and c ranks her above every other student. A matching that assigns s_j to c is stable if and only if what it does with the other students is stable in the market without s_j and one seat of c. This is because s_j already has her first choice and c ranks her above everyone else. This gives a one-to-one correspondence between the stable matchings of the two markets that leaves every other student's assignment unchanged. It therefore maps the student-optimal stable matching of one market to that of the other.

DA, feasible first choices. If no school is ranked first by more students than it has seats, no student is rejected in the first round. DA then stops, with every student at her first choice.

For TTC, two standard facts are used. First, at every step the cycles are disjoint, and removing one cycle leaves every other cycle intact: its students still point to the same schools, and these schools still point to the same students. A cycle therefore persists until it is removed. By induction on the number of students, the outcome does not depend on the order in which cycles are removed, or on whether they are removed together or one at a time. Second, TTC is Pareto efficient (Abdulkadiroğlu and Sönmez, 2003). On every cycle, each school points to the next student on the cycle, and each school loses exactly one seat, to the student who points to it.

TTC, respect for top-priority sets. Suppose that s_j receives a school outside A, or nothing. Remove the cycles one at a time. When s_j is assigned, she points to her favorite school with a free seat under $\hat{P}_j$, so at that moment every school of A is full. If she is never assigned, TTC ends with no free seat at all. Either way, every seat of every school of A is used up while s_j is still present, on cycles that do not contain her. When a seat of a school $c \in A$ is used up, c lies on the removed cycle and points to the next student z on it. Because s_j is still present and z is the highest-priority student present at c, z outranks s_j at c, and z is assigned on that cycle. Charge the seat to z. A student lies on at most one removed cycle, and the schools of a cycle point to different students, so different seats are charged to different students. These students fill every seat of A, each at a school where she outranks s_j, so A is blockable for s_j.

TTC, top-top consistency. If (s_j, c) is a top-top pair, then at the first step s_j points to c and c points to s_j. By the first fact, this cycle can be removed first, and the procedure that remains is TTC in the market without s_j and one seat of c.

TTC, feasible first choices. If the first choices fit into the seats, the assignment that gives every student her first choice Pareto dominates every other assignment. TTC is Pareto efficient, so it gives every student her first choice.

A.2 Proof of Lemma 4.10

The idea. Example 4.9 shows what can go wrong with an arbitrary covering under TTC. The rival chosen to fill a seat may not hold the strongest claim to it, and the student can then obtain the seat through a trade. The construction avoids this by choosing rivals who can be served one at a time, each of whom has the highest priority at her school among the students still present when her turn comes. By respect for top-priority sets, such a rival receives the school she ranks first. By top-top consistency, she can then be removed with her seat, which puts the next rival in the same position. Two things must be arranged for this to work, and the first two steps of the proof arrange one each. First, the rivals with the highest priorities at the schools of V must be among those chosen. Otherwise an unchosen rival would stand between a chosen one and the top of the priority order. Step 1 arranges this. Second, the chosen rivals must be matched to the seats so that each school is served first to its highest-priority chosen rival. Step 2 arranges this. Step 3 then serves the rivals in order and checks the conclusion.

Let $S = \sum_{c \in V} q_c$. If $V = \varnothing$, there is nothing to prove. Otherwise, fix a covering K of V by distinct blockers of s_i: each seat of V is matched to a different rival who has higher priority than s_i at its school.

Step 1: a closed covering. This step finds S rivals who can fill every seat of V, each at a school where she outranks s_i. They are chosen so that whenever one of them sits at a school, every rival with higher priority at that school is also chosen. Consider the market with the schools of V and all the rivals, in which every rival finds every school of V acceptable and each rival matched by K ranks her K-school first. Let μ be the outcome of school-proposing DA in this market. Each school $c \in V$ is accepted by its q_c rivals under K whenever it proposes to them, because they rank it first, so it never has to propose below the lowest-priority of them. Hence c ends full, and it holds only rivals with higher priority than s_i. Because μ is stable in this market, it is *closed*. Suppose that a rival y has higher priority at c than some rival x that μ assigns to c. Then μ assigns y to a school of V, since otherwise y, who finds c acceptable, and c would form a blocking pair. Let Y be the set of the S rivals assigned by μ.

