School Composition and Improvements over Deferred Acceptance*

Jorge Arenas†
University of Southern California

October 2026

Abstract

School choice mechanisms determine not only where students study but also who their classmates are. Deferred acceptance (DA) protects families in this regard through local non-bossiness: a student who keeps her school cannot alter its cohort by misreporting her ranking. This paper examines whether the protection persists under reforms designed to correct DA's inefficiency, and concludes that it does not. Any mechanism that yields a Pareto improvement over DA whenever one exists fails to be locally non-bossy in every problem with a separating double rotation, a configuration of priorities and capacities that arises under random priorities in most large markets. Efficiency-adjusted deferred acceptance loses the protection once two identified students consent and, in large markets, with probability approaching one when a fixed fraction of students chosen independently of priorities consents. We also show how the consequences of this manipulation for segregation depend on how student types relate to priorities.

Keywords: school choice; school composition; Pareto efficiency; local non-bossiness.
JEL: C78, D47, D82, I28.

1 Introduction

A centralized school choice system assigns students to schools, and in doing so it also decides who their classmates will be. That second decision matters for two reasons.

*I thank Josué Ortega and Juan Pablo Torres-Martínez for helpful comments and suggestions.
†jiarenas@usc.edu

Families care about it, because what they look for in a school includes who its students will be, and the rankings they submit respond to that (Abdulkadiroğlu et al., 2020; Beuermann et al., 2023; Burgess et al., 2015; Rothstein, 2006). Districts act on it, through catchment areas, reserved seats and the priority rules themselves (Correa et al., 2022; Cox et al., 2025), and a growing theoretical literature asks how the assignment mechanism itself shapes the composition and the segregation of schools (Escobar & Huerta, 2026; Ortega et al., 2026). A mechanism should then be judged not only by the school each student receives, but also by whose report decides the group she joins.

The mechanism most districts use is student-proposing deferred acceptance (DA). It is stable, so no student is denied a seat held by somebody with a weaker claim on it, and it is strategy-proof, so no student obtains a better school by misreporting her ranking. But it is not efficient: the matching it produces can be Pareto dominated for the students by an unstable one. A large literature builds mechanisms that recover those gains without leaving any student worse off than DA does in her ranking of schools, and studies what such a repair has to give up. We ask what it does to the composition of schools, and in particular whether it gives a family some control over who its child's classmates will be.

Our question takes a precise form because of two known results. A mechanism is *non-bossy* if a student who submits a different ranking and keeps her school cannot change anybody else's assignment. DA is bossy: a student can misreport, keep her own school and still rearrange the rest of the market. Under DA the property holds exactly when priorities are acyclic in the sense of Ergin (2002), which is also exactly when DA is efficient, so non-bossiness is available only where there is nothing to repair. Duque-Rosas, Pereyra, and Torres-Martínez (2026) isolate the part of the property that survives everywhere: DA is *locally non-bossy*, which means that a student who submits a different ranking and keeps her school cannot change the set of students assigned to that school. This is what a district can guarantee a family under DA, regardless of its priority structure. A family that keeps its school cannot rearrange its own cohort through the list it submits.[1] We ask whether that promise survives the repair of DA's inefficiency.

It does not. Our first result gives a school choice problem in which no mechanism that implements a Pareto improvement over DA whenever one exists is locally non-bossy (Theorem 1). The class is large. It contains every Pareto-efficient mechanism that weakly dominates DA, among them efficiency-adjusted deferred acceptance (EADA) when every student consents, and it also includes mechanisms that fall well short of efficiency, such as one that executes a single improving exchange whenever DA leaves one available, or one

[1]The promise is about the family's own report and not about the reports of others: under DA a family can see its classmates change because somebody else reported differently, and the problem Γ^4 of the Appendix shows it.

that implements an improvement whose priority violations are all justifiable in the sense of Ortega and Arribillaga (2026). The problem has two preference profiles that differ only in the order of the two schools one student ranks second and third. At each profile DA admits exactly one Pareto improvement, so every mechanism in the class has to select it, and the two selections give that student the same school with a different classmate. The conflict does not come from unacceptable schools, from unassigned students, or from the designer's freedom to choose among several corrections.

The rest of the paper asks which problems have this feature, and answers with conditions on priorities and capacities, that is, on what a district designs. The conflict needs a *rotation* through a school with several seats: a chain of denied claims that closes on such a school, with each seat along the chain held by a student with a stronger claim. When no such chain exists, every DA-dominating mechanism is locally non-bossy, efficient or not (Theorem 2).[2] The conflict arises when one student, the *manipulator*, has two such chains into the same school with several seats, the *contested school*, each starting from a different school. We call them her two *routes*. Along each route, the manipulator takes a seat at the contested school and one of its students, a *rival*, moves to the seat the manipulator leaves. For the two routes to give the manipulator different classmates, each route must be closed while the manipulator travels the other, and priorities can close a route in two ways. In a *guarded* route, the seat the rival would move into is held by a third student, the *guard*, who has lower priority than the manipulator and higher priority than the rival at that school. Only the manipulator can displace the guard, so the rival stays at the contested school unless the manipulator arrives through that route. In a *direct* route, the rival already holds that seat, and it is the manipulator's arrival that sends her back to the contested school, where she takes the seat of a classmate with lower priority, who is consequently crowded out. When both routes are guarded, the manipulator's report decides which of the two rivals remains her classmate, while when one route is direct, it decides whether the crowded-out classmate keeps her seat. We call this configuration a *separating double rotation*, and we show that it produces the impossibility for every mechanism in the class (Theorem 3).

Two further results say how large the region of conflict is. The configuration is easy to meet: it is enough that the student with the lowest priority at a school with at least two seats rank above four students at each of two other schools, in a market where those three schools together leave at least one student without a seat. If the priority orders

[2]The condition is not a restatement of acyclicity in the sense of Ergin (2002). We exhibit a problem with cyclic priorities, an oversubscribed school and Pareto-inefficient DA outcomes that still falls in the compatible region, because the improvements its priority reversal makes available never reach the school with several seats.

of the three schools are drawn independently and uniformly, that condition holds with probability at least $(1 - 4/n)^2$ in a market with n students. The bound does not depend on the number of schools or on their capacities, so the conflict becomes the rule as markets grow, also when capacities grow with them. We do not claim the converse. Our results do not characterize the problems in which the three properties are compatible, and they do not say that priorities that differ only slightly are safe: padding the problem Γ^1 of the proof of Theorem 1 with students who find no school acceptable, and who sit in the same positions at the bottom of every order, leaves the impossibility intact and makes the orders agree on all but five students.

We then look at the repair that has been studied most and that a public authority has considered. EADA repairs DA by asking students to waive priorities whose violation cannot affect their own assignment, and the Flemish Ministry of Education took the first steps towards implementing it in 2019 (Cerrone, Hermstrüwer, & Kesten, 2024). Its full-consent version belongs to the class, so it inherits the impossibility. The sharper statement is about consent. Kesten (2010) shows that waiving is costless for the student who waives, in the sense that it cannot change her own school. We show that in every problem with a separating double rotation two such waivers are enough to make the mechanism locally bossy, and we name the two students: the two guards when both routes are guarded, and the guard of the guarded route together with the classmate crowded out of the contested school when one route is direct. If nobody consents the mechanism is DA and the protection holds. The waiver that costs its owner nothing is what lets another student change her own classmates while keeping her school.

Since a district does not choose who consents, we also ask what happens when the consenting students are not determined in advance. If a fixed fraction of students consents, the consenting students are unrelated to priorities, and the three schools involved leave a growing number of students without a seat, EADA is locally bossy with probability approaching one. Neither statement depends on the details of EADA. At the profiles we construct, the improvement over DA is unique and violates only the priorities of students who consented, so both statements extend to every consent-based mechanism that is efficient whenever possible in the sense of Ortega and Arribillaga (2026).

While the manipulator gains or loses a specific classmate, a district monitoring school composition evaluates diversity through aggregate student types, such as income, ethnicity, or prior achievement. Ortega et al. (2026) show that when priorities are stratified by type, so that every school ranks all students of one type above all students of the next, every matching that dominates DA admits the same number of students of each type to each school as DA does. We show that the invariance also holds across reports: under stratified

priorities, no DA-dominating mechanism, efficient or not, lets a student who keeps her school change the number of students of each type in it, and the bossiness we identify is confined to identities within a type. Stratification is a strong assumption, however, and when types are unrelated to priorities the conclusion reverses. At the profiles we build from a separating double rotation, the two cohorts of the manipulator's school differ in exactly one student, and with uniform priorities and types fixed in advance, as long as no type contains almost every student, the manipulator can choose between classmates of different types with probability approaching one.

The deviation we exhibit is also outside what the usual criteria measure. Since Abdulkadiroğlu, Pathak, and Roth (2009) showed that no strategy-proof mechanism can Pareto dominate DA, that literature has weakened strategy-proofness instead of dropping it, and asks whether EADA is obviously manipulable, regret-free, maxmin optimal or immune to simple misreports. All of these criteria ask whether a student can improve *her own* assignment. At the profiles we construct, every mechanism in the class gives the deviating student the same school under both reports, so she gains nothing that those criteria measure, and what her report moves is her cohort. On the domain of preferences over colleagues of Duque-Rosas, Pereyra, and Torres-Martínez (2026), where a student ranks matchings first by her school and then by her schoolmates, the same deviation is strictly profitable. The criteria listed above are stated in terms of the school a student obtains, and none of them has been formulated on that larger domain.

Related literature. The efficiency cost of priority-respecting assignment is studied by Balinski and Sönmez (1999) and Abdulkadiroğlu and Sönmez (2003), and Ergin (2002) characterizes the acyclic priority structures under which it vanishes. Efficient improvements over DA are constructed and characterized by Kesten (2010), Tang and Yu (2014), Tang and Zhang (2021) and Reny (2022), and the limits of such improvements are the subject of a growing literature: they cannot be strategy-proof (Abdulkadiroğlu, Pathak, & Roth, 2009; Alva & Manjunath, 2019; Kesten & Kurino, 2019), consistent (Doğan & Yenmez, 2020), cardinally minimally unstable (Doğan & Ehlers, 2021), regret-free truth-telling once consent decisions are not part of the report (Chen & Möller, 2024), or immune to manipulations from below (Afacan et al., 2026); they cannot repair DA's rank distribution or the welfare of its worst-off student, and under stratified priorities they cannot change how many students of each type each school admits, so they cannot reduce segregation measured in those terms (Ortega et al., 2026). Additionally, they cannot improve many students at once while being minimally unstable in the set-inclusion sense (Knipe & Ortega, 2025). Ortega and Arribillaga (2026) ask which of the priority violations that an

improvement requires can be justified, show that an improvement whose violations are all justifiable exists whenever DA is inefficient, and study what consent-based mechanisms can achieve. Our result adds a requirement of a different kind, about the identity of a student's classmates with her own school held fixed rather than about the school she obtains.