Step 2: an order of the seats. This step reassigns the rivals in Y to the seats of V. It also puts them in an order in which each, when her turn comes, has the highest priority at her school among the rivals still present. Run the following elimination, in which every rival and every seat of V is present at the start. Each present member of Y points to her school under μ, and each school of V with a present seat points to its highest-priority present rival. Find a cycle, give each student on it the school that points to her, record these pairs, and remove the students of the cycle and one seat of each school on it. Then repeat.

This elimination never gets stuck. Each school on a removed cycle loses one seat and one of the rivals that μ assigns to it, namely the one who points to it. So every school of V keeps as many present seats as present rivals assigned to it by μ. Hence, if c has a present seat, some rival x assigned to c by μ is still present. The rival that c points to is then either x or someone with higher priority than x at c. By closedness, she belongs to Y. Every pointer therefore stays among the present members of Y and the schools with present seats, so a cycle always exists. The elimination

ends when all S rivals and all seats of V have been recorded. When a pair (x, c) is recorded, x has the highest priority at c among the rivals present at that moment. She also has higher priority than s_i at c, because she ranks weakly above a present rival assigned to c by μ, who is a blocker.

Step 3: the rivals keep their seats. Let each member of Y rank her recorded school first, and fix any rankings of the other students, s_i included. Take the recorded pairs in the order in which they were removed, taking the pairs of one cycle in any order. When the turn of a pair (x, c) comes, consider the market from which the students of the earlier pairs have been removed, each with one seat of her school. The students present are s_i, the rivals outside Y, and the members of Y not yet reached. When (x, c) was recorded, x had the highest priority at c among the rivals then present, and this set contains every rival still present. She also outranks s_i at c. So (x, c) is a top-top pair, and c still has a seat, because it had a present seat when the pair was recorded. By respect for top-priority sets, x receives c. By top-top consistency, every other student present receives the school she receives in the market without x and one seat of c. Chaining these equalities from the full market, every member of Y receives her recorded school. These rivals fill every seat of V, so s_i receives no school of V, whatever she reports.

A walk-through of Step 1. In Example 4.9, $V = \{c_1\}$ has one seat, and its blockers are s_1 and s_2. Start from the covering K that uses s_2, the covering that fails under TTC. In Step 1, school c_1 proposes to its highest-priority rival, s_1, who accepts, because c_1 is the only school in this small market and she finds it acceptable. So μ seats s_1 at c_1, and $Y = \{s_1\}$. This covering is closed, because nobody has higher priority than s_1 at c_1. In Step 2, s_1 points to c_1 and c_1 points to s_1, so the pair (s_1, c_1) is recorded. In Step 3, s_1 ranks c_1 first, and because she has the highest priority there, she receives c_1 under DA, TTC, and every other mechanism in the class. Step 1 has replaced s_2 by s_1: under TTC, s_3 could obtain the seat of s_2 through a trade, but nobody can obtain the seat of s_1 that way.

A walk-through of Step 2. Step 2 matters when a selected rival does not have the highest priority at her own school. Take four students, with s_4 as the focal student, and let V consist of two schools: c_1, with two seats and priorities $s_1 \succ s_2 \succ s_3 \succ s_4$, and c_2, with one seat and priorities $s_2 \succ s_1 \succ s_3 \succ s_4$. Start from the covering that seats s_1 at c_2 and seats s_2 and s_3 at c_1. Step 1 returns the same covering, which is closed. But s_1 does not have the highest priority at c_2, and s_2 does not have it at c_1, so neither of them forms a top-top pair with her school. In Step 2, s_1 points to c_2, which points to s_2, who points to c_1, which points back to s_1. This cycle records the pairs (s_2, c_2) and (s_1, c_1): the two rivals exchange their seats, and each now has the highest priority at her school. Once they are removed, s_3 points to c_1, which points to s_3, and the pair (s_3, c_1) is recorded. In Step 3, s_2 and s_1 each form a top-top pair with their recorded schools, and once they have been removed, so does s_3.