Two papers share the logical form of our conditions. Doğan and Ehlers (2021) characterize the priority profiles that admit a minimally unstable Pareto improvement over DA by conditions on Ergin cycles, trading an axiom on improvements against a restriction on priorities, as we do; their axiom concerns the instability the improvement carries and ours concerns whose composition it changes, and their characterization is stated for unit capacities, which is the case in which our question is empty. Ortega et al. (2026) show that under stratified priorities the number of students of each type at each school is the same under every rule that dominates DA, and observe that individual identities may still vary. That is the count-level counterpart of our object: what their result fixes is how many, what we study is who, and the distinction is exactly what makes the question non-trivial. Section 7 connects the two. Their invariance survives a change in the report of a single student, so that under stratified priorities the bossiness we identify moves identities and never numbers, and it fails once types and priorities come apart.

Non-bossiness goes back to Satterthwaite and Sonnenschein (1981); see Thomson (2016) for a survey and Pápai (2000) for its relation with group strategy-proofness. Raghavan (2020) reads bossiness as the ability of an agent to influence the welfare of another without changing her own, studies the structure of that influence when non-bossiness is out of reach, and shows that the consenting students that EADA considers are students who have influence under DA. Kojima (2010) shows that stability and non-bossiness are incompatible when both sides report. Afacan and Dur (2017) show that the school-proposing mechanism is the only standard one that is harmless, in the sense that no student can misreport in a way that leaves her no worse off and makes another student worse off, and Moreno-García and Torres-Martínez (2026) show that immunity to coalitional manipulation requires strategy-proofness and non-bossiness jointly, which is why the two ends of the stable lattice satisfy one property each. Local non-bossiness is what the student-optimal end retains unconditionally, and this paper asks what an efficiency adjustment does to it. The paper closest to ours is Duque-Rosas, Pereyra, and Torres-Martínez (2026), who introduce the property, prove that DA satisfies it, characterize DA with it, and show how it governs incentives when students care about their schoolmates. They do not ask whether Pareto improvements over DA preserve it.

Organization. Section 2 presents the model. Section 3 records what an improvement over DA can do and turns it into restrictions on local non-bossiness. Section 4 contains the impossibility. Section 5 bounds the problems that generate it and shows that they are the typical case. Section 6 studies EADA, consent and preferences over colleagues. Section 7 studies the composition of schools by type, and Section 8 concludes. Omitted proofs are in the Appendix.

2 Model

There is a finite set of students I and a finite set of schools C. Each school $c \in C$ has capacity $q_c \in \mathbb{N}$ and a strict priority order $\succ_c$ over I, and we write $q = (q_c)_{c \in C}$ for the capacity vector and $\succ = (\succ_c)_{c \in C}$ for the priority profile.[3] Each student $i \in I$ has a strict preference relation P_i over $C \cup \{\emptyset\}$, where $\emptyset$ is the outside option, and R_i denotes the associated weak relation. A school c is *acceptable* to i if $c \, P_i \, \emptyset$. We write $P = (P_i)_{i \in I}$ for a preference profile, and $P_{-i} = (P_h)_{h \neq i}$ for the reports of the students other than i; $\mathcal{P}$ is the unrestricted domain of preference profiles. A tuple $\Gamma = (I, C, q, \succ)$ is a *school choice problem*.[4]

An assignment $\mu : I \to C \cup \{\emptyset\}$ is *feasible* if no school exceeds its capacity. We use the term *matching* for a feasible assignment, we write $\mu_i \in C \cup \{\emptyset\}$ for the assignment of student i, and for every $x \in C \cup \{\emptyset\}$ we write

$$\mu_x = \{i \in I : \mu_i = x\}$$

for the set of students assigned to x. Thus $|\mu_c| \leq q_c$ for every $c \in C$ and $\mu_\emptyset$ is the set of unassigned students. Let $\mathcal{M}$ denote the set of matchings. Since I and C are finite, so is $\mathcal{M}$. A matching is *individually rational* if $\mu_i \, R_i \, \emptyset$ for every i, and *non-wasteful* if there is no student-school pair (i, c) such that i prefers c to μ_i while c has an empty seat. We say that school c is *full* at μ if $|\mu_c| = q_c$ and *undersubscribed* otherwise.

A matching μ is *stable* if it is individually rational and there is no pair (i, c) such that $c \, P_i \, \mu_i$ and either $|\mu_c| < q_c$ or $i \succ_c j$ for some $j \in \mu_c$. Such a pair is a *blocking pair*: student i prefers c to her assignment and can claim either an empty seat or a seat held

[3] Priorities are not preferences. Schools are objects to be consumed rather than agents, and every student is acceptable to every school: leaving a seat vacant when some student would rather occupy it is ruled out by the designer and not by the school, and $\succ_c$ only says who has the stronger claim when its seats are scarce. This is the standard convention in school choice (Abdulkadiroğlu & Sönmez, 2003; Balinski & Sönmez, 1999), and it is what non-wastefulness formalizes below.

[4] Preferences are not part of it. A mechanism is defined below on a fixed Γ and assigns a matching to every profile in $\mathcal{P}$, so each property we study is a property of a mechanism in a given school choice problem.

by a lower-priority student. Thus, stability is equivalent here to individual rationality, non-wastefulness, and absence of justified envy.[5]

The following algorithm is the standard way of computing a stable matching, and its outcome is the benchmark against which the mechanisms defined below are compared.

Algorithm 1 Student-proposing deferred acceptance, $\mathrm{DA}(P)$

1: Each student applies to her most-preferred acceptable school. Each school tentatively keeps up to q_c highest-priority applicants and rejects the rest.

2: Every rejected student applies to her most-preferred acceptable school that has not yet rejected her. Each school again keeps the q_c highest-priority students among its current tentative assignees and new applicants, rejecting the rest.

3: Repeat Step 2 until no rejection occurs. Tentative assignments then become final.

The resulting matching is denoted $\mathrm{DA}(P)$. It is the student-optimal stable matching: every student weakly prefers it to every other stable matching at P (Gale & Shapley, 1962).

The algorithm has an elementary property that we use at several points below. Once a school rejects an applicant it holds a full set of tentative assignees, and the priority of the lowest-ranked student it holds only improves thereafter. Hence, if c rejects i at some round of the run at P, then c is full at $\mathrm{DA}(P)$ and every member of $\mathrm{DA}_c(P)$ has higher priority than i at c. A rejection is permanent evidence that the school ends up full with students who have a stronger claim.

Given P, a matching ν *weakly Pareto dominates* μ if $\nu_i \, R_i \, \mu_i$ for every $i \in I$. It *Pareto dominates* μ if, in addition, $\nu_j \, P_j \, \mu_j$ for some j. A matching is *Pareto efficient* if no matching Pareto dominates it.

A *mechanism* is a map $\varphi : \mathcal{P} \to \mathcal{M}$. We write $\varphi_i(P)$ for student i's assignment and $\varphi_x(P) = \{i \in I : \varphi_i(P) = x\}$ for $x \in C \cup \{\emptyset\}$. A mechanism is *strategy-proof* if no student can obtain a strictly preferred assignment by changing only her own report: there are no $i \in I$, $P \in \mathcal{P}$, and P_i' such that $\varphi_i(P_i', P_{-i}) \, P_i \, \varphi_i(P)$. The *student-optimal stable mechanism*, denoted DA, assigns to every profile P the matching $\mathrm{DA}(P)$ produced by the algorithm above. It is strategy-proof for students (Dubins & Freedman, 1981; Roth, 1982).

Definition 1. A mechanism φ is *DA-dominating* if $\varphi(P)$ weakly Pareto dominates $\mathrm{DA}(P)$ for every $P \in \mathcal{P}$, and *DA-improving* if, in addition, $\varphi(P) \neq \mathrm{DA}(P)$ at every profile P at which $\mathrm{DA}(P)$ is not Pareto efficient. It is *efficient* if $\varphi(P)$ is Pareto efficient for every $P \in \mathcal{P}$.

[5] A student has *justified envy* at a matching when she prefers some school to her own assignment and that school admits a student with lower priority for it. Priorities are respected in this weak sense: a student may envy a classmate of a school she cannot claim, but never one who ranks below her.

The mechanism DA is DA-dominating but not efficient. Stability and Pareto efficiency for students are incompatible in general, and the student-optimal stable matching may be Pareto dominated by an unstable one (Abdulkadiroğlu & Sönmez, 2003; Balinski & Sönmez, 1999; Roth, 1982). Ergin (2002) characterizes the priority structures under which this never happens. DA-dominance says that a mechanism never leaves a student worse off than the benchmark does, and the other two properties say how much of the foregone gains it must recover. Since preferences are strict, a matching that weakly Pareto dominates $\mathrm{DA}(P)$ and differs from it Pareto dominates it. A DA-improving mechanism is therefore required only to implement some improvement over DA whenever there is one, while an efficient one must leave no improvement unexploited.

Every efficient DA-dominating mechanism is DA-improving, but the class of DA-improving mechanisms is much larger. It contains a mechanism that, whenever DA is inefficient, executes a single trading cycle among students who all prefer the seat they receive and leaves every other student where DA placed her. Such a cycle exists whenever DA is inefficient, by Lemma 1(iv) below. It also contains a mechanism that returns DA when DA is efficient and otherwise an improvement whose priority violations are all justifiable in the sense of Ortega and Arribillaga (2026), since by their Theorem 1 such an improvement exists whenever DA is inefficient. These mechanisms need not be efficient, and they are the natural benchmark for a district that wants to improve upon DA without committing to exhaust every gain. The impossibility theorems below apply to every DA-improving mechanism, and the compatibility results to every DA-dominating one.

Definition 2 (Duque-Rosas, Pereyra, and Torres-Martínez (2026))**.** A mechanism φ is *locally non-bossy* if for every $i \in I$, every $P \in \mathcal{P}$, every report P_i', and every $x \in C \cup \{\emptyset\}$,

$$\varphi_i(P) = \varphi_i(P_i', P_{-i}) = x \quad \Longrightarrow \quad \varphi_x(P) = \varphi_x(P_i', P_{-i}).$$

In words, a student who reports different preferences and is still assigned to the same school cannot change the composition of that school, and a student who reports different preferences and remains unassigned cannot change the set of unassigned students. The latter condition makes local non-bossiness non-vacuous even in one-to-one problems (Duque-Rosas, Pereyra, & Torres-Martínez, 2026). Local non-bossiness is weaker than *non-bossiness*, which requires the entire matching to remain unchanged whenever the reporting student keeps her assignment.[6]

[6]Formally, φ is non-bossy if $\varphi_i(P) = \varphi_i(P_i', P_{-i})$ implies $\varphi(P) = \varphi(P_i', P_{-i})$ (Satterthwaite & Sonnenschein, 1981). Local non-bossiness places no restriction on what the report of i does at schools other than $\varphi_i(P)$. We call a report *bossy*, or *locally bossy*, when it does what the corresponding property forbids, and

Proposition 1 (Duque-Rosas, Pereyra, and Torres-Martínez (2026), Theorem 1)**.** *The student-optimal stable mechanism is locally non-bossy.*

How much the weakening gives up is itself a question about priorities rather than about reports, since the stronger property is available only at priority structures of a particular shape. We state the condition that describes them now, both because it is the benchmark against which the rest of the paper measures its own conditions and because the objects it is built from, a priority reversal between two schools and the students who fill their seats, reappear throughout.

Definition 3 (Ergin (2002))**.** An *Ergin cycle* consists of distinct schools $a, b \in C$ and distinct students $i, j, k \in I$ such that

$$i \succ_a j \succ_a k \qquad \text{and} \qquad k \succ_b i,$$

together with disjoint sets $N_a, N_b \subseteq I \setminus \{i, j, k\}$ with $|N_a| = q_a - 1$ and $|N_b| = q_b - 1$ such that every member of N_a has higher priority than j at a and every member of N_b has higher priority than i at b. The priority profile is *acyclic* if it has no Ergin cycle.

The weakening is what makes the property available. Under DA full non-bossiness is equivalent to Ergin acyclicity, and hence to the efficiency of DA itself, so it holds only where there is nothing to repair.[7] Proposition 1 says that the local part costs nothing: it holds at every priority structure and every capacity vector. This paper asks whether it also costs nothing once the designer insists on repairing the inefficiency, and the answer turns out to be a condition on priorities and capacities again.

Remark 1. The two ends of the stable lattice divide the two properties between them. The student-optimal mechanism is strategy-proof for students and locally non-bossy, and it is non-bossy only under acyclicity. The school-proposing stable mechanism is non-bossy, and indeed no misreport under it can leave the deviating student no worse off while harming somebody else (Afacan & Dur, 2017), but it is not strategy-proof. Moreno-García and Torres-Martínez (2026) compare how manipulable by coalitions the stable mechanisms are, and immunity to coalitional manipulation requires the two properties jointly, which no stable mechanism can guarantee at every priority structure. The school-optimal end is nevertheless outside the class studied here, since it is the worst stable matching for students

a mechanism is bossy, or locally bossy, when some report of some student at some profile is.

[7] Ergin (2002) shows that DA is group strategy-proof if and only if $\succ$ is acyclic, and a strategy-proof mechanism is group strategy-proof exactly when it is non-bossy (Pápai, 2000; Thomson, 2016). Since DA is strategy-proof, DA is non-bossy if and only if $\succ$ is acyclic, which by Ergin's theorem is also the condition under which DA is efficient.

and therefore does not dominate DA. Local non-bossiness is what the student-optimal end retains at every priority structure, and what the rest of the paper compromises.

3 Improvements over DA

This section collects what a matching that Pareto improves upon a stable matching can do, and turns it into restrictions on local non-bossiness. A DA-dominating mechanism can only move students among the seats that DA already fills. This keeps the conflict at schools with several seats and fixes the position of the student who can create it. Once efficiency and DA-dominance are imposed, choosing a mechanism means choosing one welfare correction at each profile, and local non-bossiness becomes a restriction that links the corrections chosen at two profiles that differ in a single report. The three properties are then incompatible as soon as two such profiles admit no compatible pair of corrections, which is what we call a rigid pair (Corollary 1).

Lemma 1. *Let μ be individually rational and non-wasteful at P, and let ν weakly Pareto dominate μ at P. Then*

(i) $|\nu_c| = |\mu_c|$ *for every* $c \in C$, *and* $\nu_\emptyset = \mu_\emptyset$;

(ii) $\nu_c = \mu_c$ *at every school* c *that is undersubscribed at* μ;

(iii) every student i *with* $\nu_i \neq \mu_i$ *satisfies* $\mu_i \in C$, $\nu_i \in C$ *and* $\nu_i \, P_i \, \mu_i$;

(iv) if $\nu_c \neq \mu_c$ *for some* $c \in C$, *then there are an integer* $L \geq 2$, *distinct schools* $a_1 = c, a_2, \ldots, a_L$ *and distinct students* $k_1, \ldots, k_L$ *such that*

$$\mu_{k_\ell} = a_\ell \quad \text{and} \quad \nu_{k_\ell} = a_{\ell+1} \qquad \text{for every } \ell,$$

indices taken modulo L; *moreover* k_L *may be taken to be any prescribed student of* $\nu_c \setminus \mu_c$.

Proof. (i) If $\mu_i \in C$, individual rationality and strict preferences give $\mu_i \, P_i \, \emptyset$, so weak Pareto dominance implies $\nu_i \neq \emptyset$ and hence $\nu_\emptyset \subseteq \mu_\emptyset$. Suppose $|\nu_c| > |\mu_c|$ for some c. Feasibility of ν implies $|\mu_c| < q_c$, and there is $i \in \nu_c \setminus \mu_c$; since $\nu_i = c \neq \mu_i$ and preferences are strict, $c \, P_i \, \mu_i$, so i prefers a school with a vacant seat at μ, contradicting non-wastefulness. Hence $|\nu_c| \leq |\mu_c|$ at every school. Summing gives $\sum_c |\nu_c| \leq \sum_c |\mu_c|$, while $\nu_\emptyset \subseteq \mu_\emptyset$ gives the reverse inequality for the numbers of assigned students. Therefore equality holds at every school and $\nu_\emptyset = \mu_\emptyset$.

(ii) The argument in the previous paragraph shows that no student can move into a school that has a vacant seat at μ, so $\nu_c \subseteq \mu_c$ whenever c is undersubscribed, and part (i) gives equality.

(iii) By part (i) no student of $\mu_\emptyset$ moves and no assigned student becomes unassigned, and strict preferences with weak dominance give $\nu_i \, P_i \, \mu_i$.

(iv) Let i^* be any student of $\nu_c \setminus \mu_c$, which is non-empty because $\nu_c \neq \mu_c$ while the two sets have the same number of elements by part (i). Say that a school b *feeds* a school b' if some student has $\mu_h = b$ and $\nu_h = b'$; by part (iii) both are schools. Part (i) gives $|\mu_b \setminus \nu_b| = |\nu_b \setminus \mu_b|$ at every school, so at every school the number of students who leave equals the number who arrive.

Let B be the set of schools that can be reached from c by a chain of feedings, with $c \in B$. Every school of B feeds only schools of B, so every student who leaves a school of B arrives at a school of B. Adding the equalities above over the schools of B, the number of students who leave a school of B equals the number who arrive at one; since the departures from B already account for that many arrivals, no student arrives at a school of B from outside B. Student i^* arrives at $c \in B$ from μ_{i^*}, hence $\mu_{i^*} \in B$.

Take a shortest chain of feedings $c = b_0, b_1, \ldots, b_r = \mu_{i^*}$. Its schools are distinct, and $r \geq 1$ because $\mu_{i^*} \neq c$. For every $t < r$ choose a student k_{t+1} with $\mu_{k_{t+1}} = b_t$ and $\nu_{k_{t+1}} = b_{t+1}$, put $L = r + 1$, $k_L = i^*$ and $a_{t+1} = b_t$ for $t \leq r$. Then $\mu_{k_\ell} = a_\ell$ and $\nu_{k_\ell} = a_{\ell+1}$ for $\ell < L$, while $\mu_{k_L} = b_r = a_L$ and $\nu_{k_L} = c = a_1$. The schools are distinct by construction and the students are distinct because they occupy distinct schools under μ. □

Parts (i) and (ii) are known. Part (i) appears, in various formulations, in Abdulkadiroğlu, Pathak, and Roth (2009), Erdil (2014) and Kesten and Kurino (2019); see also the discussion in Tang and Yu (2014) and Ortega et al. (2026). Troyan, Delacrétaz, and Kloosterman (2020, Proposition 3) prove an analogous invariance on a different domain, showing that any two essentially stable matchings agree at every undersubscribed school and leave the same students unassigned.

Applied to $\mu = \mathrm{DA}(P)$, which is stable and therefore individually rational and non-wasteful, the lemma says three things about a welfare correction over DA. First, the correction never expands access: it leaves the set of unassigned students and the number of students at each school exactly as DA leaves them, so it can only permute students among seats that DA already fills. Second, and more sharply, it leaves the entire cohort of any school that DA does not fill, not merely its size. A school with a vacant seat is one whose priorities do not bind, since no student wants a seat there that she does not already have. Then, welfare corrections redistribute *contested* seats, those belonging to schools that DA fills and for which some rejected student still has a claim, and only those. Third,

every correction is a collection of trading cycles among assigned students, and part (iv) says that a correction which changes the cohort of a school c contains a cycle through c: a list of schools starting and ending at c along which each student moves to the school the next one vacates. Section 5 reads that cycle back into the priority structure.

Applying the lemma to $\nu = \varphi(P)$ for a DA-dominating mechanism φ delivers the consequences for local non-bossiness.

Proposition 2. *Let φ be DA-dominating. Fix $P \in \mathcal{P}$, $i \in I$, and another report P_i', and let $P' = (P_i', P_{-i})$.*

(i) If $\varphi_i(P) = \varphi_i(P') = \emptyset$, then $\varphi_\emptyset(P) = \varphi_\emptyset(P')$.

(ii) If $\varphi_i(P) = \varphi_i(P') = c \in C$ and $\varphi_c(P) \neq \varphi_c(P')$, then $q_c \geq 2$.

(iii) If $\varphi_i(P) = \varphi_i(P') = c \in C$, then $\mathrm{DA}_i(P) = c$ if and only if $\mathrm{DA}_i(P') = c$.

Proof. For part (i), suppose $\varphi_i(P) = \emptyset$. Because $\varphi(P)$ weakly Pareto dominates the individually rational matching $\mathrm{DA}(P)$, student i must also be unassigned under $\mathrm{DA}(P)$; otherwise $\mathrm{DA}_i(P)\, P_i\, \emptyset = \varphi_i(P)$. The same argument gives $\mathrm{DA}_i(P') = \emptyset$. Proposition 1, applied to the outside option, therefore gives $\mathrm{DA}_\emptyset(P) = \mathrm{DA}_\emptyset(P')$, and Lemma 1(i) gives $\varphi_\emptyset(P) = \mathrm{DA}_\emptyset(P)$ and $\varphi_\emptyset(P') = \mathrm{DA}_\emptyset(P')$.