A.3 Proof of Lemma 5.1

The reduced market has $n - 1$ students and $Q \geq n$ seats, so some school v has an empty seat at η. Let each rival rank her school under η first and, if that school is not v, rank v second. Complete the rest of each ranking arbitrarily. The matching η is stable, because every student receives her first choice.

Under responsive choice and complete acceptability, the rural hospitals theorem implies that a school with an empty seat in one stable matching has exactly the same students in every stable matching (Roth, 1986). So v has the same students in every stable matching. Consider any student who is not at v under η. If another stable matching moved her away from her first choice, she would prefer v to her new school. Since v would still have an empty seat, she and v would form a blocking pair. The students at v stay there, because the set of students at v is the same in every stable

matching. Hence nobody moves, and η is the unique stable matching.

For the converse, every realizable matching leaves a seat empty, because $Q > n - 1$. An unassigned rival and a school with an empty seat would form a blocking pair, since every school is acceptable and schools prefer any student to an empty seat. A stable matching therefore seats every rival.

B Proofs for Balanced Markets

This appendix proves Proposition 5.3 and Corollary 5.4. Throughout, the mechanism is DA and $Q = n$.

Proof of Proposition 5.3. By Lemma 5.1, the realizable matchings are exactly the ways of seating all $n - 1$ rivals. With $Q = n$, each of them leaves exactly one seat empty. By Lemma 5.2, the pair (r, t) is decisive if and only if some such matching η satisfies (W1) and (W2).

Only if. Let η seat all the rivals and satisfy (W1) and (W2). By (W1), the rivals that η assigns to the schools of $H_i(r-1)$ fill all their seats, each at a school where she has higher priority than s_i, so they form a covering of $H_i(r-1)$. Only one seat is empty, so at most one of $c_{i,r}$ and $c_{i,t}$ has it. The other one, call it $c_{i,x}$, is full, and by (W2) it holds a rival w whom s_i outranks there, that is, $w \in L_i(c_{i,x})$. This rival sits at $c_{i,x}$, which is not in $H_i(r-1)$, so the covering leaves her unused.

If. Suppose that a covering of $H_i(r-1)$ by distinct blockers leaves unused a rival $w \in L_i(c_{i,x})$, where x is r or t, and let y be the other one. Let S be the number of seats of $H_i(r-1)$. Build a matching η in four steps. Seat the rivals of the covering as it does, and seat w at $c_{i,x}$. Keep one seat of $c_{i,y}$ empty, and seat the remaining $n - S - 2$ rivals in the remaining $n - S - 2$ seats. This matching seats every rival, so it is realizable. It satisfies (W1) by construction, and it satisfies (W2), because $c_{i,y}$ has an empty seat and $c_{i,x}$ holds w. So (r, t) is decisive.

Finally, a covering of $H_i(r-1)$ that leaves w unused exists if and only if $H_i(r-1)$ is blockable, that is, $r \le \kappa_i$, and w is not essential for rank r. This gives (5). □

Proof of Corollary 5.4. (i) Suppose that $L_i(c_{i,r}) \not\subseteq \mathrm{Ess}_i(r)$ for every $r \le \kappa_i$. Then (5) makes every pair (r, t) with $r \le \kappa_i$ decisive. By Lemma 4.11(i), no other pair is decisive, so the decisive pairs are exactly the top-segment order O_{κ_i}, and Corollary 2.6 gives the dominant set.

(ii) If s_i has the lowest priority at every school, every set $L_i(c)$ is empty, so no pair is decisive and every report is dominant. Conversely, suppose that s_i outranks some rival w at a school $c_{i,t}$. The empty set $H_i(0)$ needs no blockers, so w is not essential for rank 1. If $t \ge 2$, the pair $(1, t)$ is therefore decisive by (5). If $t = 1$, the pair $(1, 2)$ is decisive for the same reason. In both cases, some report is not dominant.