For part (ii), student i belongs to both $\varphi_c(P)$ and $\varphi_c(P')$. If $q_c = 1$, feasibility forces both sets to equal $\{i\}$, contradicting $\varphi_c(P) \neq \varphi_c(P')$.

For part (iii), suppose $\mathrm{DA}_i(P) = c$. Strategy-proofness of DA, applied to student i with preference P_i' against the fixed reports P_{-i}, gives $\mathrm{DA}_i(P')\, R_i'\, \mathrm{DA}_i(P_i, P_{-i}) = c$. Since $\varphi(P')$ weakly Pareto dominates $\mathrm{DA}(P')$ at P', we also have $c = \varphi_i(P')\, R_i'\, \mathrm{DA}_i(P')$. Strict preferences give $\mathrm{DA}_i(P') = c$, and the converse follows by exchanging the roles of P_i and P_i'. □

Regarding part (i), local non-bossiness restricts the outside option as well as the schools, and Duque-Rosas, Pereyra, and Torres-Martínez (2026) stress that this is what makes the property non-trivial in one-to-one problems. For DA-dominating mechanisms that part never binds, because a student who remains unassigned cannot change who else is unassigned, whatever the capacities. Part (ii) places the conflict at schools with several seats, where students compete for a scarce resource and priorities do the rationing. Part (iii) places the manipulator relative to such a school. A student who receives the same school under two reports either holds a seat there under DA at both profiles, so that the correction leaves her where she was, or is denied it under DA at both profiles, so that the correction admits her. The intermediate case, in which DA grants her the school under one report and the correction has to fetch it under the other, cannot arise.

Remark 2. Part (ii) already settles the one-to-one case. If $q_c = 1$ for every $c \in C$, then every DA-dominating mechanism is locally non-bossy, since the property has no bite at the outside option by part (i) and a school with a single seat cannot change its composition while keeping the same student. Efficient DA-dominating mechanisms exist in every problem, as Lemma 2 records, so efficiency, DA-dominance and local non-bossiness are then compatible. The case is of limited interest for school choice: when every school admits one student, the composition of a school is the student assigned to it, so local non-bossiness is satisfied automatically in this class, while ordinary non-bossiness remains restrictive.

3.1 The design problem is a selection problem

Efficiency and DA-dominance leave the designer with a choice. At a given profile they rule out every matching that makes some student worse off than DA does and every matching that still leaves gains on the table, but in general several matchings survive both tests, and the designer must pick one of them. For a profile P, let $\mathcal{D}(P)$ be the set of matchings that weakly Pareto dominate $\mathrm{DA}(P)$ at P and let

$$\mathcal{E}(P) = \{\nu \in \mathcal{D}(P) : \ \nu \text{ is Pareto efficient at } P\}$$

be the set of welfare corrections a designer may implement at P once she has committed to leaving no student worse off than under DA and to exhausting the remaining gains. The class of efficient DA-dominating mechanisms is exactly the class of selections from $\mathcal{E}$, and it is non-empty and easy to describe constructively. Fix a linear order $\rhd$ on I, to be read as an order of precedence that the designer uses to break ties between welfare corrections, and refine $\mathcal{E}(P)$ by keeping the elements that are best for the $\rhd$-first student, then those best for the second, and so on. Call the resulting mechanism the *serial refinement* $\varphi^{\rhd}$.

Lemma 2. *$\mathcal{E}(P) \neq \emptyset$ for every P, and a mechanism φ is efficient and DA-dominating if and only if $\varphi(P) \in \mathcal{E}(P)$ for every $P \in \mathcal{P}$. Moreover, for every linear order $\rhd$ on I the serial refinement $\varphi^{\rhd}$ is well defined, efficient and DA-dominating.*

Proof. The set $\mathcal{D}(P)$ is non-empty and finite, so it has an element that no other element Pareto dominates; if some matching Pareto dominated it, transitivity would place that matching in $\mathcal{D}(P)$ and contradict the choice, so the element is Pareto efficient and belongs to $\mathcal{E}(P)$. If φ is efficient and DA-dominating then $\varphi(P)$ lies in $\mathcal{D}(P)$ and is Pareto efficient, hence in $\mathcal{E}(P)$; the converse is immediate. Finally, the refinement defining $\varphi^{\rhd}$ operates on a finite non-empty set and each step preserves non-emptiness, and any two elements

surviving all steps give every student the same assignment and therefore coincide, so $\varphi^{\triangleright}(P)$ is a well-defined element of $\mathcal{E}(P)$. □

Local non-bossiness is then a restriction on the *selection*, not on any single profile: it links the corrections chosen at profiles that differ in one report, which is what makes the joint requirement demanding.

Definition 4. A pair of profiles (P, P') with $P' = (P'_i, P_{-i})$ is a *rigid pair* for i if there is $x \in C \cup \{\emptyset\}$ such that every $\nu \in \mathcal{E}(P)$ and every $\nu' \in \mathcal{E}(P')$ satisfy

$$\nu_i = \nu'_i = x \qquad \text{and} \qquad \nu_x \neq \nu'_x.$$

At a rigid pair the two reports lead to welfare corrections that put i at the same school with different classmates, whichever admissible correction the designer implements at each profile. The designer has no way out, and the following is immediate.

Corollary 1. *If a school choice problem Γ admits a rigid pair, then no mechanism is simultaneously efficient, DA-dominating and locally non-bossy.*

Proof. Let φ be efficient and DA-dominating and let (P, P') be a rigid pair for i. By Lemma 2, $\varphi(P) \in \mathcal{E}(P)$ and $\varphi(P') \in \mathcal{E}(P')$. Rigidity gives $\varphi_i(P) = \varphi_i(P') = x$ and $\varphi_x(P) \neq \varphi_x(P')$ for some x, so local non-bossiness fails. □

To show that efficiency, DA-dominance and local non-bossiness are incompatible in a given problem it is therefore enough to exhibit a rigid pair, and every impossibility below is established in this way. The rigid pairs behind Theorems 1 and 3 have an additional feature, which extends the conclusion to the larger class of DA-improving mechanisms: at each of the two profiles, the admissible correction is not merely the only efficient matching that weakly dominates DA, but the only matching other than DA that does.

Definition 5. A rigid pair (P, P') is *tight* if $\mathcal{D}(P) = \{\mathrm{DA}(P), \nu\}$ and $\mathcal{D}(P') = \{\mathrm{DA}(P'), \nu'\}$ for some matchings $\nu \neq \mathrm{DA}(P)$ and $\nu' \neq \mathrm{DA}(P')$.

Corollary 2. *If a school choice problem Γ admits a tight rigid pair, then no DA-improving mechanism is locally non-bossy. In particular, no mechanism is simultaneously efficient, DA-dominating and locally non-bossy.*

Proof. Let (P, P') be a tight rigid pair for i, with ν and ν' as in Definition 5. Since ν weakly Pareto dominates $\mathrm{DA}(P)$ and differs from it, it Pareto dominates it, so $\mathrm{DA}(P)$ is not Pareto efficient. Moreover ν is Pareto efficient, because a matching that Pareto dominated ν would weakly Pareto dominate $\mathrm{DA}(P)$ and would therefore be either $\mathrm{DA}(P)$,

which ν Pareto dominates, or ν itself. Hence $\mathcal{E}(P) = \{\nu\}$, and likewise $\mathcal{E}(P') = \{\nu'\}$, so rigidity gives $\nu_i = \nu'_i = x$ and $\nu_x \neq \nu'_x$ for some x. Let φ be DA-improving. Then $\varphi(P) \in \mathcal{D}(P)$ and $\varphi(P) \neq \mathrm{DA}(P)$, so $\varphi(P) = \nu$, and likewise $\varphi(P') = \nu'$. Hence φ is not locally non-bossy. The last statement holds because every efficient DA-dominating mechanism is DA-improving. □

3.2 Which reports can be bossy

A last restriction concerns the manipulation itself. Under DA a student's report matters only through the schools she applies to, hence only through the schools she ranks above the one she obtains.

Proposition 3. *Let (P, P') be a rigid pair for i and let $e = \mathrm{DA}_i(P)$. Then*

$$\{y \in C \cup \{\emptyset\} : \ y \, P_i \, e\} \neq \{y \in C \cup \{\emptyset\} : \ y \, P'_i \, e\}.$$

Proof. Suppose the two sets coincide. We first claim that $\mathrm{DA}(P') = \mathrm{DA}(P) =: \mu$.[8] The school e is acceptable under P'_i, since otherwise $\emptyset \, P'_i \, e$ and hence $\emptyset \, P_i \, e$, contradicting individual rationality of μ at P. The matching μ is then stable at P': individual rationality holds because $\mu_i = e$ is acceptable under P'_i and the other reports are unchanged, a blocking pair (i, y) at P' would require $y \, P'_i \, e$ and hence $y \, P_i \, e$ and would block at P as well, and blocking pairs not involving i are the same at both profiles. Student optimality at P' therefore gives $\mathrm{DA}_i(P') \, R'_i \, e$, so $\mathrm{DA}_i(P')$ lies in $\{y : y \, P'_i \, e\} \cup \{e\} = \{y : y \, P_i \, e\} \cup \{e\}$ and hence $\mathrm{DA}_i(P') \, R_i \, e$; strategy-proofness of DA gives $e = \mathrm{DA}_i(P) \, R_i \, \mathrm{DA}_i(P')$, so $\mathrm{DA}_i(P') = e$. The hypothesis is symmetric once this is known, so $\mathrm{DA}(P')$ is stable at P as well, and student optimality at each profile gives $\mathrm{DA}_h(P) = \mathrm{DA}_h(P')$ for every h.

The set of schools weakly preferred to μ_h is now the same under P and P' for every student, including i, so $\mathcal{D}(P) = \mathcal{D}(P') =: \mathcal{D}$. Rigidity implies that all elements of $\mathcal{E}(P)$ and all elements of $\mathcal{E}(P')$ assign i to one and the same school x, since otherwise a pair with $\nu_i \neq \nu'_i$ would be available. Let $\nu \in \mathcal{E}(P)$, so $\nu_i = x$, and suppose $\nu \notin \mathcal{E}(P')$; let $w \in \mathcal{D}$ Pareto dominate ν at P'. For every $h \neq i$ we have $w_h \, R_h \, \nu_h$, because those students report the same preferences at both profiles. If we also had $w_i \, R_i \, x$, then w would weakly Pareto dominate ν at P, and the student for whom the dominance at P' is strict would make it strict at P as well, contradicting the efficiency of ν at P. Hence $x \, P_i \, w_i$, and $w_i \, P'_i \, x$. Choose $\nu' \in \mathcal{E}(P')$ weakly Pareto dominating w at P'. Then $\nu'_i \, R'_i \, w_i \, P'_i \, x$, so $\nu'_i \neq x$, a

[8] The claim is a standard consequence of student optimality and strategy-proofness, and it is the same device, applied to arbitrary reports rather than to truncations, that Duque-Rosas, Pereyra, and Torres-Martínez (2026) use in the proof of their Theorem 1.

contradiction. Hence $\mathcal{E}(P) \subseteq \mathcal{E}(P')$ and, symmetrically, $\mathcal{E}(P') \subseteq \mathcal{E}(P)$. Taking $\nu' = \nu$ then contradicts rigidity. □

The proof shows more than the statement needs, and the two should be kept apart. Reordering, adding or deleting schools that a family declares *worse* than the school DA gives it leaves the DA outcome and the whole set $\mathcal{D}$ of admissible corrections untouched, because the same applications are made and the same rejections are triggered, so the designer faces the same menu. It does not follow that every mechanism in the class ignores that part of the report, since a selection from $\mathcal{E}$ may respond to it. Proposition 3 says that a rigid pair cannot be built that way. A report that produces one must change the set of schools that the family declares better than its DA assignment, that is, it must change which rejections the family provokes on its way down its list. In every problem of Sections 4 and 5 the manipulating student does exactly this, by moving one school from below to above the school DA assigns her, as Figure 1 shows.

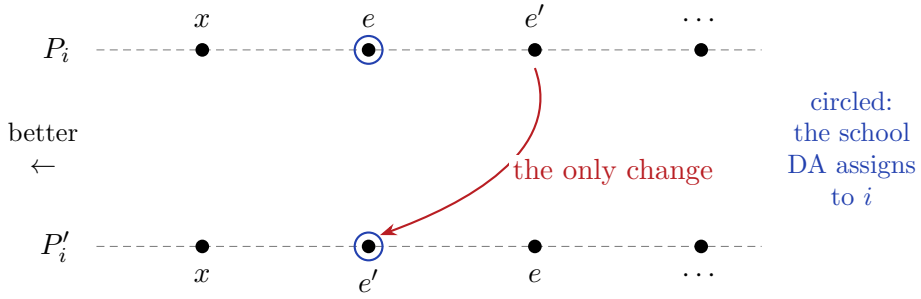


Figure 1: The shape of a bossy deviation. Student i is assigned $e = \mathrm{DA}_i(P)$ at P and $e' = \mathrm{DA}_i(P')$ at P', while every admissible welfare correction assigns her the same x at both profiles. Proposition 3 forces the two reports to differ above e, and Proposition 5 forces e and e' to be different schools, both distinct from x. In the problems of this paper the reports differ in the smallest way compatible with this, by moving one school from just below to just above the school DA assigns.

4 The impossibility

We now show that implementing an available improvement over DA is already incompatible with local non-bossiness in a small problem containing a school with two seats. We exhibit a school choice problem with two preference profiles that differ in the report of a single student. At each profile exactly one matching other than DA weakly Pareto dominates it, so every DA-improving mechanism, and in particular every efficient DA-dominating one, is forced to select it, and the two forced selections give the deviating student the same

school with a different classmate. The problem is balanced and every school is acceptable to every student, so neither unassigned students nor outside options play any role.

Theorem 1. *There is a school choice problem in which no DA-improving mechanism is locally non-bossy. In particular, no mechanism is simultaneously efficient, DA-dominating and locally non-bossy in it.*

Proof. Let Γ^1 be the school choice problem with

$$I = \{i_1, i_2, i_3, i_4, i_5\}, \qquad C = \{c_1, c_2, c_3, c_4\}, \qquad q_{c_1} = 2, \quad q_{c_2} = q_{c_3} = q_{c_4} = 1,$$

the priorities of Table 1, and the preference profile P of the same table. The profile $P' = (P'_{i_1}, P_{-i_1})$ differs from $P = (P_{i_1}, P_{-i_1})$ only in the report of i_1.

Priorities		**Preferences**	
$\succ_{c_1}$:	i_2, i_3, i_1, i_4, i_5	P_{i_1}:	c_1, c_2, c_3, c_4
$\succ_{c_2}$:	i_1, i_4, i_2, i_3, i_5	P'_{i_1}:	c_1, c_3, c_2, c_4
$\succ_{c_3}$:	i_1, i_5, i_3, i_2, i_4	P_{i_2}:	c_2, c_1, c_3, c_4
$\succ_{c_4}$:	i_4, i_5, i_1, i_2, i_3	P_{i_3}:	c_3, c_1, c_2, c_4
		P_{i_4}:	c_2, c_4, c_1, c_3
		P_{i_5}:	c_3, c_4, c_1, c_2

Table 1: The problem Γ^1 and the two preference profiles.

Total capacity equals the number of students and all schools are acceptable to all students. The two reports of i_1 differ only in the order of c_2 and c_3, the two schools immediately below her top choice.

The two DA outcomes. At P, school c_2 holds i_4 and rejects i_2, school c_3 holds i_5 and rejects i_3, and c_1 then receives i_1, i_2 and i_3 and keeps its two highest-priority applicants. Student i_1 next applies to c_2, displaces i_4 because $i_1 \succ_{c_2} i_4$, and i_4 is accepted at c_4. At P' the run is identical until i_1 is rejected by c_1: she then applies to c_3, displaces i_5 because $i_1 \succ_{c_3} i_5$, and i_5 is accepted at c_4, while i_4 keeps c_2. Hence

$$\mathrm{DA}(P): \quad c_1 = \{i_2, i_3\}, \quad c_2 = \{i_1\}, \quad c_3 = \{i_5\}, \quad c_4 = \{i_4\},$$

$$\mathrm{DA}(P'): \quad c_1 = \{i_2, i_3\}, \quad c_2 = \{i_4\}, \quad c_3 = \{i_1\}, \quad c_4 = \{i_5\}.$$

The unique efficient improvement at P. By Lemma 1(i) every matching in $\mathcal{D}(P)$ assigns two students to c_1 and one to each other school. Student i_5 holds her top choice c_3 and is fixed. Consequently, i_3, whose only improvement over c_1 is c_3, is fixed at c_1 and

occupies one of its two seats. The remaining students i_1, i_2 and i_4 fill the free seat of c_1 and the seats of c_2 and c_4, and their upper contour sets allow i_1 only c_1 or c_2, i_2 only c_2 or c_1, and i_4 only c_2 or c_4. If i_1 takes the seat at c_1, then i_4 cannot take c_2, since i_2 would then have no assignment weakly preferred to c_1. Then, i_2 takes c_2 and i_4 stays at c_4. If i_1 does not take that seat, i_2 does, and i_1 and i_4 stay where DA left them. Therefore $\mathcal{D}(P) = \{\mathrm{DA}(P), \nu\}$ with

$$\nu: \quad c_1 = \{i_1, i_3\},\ c_2 = \{i_2\},\ c_3 = \{i_5\},\ c_4 = \{i_4\},$$

and since ν makes i_1 and i_2 strictly better off, $\mathcal{E}(P) = \{\nu\}$.

The unique efficient improvement at P'. The argument is symmetric. Student i_4 holds her top choice c_2 and is fixed, which leaves i_2 with no improvement over c_1 and fixes her at c_1. The remaining students i_1, i_3 and i_5 fill the free seat of c_1 and the seats of c_3 and c_4, with i_1 allowed only c_1 or c_3, i_3 only c_3 or c_1, and i_5 only c_3 or c_4. If i_1 takes the seat at c_1, then i_5 cannot take c_3, since i_3 would then have no assignment weakly preferred to c_1. Therefore $\mathcal{D}(P') = \{\mathrm{DA}(P'), \nu'\}$ with

$$\nu': \quad c_1 = \{i_1, i_2\},\ c_2 = \{i_4\},\ c_3 = \{i_3\},\ c_4 = \{i_5\},$$

and $\mathcal{E}(P') = \{\nu'\}$.

Conclusion. We have $\nu_{i_1} = c_1 = \nu'_{i_1}$ while $\nu_{c_1} = \{i_1, i_3\} \neq \{i_1, i_2\} = \nu'_{c_1}$, so (P, P') is a rigid pair for i_1. Since $\mathcal{D}(P)$ and $\mathcal{D}(P')$ have two elements each, the pair is tight, and Corollary 2 applies. □

There are three relevant aspects in the argument of the proof. First, dominance is evaluated at the reported profile, and local non-bossiness is also a property of reports, so no separate notion of true preferences is needed. Second, efficiency plays no role beyond ruling out DA itself. At each profile of Γ^1 there is a single improvement over DA, so a mechanism that implements some improvement whenever one exists has no choice to make, and this is why the theorem is stated for DA-improving mechanisms. Finally, the deviating student moves a single school from below to just above the school DA assigns her, which is the smallest manipulation that Proposition 3 allows.

5 Bounding the impossibility

This section asks which problems exhibit the conflict between improving upon DA and local non-bossiness, and bounds them from both sides by conditions on priorities and

capacities, that is, on the objects a district designs rather than on the reports it receives. The necessary condition applies to every DA-dominating mechanism, and the sufficient condition to every DA-improving one.

Two benchmarks delimit the question. The first is the acyclicity condition of Definition 3.

Proposition 4. *If $\succ$ is acyclic, then DA is the only DA-dominating mechanism, and it is locally non-bossy.*

Proof. By Ergin (2002), acyclicity implies that $\mathrm{DA}(P)$ is Pareto efficient at every preference profile. Let φ be DA-dominating and fix P. Then $\varphi(P)$ weakly Pareto dominates the Pareto-efficient matching $\mathrm{DA}(P)$, so no student is strictly better off under $\varphi(P)$, and since preferences are strict every student receives the same assignment. Hence $\varphi(P) = \mathrm{DA}(P)$, and local non-bossiness follows from Proposition 1. □

The second benchmark is Remark 2. Unit capacities also leave no room for the conflict, whether or not the priority profile is acyclic. The two benchmarks are degenerate in opposite ways. Under acyclicity there is nothing to correct, because DA is already efficient. Under unit capacities there is nothing to be bossy about, because a student who keeps her school is the whole cohort of that school. A condition that describes the conflict must then restrict priorities and capacities together.

5.1 Rotations

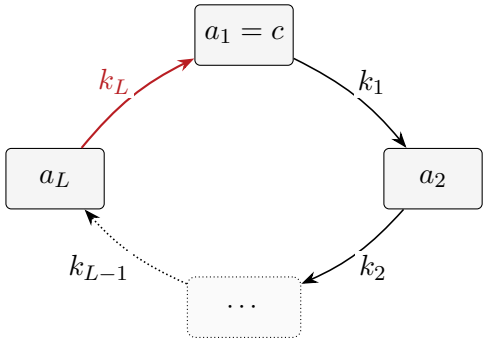
A welfare correction moves students among occupied seats along trading cycles, and a cycle that passes through a school c leaves a mark on the priority structure. At every school of the cycle, the student who holds the seat under DA outranks the student who arrives there through the correction, and so does every other student the school admits. We record this as a condition on $(\succ, q)$.