(iii) With one priority order, every school has the same blockers, the $p - 1$ rivals ranked above s_i, and the same lower-priority rivals, the $n - p$ rivals ranked below her. The seats of $H_i(h)$ can be filled by distinct blockers if and only if they number at most $p - 1$, which gives the stated κ_i. The essential rivals are always among the rivals above s_i, so no set $\mathrm{Ess}_i(r)$ meets any set $L_i(c)$. If $p < n$, every set $L_i(c)$ is nonempty, and part (i) applies. If $p = n$, part (ii) applies. □

C Dominance before a Lottery

This appendix proves Proposition 6.1 and Lemma 6.2.

Proof of Proposition 6.1. Write u_i for the utility function and $\pi(\omega) > 0$ for the probability of state ω. Take a report $\hat{P}_i$ that is dominant in every state. Against every profile $\hat{P}_{-i}$ of rival reports, and

in every state, it gives the same school as the truth, so it gives the same expected utility. The truth is a best report in every state, by strategy-proofness, so its expected utility is at least that of every other report. Hence $\hat{P}_i$ is dominant before the draw.

Conversely, take a report $\hat{P}_i$ that is not dominant in some state ω^*. By Lemma 2.2, there is a profile $\hat{P}^*_{-i}$ of rival reports against which $\hat{P}_i$ gives, in state ω^*, a school different from the truthful one. Against $\hat{P}^*_{-i}$, the truth is at least as good as $\hat{P}_i$ in every state, by strategy-proofness, and strictly better in state ω^*, because P_i is strict. As $\pi(\omega^*) > 0$, the truth gives a strictly higher expected utility than $\hat{P}_i$ against $\hat{P}^*_{-i}$, so $\hat{P}_i$ is not dominant before the draw.

Finally, by Proposition 2.5, a ranking is dominant in state ω if and only if it keeps the true order of every decisive pair of that state. It is therefore dominant in every state if and only if it keeps the true order of every pair in $\bigcup_\omega E_i^{\varphi^\omega}(P_i)$. The rankings that do so are exactly the linear extensions of the order that these pairs generate. Neither the probabilities nor u_i appear in this description. □

Proof of Lemma 6.2. Fix a student s_i and her true ranking, and let $e = c_{i,d}$ be her district school, at true rank d. Every draw produces a balanced market, as all districts together contain all the students and all the seats. By Proposition 6.1 and Corollary 2.6, it is enough to show that the decisive pairs of all draws together are exactly the pairs of the top-segment order O_d.

No pair outside O_d is ever decisive. In every draw, the only students with higher priority than s_i at e are other residents of her district, and there are at most $q_e - 1$ of them. So e is individually guaranteed to her (Remark 4.4), and her guarantee rank is at most d in every draw. By Lemma 4.11(i), every menu of s_i in that draw contains one of her top d schools, so every decisive pair (r, t) has $r \leq d$.

Every pair of O_d is decisive in one draw. Consider the draw that puts s_i first in the lottery. At e, nobody has higher priority than she does. At every other school c, the students with higher priority than s_i are exactly the residents of c, and there are as many of them as c has seats. Every other student has lower priority than s_i there. Fix true ranks $r < t$ with $r \leq d$. The schools she likes more than $c_{i,r}$ are all different from e, so their residents fill their seats, and these residents are their only blockers. So $r \leq \kappa_i$ in this draw, the covering of $H_i(r-1)$ is unique, and the essential rivals for rank r are exactly the residents of the schools in $H_i(r-1)$. Now take the residents of $c_{i,r}$ if $c_{i,r} \neq e$, and the residents of $c_{i,t}$ if $c_{i,r} = e$. They are not essential for rank r, because their school is not in $H_i(r-1)$. And s_i outranks them at the other school of the pair: either they are not residents of that school, or that school is e. By Proposition 5.3, the pair (r, t) is decisive in this draw. □