Definition 6. The school choice problem Γ admits a *rotation through* $c \in C$ if there are an integer $L \geq 2$, distinct schools $a_1 = c, a_2, \ldots, a_L$, distinct students $k_1, \ldots, k_L$, and pairwise disjoint sets $N_{a_1}, \ldots, N_{a_L} \subseteq I \setminus \{k_1, \ldots, k_L\}$ with $|N_{a_\ell}| = q_{a_\ell} - 1$, such that for every ℓ, indices taken modulo L,

$$\text{every member of } N_{a_\ell} \cup \{k_\ell\} \text{ has higher priority than } k_{\ell-1} \text{ at } a_\ell.$$

We call k_L the *terminal student* of the rotation and a_L its *origin*.

The reading is direct. We can think of k_ℓ as a student holding a seat at a_ℓ who would rather be at $a_{\ell+1}$, and of N_{a_ℓ} as the students who fill the remaining seats of a_ℓ. The displayed requirement says that a_ℓ is filled by students who all outrank $k_{\ell-1}$ there, so her claim on a_ℓ is denied, and that the seat she wants is held by k_ℓ, whom she can replace only if k_ℓ moves on. Thus, a rotation is a chain of denied claims that closes on itself, written in terms of priorities. The terminal student is the one who ends up admitted to c, and the origin is the school she comes from. For $L = 2$ the chain is a single reversal between c and one other school, and longer rotations let it travel through several schools before returning. Figure 2 displays the object.



k_ℓ holds a_ℓ under DA and $a_{\ell+1}$ under ν, indices modulo L
$N_{a_\ell} \cup \{k_\ell\} \succ_{a_\ell} k_{\ell-1}$, $|N_{a_\ell}| = q_{a_\ell} - 1$, the N_{a_ℓ} pairwise disjoint
k_L is the terminal student and a_L the origin

Figure 2: A rotation through c. Each arrow is a student who holds the school at its tail under DA and the school at its head under an improvement ν, so the cycle of arrows is the trading cycle of Lemma 1(iv). What Definition 6 adds is the line of priority requirements: at every school of the cycle the seats are held by students who outrank the one who wants to move in, which is why DA denies each link.

The next lemma reads a welfare correction back into the priority structure. Whenever a correction changes the cohort of a school, the problem contains a rotation through it. Theorem 2 then turns this into a condition under which the three properties are compatible.

Lemma 3. *Let P be a preference profile, let $\mu = \mathrm{DA}(P)$, and let ν weakly Pareto dominate μ at P. If $\nu_c \neq \mu_c$ for some school c, then Γ admits a rotation through c, and its terminal student may be taken to be any prescribed student of $\nu_c \setminus \mu_c$.*

Proof. Since $\mathrm{DA}(P)$ is stable, it is individually rational and non-wasteful, so Lemma 1(iv) applies and delivers distinct schools $a_1 = c, \ldots, a_L$ with $L \geq 2$ and distinct students

$k_1, \ldots, k_L$ with $\mu_{k_\ell} = a_\ell$ and $\nu_{k_\ell} = a_{\ell+1}$, indices modulo L, and with k_L any prescribed student of $\nu_c \setminus \mu_c$. Part (iii) of the same lemma gives $a_{\ell+1} \, P_{k_\ell} \, a_\ell$.

Fix ℓ. Student $k_{\ell-1}$ prefers a_ℓ to her assignment $\mu_{k_{\ell-1}} = a_{\ell-1}$, and μ is stable at P, so $(k_{\ell-1}, a_\ell)$ is not a blocking pair. Hence a_ℓ is full at μ and every member of μ_{a_ℓ} has higher priority than $k_{\ell-1}$ at a_ℓ. Since $k_\ell \in \mu_{a_\ell}$, the set $N_{a_\ell} := \mu_{a_\ell} \setminus \{k_\ell\}$ has $q_{a_\ell} - 1$ elements and every member of $N_{a_\ell} \cup \{k_\ell\} = \mu_{a_\ell}$ ranks above $k_{\ell-1}$ at a_ℓ. Finally, the sets $N_{a_1}, \ldots, N_{a_L}$ are contained in the cohorts of distinct schools at the single matching μ, hence pairwise disjoint, and they contain no k_m, because $k_m \in \mu_{a_m}$ and $k_m \notin N_{a_m}$. □

Theorem 2. *If a school choice problem Γ admits no rotation through any school with at least two seats, then every DA-dominating mechanism in Γ is locally non-bossy, whether or not it is efficient. Consequently, if no mechanism in Γ is simultaneously efficient, DA-dominating and locally non-bossy, as happens in particular when no DA-improving mechanism in Γ is locally non-bossy, then Γ admits a rotation through some school with at least two seats.*

Proof. Let φ be DA-dominating. Let c be any school with $q_c \geq 2$ and let P be any preference profile. Since $\varphi(P)$ weakly Pareto dominates $\mathrm{DA}(P)$ and Γ admits no rotation through c, Lemma 3 gives $\varphi_c(P) = \mathrm{DA}_c(P)$.

Suppose local non-bossiness failed, so that $\varphi_i(P) = \varphi_i(P') = x$ and $\varphi_x(P) \neq \varphi_x(P')$ for some i, some P and some $P' = (P'_i, P_{-i})$. By Proposition 2(i) we have $x \in C$, and by Proposition 2(ii) we have $q_x \geq 2$. The previous paragraph then gives $\varphi_x(P) = \mathrm{DA}_x(P)$ and $\varphi_x(P') = \mathrm{DA}_x(P')$. Since $i \in \varphi_x(P) = \mathrm{DA}_x(P)$ we have $\mathrm{DA}_i(P) = x$, and likewise $\mathrm{DA}_i(P') = x$. Proposition 1 therefore gives $\mathrm{DA}_x(P) = \mathrm{DA}_x(P')$, so $\varphi_x(P) = \varphi_x(P')$, a contradiction. The second statement follows because efficient DA-dominating mechanisms exist by Lemma 2, and they are DA-improving and in particular DA-dominating. □

The problem Γ^1 of Theorem 1 illustrates what the theorem requires. Take $L = 2$, $a_1 = c_1$, $a_2 = c_2$, $k_1 = i_2$, $k_2 = i_1$, $N_{c_1} = \{i_3\}$ and $N_{c_2} = \emptyset$. At c_1 both i_3 and i_2 have higher priority than i_1, and at c_2 student i_1 has higher priority than i_2, so this is a rotation through c_1 with terminal student i_1 and origin c_2. It is the footprint of the correction that produces the impossibility: DA leaves i_1 at c_2 and fills c_1 with i_2 and i_3, and the unique efficient improvement admits i_1 to c_1 in exchange for the seat that i_2 takes at c_2. Replacing c_2 by c_3 and i_2 by i_3, with $N_{c_1} = \{i_2\}$, gives a second rotation with the same terminal student and origin c_3, which is the correction forced at the other profile.

Theorem 2 also contains Remark 2 as a special case: under unit capacities there is no school with two seats, so there is nothing to rotate through. It is strictly more general,

and the gain is not vacuous. Let Γ^2 be the problem with

$$I = \{i_1, \dots, i_5\}, \qquad C = \{c_1, c_2, c_3\}, \qquad q_{c_1} = 3, \quad q_{c_2} = q_{c_3} = 1,$$

$$\succ_{c_1}\colon\ i_1, i_2, i_3, i_4, i_5, \qquad \succ_{c_2}\colon\ i_1, i_2, i_3, i_4, i_5, \qquad \succ_{c_3}\colon\ i_2, i_3, i_1, i_4, i_5.$$

Its priority profile is not acyclic, since $a = c_2$, $b = c_3$ and $(i, j, k) = (i_1, i_2, i_3)$ form an Ergin cycle, and DA is Pareto inefficient at some preference profiles: if i_1 ranks c_3 then c_2, student i_3 ranks c_2 then c_3, student i_2 finds only c_2 acceptable and i_4 and i_5 find no school acceptable, then DA assigns c_2 to i_1 and c_3 to i_3, and exchanging their assignments makes both strictly better off. Lemma 8 in the Appendix shows that Γ^2 admits no rotation through c_1, so by Theorem 2 every efficient DA-dominating mechanism in Γ^2 is locally non-bossy.

The priority reversal that makes DA inefficient in Γ^2 involves only the two single-seat schools, and the improvements it makes available never reach c_1. For a student to be admitted to c_1 after being rejected there, two further students would have to hold its remaining seats with a stronger claim than hers and a third would have to be willing to leave, and the priority profile does not supply them. Here capacity protects the composition of c_1 rather than exposing it. Welfare corrections and oversubscribed schools can therefore coexist, as long as the corrections do not reach the oversubscribed schools, and the region of compatibility is larger than the two benchmarks above.

For every conflict that this paper produces the condition can be sharpened. All of them come from rigid pairs, and at a rigid pair the position of the manipulator is fixed. DA denies her the contested school under both reports, and the two reports send her to two different schools. Her report chooses which seat she vacates, and that is what chooses her classmate. The next result proves this, and it is what allows the single rotation of Theorem 2 to be replaced by two.

Proposition 5. *Let (P, P') be a rigid pair for i, and let x be the school or outside option that every admissible welfare correction assigns to i at both profiles, as Definition 4 requires. Then $x \in C$, $q_x \geq 2$,*

$$\mathrm{DA}_i(P) \neq x, \qquad \mathrm{DA}_i(P') \neq x, \qquad \textit{and} \qquad \mathrm{DA}_i(P) \neq \mathrm{DA}_i(P').$$

Proof. Let φ be efficient and DA-dominating, so that $\varphi(P) \in \mathcal{E}(P)$ and $\varphi(P') \in \mathcal{E}(P')$ by Lemma 2. Rigidity gives $\varphi_i(P) = \varphi_i(P') = x$ and $\varphi_x(P) \neq \varphi_x(P')$, so $x \in C$ by Proposition 2(i) and $q_x \geq 2$ by Proposition 2(ii). By Proposition 2(iii) the two profiles agree on whether DA assigns x to i, so it remains to rule out $\mathrm{DA}_i(P) = \mathrm{DA}_i(P') =: e$,

which covers both $e = x$ and $e \neq x$.

Suppose that is the case, and let P_i^* be the report whose acceptable set is $\{x, e\}$ with x ranked first, which is the report declaring only x acceptable when $e = x$. Write $P^* = (P_i^*, P_{-i})$, $\mu = \mathrm{DA}(P)$, $\mu' = \mathrm{DA}(P')$ and $\lambda = \mathrm{DA}(P^*)$. Every $\nu \in \mathcal{E}(P)$ has $\nu_i = x$ and weakly Pareto dominates μ, so $x \, R_i \, e$, and likewise $x \, R_i' \, e$. Moreover, when $e \neq x$ both preferences are strict.

Step 1: μ and μ' are stable at P^*. Individual rationality holds because e is acceptable under P_i^* and the other reports are unchanged. The only pair involving i that can block at P^* is (i, x), since nothing else is ranked above e by P_i^*. Now (i, x) blocks μ at P^* if and only if $x \, P_i^* \, e$ and either x has a vacant seat at μ or i outranks some member of μ_x at x. The second requirement does not involve preferences, and it fails, because $x \, P_i \, e$ and μ is stable at P. The same argument applies to μ', using $x \, P_i' \, e$ and the stability of μ' at P'. Pairs not involving i are unaffected because the other students report the same preferences at P, P' and P^*.

Step 2: $\lambda_i = e$, and $\lambda_h \, R_h \, \mu_h$ and $\lambda_h \, R_h \, \mu_h'$ for every $h \neq i$. Student optimality at P^*, applied to the stable matchings μ and μ', gives the second claim and $\lambda_i \, R_i^* \, e$, so $\lambda_i \in \{x, e\}$. Strategy-proofness of DA, applied to student i with preference P_i, gives $e = \mathrm{DA}_i(P_i, P_{-i}) \, R_i \, \mathrm{DA}_i(P_i^*, P_{-i}) = \lambda_i$, which rules out $\lambda_i = x$ when $x \, P_i \, e$. Hence $\lambda_i = e$.

Step 3: $\mathcal{D}(P^*) \subseteq \mathcal{D}(P) \cap \mathcal{D}(P')$. Let $\nu \in \mathcal{D}(P^*)$. For $h \neq i$, $\nu_h \, R_h \, \lambda_h \, R_h \, \mu_h$ and $\nu_h \, R_h \, \lambda_h \, R_h \, \mu_h'$. For i, $\nu_i \, R_i^* \, \lambda_i = e$ gives $\nu_i \in \{x, e\}$, and both are weakly preferred to e under P_i and under P_i'.

Step 4: $\mathcal{E}(P^*) \subseteq \mathcal{E}(P) \cap \mathcal{E}(P')$. Let $\hat{\nu} \in \mathcal{E}(P^*)$, so $\hat{\nu} \in \mathcal{D}(P)$ by Step 3, and suppose $\hat{\nu} \notin \mathcal{E}(P)$. Some matching Pareto dominates $\hat{\nu}$ at P and therefore belongs to $\mathcal{D}(P)$; let $\nu \in \mathcal{E}(P)$ weakly Pareto dominate it at P, so that ν Pareto dominates $\hat{\nu}$ at P, and $\nu_i = x$ by rigidity. Then $\nu \in \mathcal{D}(P^*)$, because $\nu_h \, R_h \, \hat{\nu}_h \, R_h \, \lambda_h$ for $h \neq i$ and $\nu_i = x \, R_i^* \, e = \lambda_i$. Moreover ν Pareto dominates $\hat{\nu}$ at P^*: the students other than i report the same preferences at P and P^*, and $\nu_i = x$ is the best school under P_i^*, so the comparison is weak for every student and strict for at least one, namely for a student other than i if the dominance at P is strict there, and otherwise for i, in which case $\hat{\nu}_i = e$ and $x \, P_i^* \, e$. This contradicts $\hat{\nu} \in \mathcal{E}(P^*)$. The argument for $\mathcal{E}(P')$ is identical.

Step 5: conclusion. $\mathcal{E}(P^*)$ is non-empty by Lemma 2, so any $\hat{\nu} \in \mathcal{E}(P^*)$ lies in $\mathcal{E}(P) \cap \mathcal{E}(P')$. Taking $\nu = \nu' = \hat{\nu}$ in Definition 4 gives $\nu_i = \nu_i'$ and $\nu_x = \nu_x'$, which rigidity forbids. □

The report P_i^* used in the proof is the smallest one consistent with what the manipulator wants, since it declares acceptable only the school she ends up with and the school DA gives her. The argument shows that this report leaves the designer at least one welfare

correction that is admissible at both original profiles, which is what a rigid pair says cannot happen.

The two routes leave a footprint of their own on the priority structure. Each of them is a rotation through the contested school, the two admit the same student, and they start from different places, which is the configuration we record in the next definition.

Definition 7. The school choice problem Γ admits a *double rotation* at $c \in C$ if $q_c \geq 2$ and there are a student i and two rotations through c, both with terminal student i, whose origins are distinct schools.

Corollary 3. *If Γ admits a rigid pair, then Γ admits a double rotation.*

Proof. Let (P, P') be a rigid pair for i, let c be the assignment that i receives under every admissible welfare correction at both profiles, and let φ be efficient and DA-dominating. By Proposition 5, $q_c \geq 2$ and $\mathrm{DA}_i(P) \neq c$, so $i \in \varphi_c(P) \setminus \mathrm{DA}_c(P)$ and hence $\varphi_c(P) \neq \mathrm{DA}_c(P)$. Lemma 3, applied at P with prescribed terminal student i, yields a rotation through c with terminal student i whose origin is $\mathrm{DA}_i(P)$, since the origin is the school i occupies under $\mathrm{DA}(P)$. The same argument at P' yields a rotation through c with terminal student i and origin $\mathrm{DA}_i(P')$, and the two origins differ by Proposition 5. □

Corollary 3 sharpens Theorem 2 for every impossibility exhibited in this paper, since all of them are established by exhibiting a rigid pair. It says that a double rotation is necessary for an impossibility that comes from a rigid pair, which is not the same as a necessary condition for every possible failure of the three properties.

Remark 3. Neither condition is sufficient, and one problem shows it for both. Let Γ^3 have $I = \{i_1, \ldots, i_5\}$, $C = \{c_1, c_2, c_3\}$, $q_{c_1} = 2$, $q_{c_2} = q_{c_3} = 1$, and

$$\succ_{c_1}\colon\ i_1, i_2, i_3, i_4, i_5, \qquad \succ_{c_2} = \succ_{c_3}\colon\ i_1, i_2, i_3, i_5, i_4.$$

Taking $a_1 = c_1$, $a_2 = c_2$, $k_1 = i_4$, $k_2 = i_5$ and $N_{c_1} = \{i_1\}$ gives a rotation through c_1 with terminal student i_5 and origin c_2, since $i_4 \succ_{c_1} i_5$, $i_5 \succ_{c_2} i_4$ and $i_1 \succ_{c_1} i_5$. Replacing c_2 by c_3 gives a second rotation with the same terminal student and origin c_3, so Γ^3 admits a double rotation at c_1. Nevertheless, the priority profile is acyclic: an Ergin cycle requires two students ranked in opposite order by two schools, and the only such pair here is $\{i_4, i_5\}$. It also requires a third student ranked strictly between them at one of the two schools, and i_4 and i_5 are adjacent in all three orders. Proposition 4 therefore applies and the three properties are compatible in Γ^3.

5.2 Separating double rotations

A double rotation gives the manipulator two routes into the contested school, and Remark 3 shows that this is not enough to generate the impossibility between the three properties. What is missing is that the two routes must deliver different classmates. Along each route one incumbent of c leaves for the school the manipulator vacates, so the route she does not take has to be closed while she travels the other one. There are two ways of closing it. In the first, the seat the departing classmate would move into is held by a student whose priority lies strictly between the manipulator's and the classmate's, so only the manipulator can empty it. In the second, the classmate already holds that seat, and it is the arrival of the manipulator that sends her to c, where she displaces a classmate with lower priority. We call the first route guarded and the second direct.

Definition 8. The school choice problem Γ has a *separating double rotation* if there are a school c with $q_c \geq 2$, a student i, distinct students j_1, j_2 and distinct schools $d_1, d_2 \in C \setminus \{c\}$ with

$$j_1 \succ_c i, \qquad j_2 \succ_c i, \qquad i \succ_{d_1} j_1, \qquad i \succ_{d_2} j_2,$$

such that each route $m \in \{1, 2\}$ is either

- *guarded*: there is a student k_m with $i \succ_{d_m} k_m \succ_{d_m} j_m$, or
- *direct*: no such student is required,

at least one route is guarded, the guards of the guarded routes are distinct from each other and from i, j_1 and j_2, and, writing G for the set of guards and $\beta \in \{0, 1\}$ for the number of direct routes, there are pairwise disjoint sets $N_c, N_{d_1}, N_{d_2} \subseteq I \setminus (\{i, j_1, j_2\} \cup G)$ with

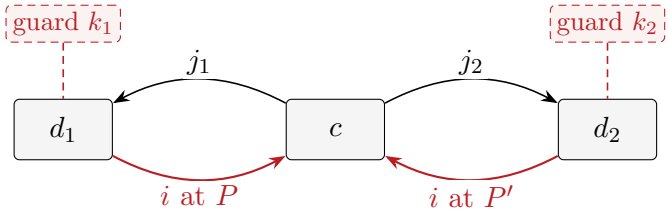
$$|N_c| = q_c - 2 + \beta, \qquad |N_{d_1}| = q_{d_1} - 1, \qquad |N_{d_2}| = q_{d_2} - 1,$$

such that every member of N_c has higher priority than i at c, every member of N_{d_m} has higher priority than k_m at d_m when route m is guarded and than j_m when it is direct, and, when $\beta = 1$, at least one member of N_c has lower priority than both j_1 and j_2 at c.

The configuration describes a student i who is outranked at an oversubscribed school c by two students j_1 and j_2, and who in turn outranks each of them at a different school. Student i has two routes out of c, and along each of them she can trade her seat back with a different one of the two students who displaced her. The scarcity sets say that each of the three schools involved is filled by students whose claims cannot be displaced, so these trades are the only ones available. The guarded and the direct case differ in who empties the seat that the departing classmate moves into. It is an intermediate student k_m in the

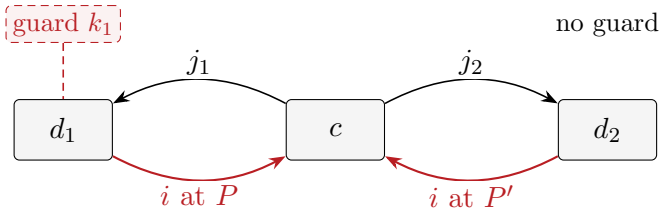
first case and the manipulator herself in the second. The extra classmate that a direct route requires at c is the student who loses her seat there when the departing classmate is pushed back, and it is her presence or absence that makes the two cohorts differ. Figure 3 displays the two forms the definition allows.

(a) both routes guarded



$N_c \cup \{j_1, j_2\} \succ_c i, \quad |N_c| = q_c - 2, \quad q_c \geq 2$
$N_{d_1} \cup \{i\} \succ_{d_1} k_1 \succ_{d_1} j_1$
$N_{d_2} \cup \{i\} \succ_{d_2} k_2 \succ_{d_2} j_2$

(b) route 2 direct



$N_c \cup \{j_1, j_2\} \succ_c i, \quad |N_c| = q_c - 1, \quad j_1, j_2 \succ_c g$ for some $g \in N_c$
$N_{d_1} \cup \{i\} \succ_{d_1} k_1 \succ_{d_1} j_1$
$N_{d_2} \cup \{i\} \succ_{d_2} j_2$

Figure 3: A separating double rotation at c, in the two forms Definition 8 allows. The arrows are the moves of the welfare correction and the lines below are the priority requirements. In (a) each route carries a guard, and the configuration is the twin cycle of Definition 9. In (b) the second route is direct: no guard is required on it, school c carries one more of the students of N_c, and one of them, g, is outranked by both j_1 and j_2. The labels on the red arrows anticipate the proof of Theorem 3, where i travels route 1 when she ranks d_1 second and route 2 when she ranks d_2 second. In (a) the classmate she keeps is the rival of the route she did not travel; in (b) it is g under one report and j_1 under the other.

Theorem 3. *If Γ has a separating double rotation, then Γ admits a tight rigid pair. Hence no DA-improving mechanism is locally non-bossy, and in particular no mechanism is simultaneously efficient, DA-dominating and locally non-bossy.*

The proof is in the Appendix, and it is worth describing the construction here. It involves four kinds of student. The *manipulator* is the student i of Definition 8, who wants

the contested school and reports it first, followed by one of her two origins. The origin she names second is the route she takes. The two *rivals* are j_1 and j_2, who hold seats at the contested school and each of whom prefers one of the two origins to it. A *guard* is the intermediate student of a guarded route, who occupies the last seat of that origin and finds nothing else acceptable. The remaining students are *fillers*, who find only one school acceptable and occupy the seats that the scarcity sets describe.

The run of DA is then easy to follow. The contested school fills with students who all outrank the manipulator, so it rejects her, and she goes to the origin she named, where she displaces whoever holds its last seat. The only exchange left is the one in which she returns to the contested school and the rival who wants that origin takes the seat she has just vacated, so the welfare correction is unique and every mechanism in the class is forced to implement it. The two kinds of route differ at one point only, namely in what keeps the route she did not take closed. If that route is guarded, its last seat is held by a guard whom only the manipulator can displace, so the rival of that route cannot move and stays a classmate. If it is direct, no guard stands in the way, the rival is already at her origin and never reaches the contested school, and its last seat goes to the extra classmate that Definition 8 requires. When the manipulator switches routes, her arrival displaces that rival, who returns to the contested school and pushes the extra classmate out of it. Either way the two reports leave the manipulator at the same school with different classmates, which is a rigid pair.

The classmate who changes is not the same in the two cases, and this is what the distinction is about. When both routes are guarded, the report of the manipulator decides which of her two rivals leaves the school and which one stays, so the cohort differs in the identity of the rival who remains. When one route is direct, her report decides whether that rival is pushed back into c at all, and the seat she would occupy there goes to a student with lower priority instead, so the cohort differs in who holds the last seat. Both routes cannot be direct. If they were, the manipulator's arrival at d_m would push j_m back to c along either route, the student displaced at c would be the same in both, and the two cohorts would coincide.

The special case in which both routes are guarded has a transparent description in terms of Ergin cycles, and is the configuration behind Γ^1.

Definition 9. A *twin cycle* is a separating double rotation in which both routes are guarded, that is, one with distinct students i, j_1, j_2, k_1, k_2 and distinct schools c, d_1, d_2 such that

$$j_1 \succ_c i, \qquad j_2 \succ_c i, \qquad i \succ_{d_1} k_1 \succ_{d_1} j_1, \qquad i \succ_{d_2} k_2 \succ_{d_2} j_2,$$

with $|N_c| = q_c - 2$, $|N_{d_1}| = q_{d_1} - 1$ and $|N_{d_2}| = q_{d_2} - 1$ as in Definition 8.

Lemma 4. *Every separating double rotation contains an Ergin cycle, and a twin cycle contains two of them, namely* (d_1, c) *with students* (i, k_1, j_1) *and* (d_2, c) *with students* (i, k_2, j_2)*. Consequently an acyclic problem has neither. The converse fails: a problem with unit capacities has no separating double rotation whether or not it is acyclic.*

Proof. Relabel so that route 1 is guarded, and take $a = d_1$, $b = c$ and the triple (i, k_1, j_1) in Definition 3. The priority requirements $i \succ_{d_1} k_1 \succ_{d_1} j_1$ and $j_1 \succ_c i$ hold by hypothesis. For the scarcity sets take $N_a = N_{d_1}$, which has $q_{d_1} - 1$ elements all ranked above k_1 at d_1, and $N_b = N_c \cup \{j_2\}$ if $\beta = 0$ and $N_b = N_c$ if $\beta = 1$; in both cases N_b has $q_c - 1$ elements, all ranked above i at c and none of them equal to i, k_1 or j_1, and the two sets are disjoint. When both routes are guarded the same argument applies to (d_2, c) with (i, k_2, j_2). The final claim holds because Definition 8 requires $q_c \geq 2$. □

Lemma 5. *A separating double rotation at* c *contains a double rotation at* c*.*

Proof. Set $L = 2$, $a_1 = c$, $a_2 = d_1$, $k'_1 = j_1$ and $k'_2 = i$ in Definition 6, so that the terminal student is i and the origin is d_1. At $a_1 = c$ take $N_{a_1} = N_c \cup \{j_2\}$ if $\beta = 0$ and $N_{a_1} = N_c$ if $\beta = 1$, which has $q_c - 1$ elements in both cases, all ranked above i at c and none equal to j_1 or i; the requirement at a_1 also asks j_1 to rank above i at c, which holds. At $a_2 = d_1$ take $N_{a_2} = N_{d_1}$, which has $q_{d_1} - 1$ elements, all ranked above j_1 at d_1, and the requirement $i \succ_{d_1} j_1$ holds by hypothesis. The two sets are disjoint. Replacing (d_1, j_1) by (d_2, j_2) gives a second rotation with the same terminal student and origin $d_2 \neq d_1$. □

These conditions are nested, as Figure 4 shows, and each of the three implications is strict. The problem Γ^4 of the Appendix has a separating double rotation with one direct route, so by Theorem 3 the three properties are incompatible in it, as Proposition 12 records, and Lemma 9 shows that it has no twin cycle. The problem Γ^3 of Remark 3 has a double rotation and no separating double rotation: its only school with several seats is c_1, and since $\succ_{c_2}$ and $\succ_{c_3}$ coincide and differ from $\succ_{c_1}$ only in the relative order of i_4 and i_5, the requirements $j \succ_{c_1} i$ and $i \succ_d j$ force $(i, j) = (i_5, i_4)$ for both $d = c_2$ and $d = c_3$, so j_1 and j_2 cannot be distinct. For the last implication, take two schools c and d with $q_c = 2$ and $q_d = 1$, three students a, b, i and the orders $\succ_c$: a, b, i and $\succ_d$: i, b, a. Setting $L = 2$, $k_1 = a$, $k_2 = i$, $N_c = \{b\}$ and $N_d = \emptyset$ gives a rotation through c, while a double rotation is impossible because it needs two distinct origins and there is only one school other than c. The problem Γ^2 of Section 5.1 is outside all four conditions, since by Lemma 8 it has no rotation through its only school with several seats.

The problem Γ^1 of Section 4 contains a twin cycle with $c = c_1$, $i = i_1$, $j_1 = i_2$, $j_2 = i_3$, $d_1 = c_2$, $d_2 = c_3$, $k_1 = i_4$, $k_2 = i_5$ and empty scarcity sets, so Theorem 1 is a special case of Theorem 3. We keep the example because its preferences are complete and its market

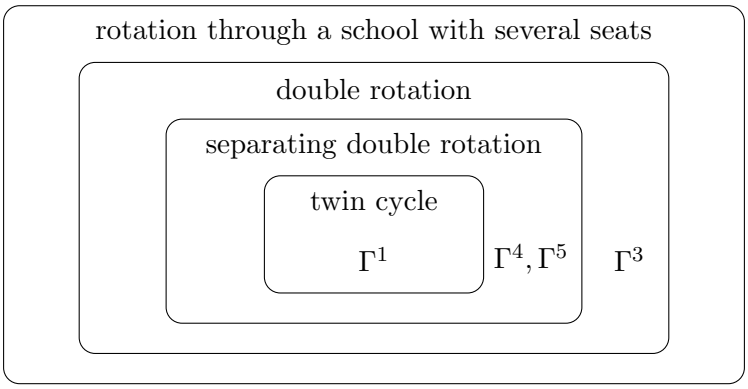


Figure 4: The four conditions and the problems of the paper. No DA-improving mechanism is locally non-bossy in the two innermost regions, by Theorem 3, and every DA-dominating mechanism is locally non-bossy outside the outermost one, by Theorem 2, where Γ^2 lies.

is balanced, which shows that the conflict does not rely on unacceptable schools or on unassigned students, whereas the profiles used in the proof of Theorem 3 do use short lists.

5.3 How common is the conflict

A condition met only by carefully built examples would say little about the systems that motivate the question. The configuration of Definition 8 is not of that kind, and the next observation is enough to see it. All it asks is that the student ranked last at a school with at least two seats not be ranked near the bottom at two other schools as well.

Proposition 6. *Let c be a school with $q_c \geq 2$ and let i be the student with the lowest priority at c. Suppose there are two distinct schools $d_1, d_2 \in C \setminus \{c\}$ such that at least four students rank below i at d_1 and at least four students rank below i at d_2, and suppose that*

$$|I| \;\geq\; q_c + q_{d_1} + q_{d_2} + 1.$$

Then Γ admits a separating double rotation, and hence no DA-improving mechanism is locally non-bossy.

Proof. Let j_1 and k_1 be the two students with the lowest priority at d_1, labelled so that $k_1 \succ_{d_1} j_1$. Since at least four students rank below i at d_1, both rank below i there, so $i \succ_{d_1} k_1 \succ_{d_1} j_1$ and route 1 is guarded with intermediate student k_1. At d_2, at least four students rank below i, so at least two of them lie outside $\{j_1, k_1\}$. Let j_2 and k_2 be the two students with the lowest priority at d_2 among those outside $\{i, j_1, k_1\}$, labelled so that $k_2 \succ_{d_2} j_2$. Both rank below i at d_2, so route 2 is guarded with intermediate student k_2.

The five students i, j_1, j_2, k_1, k_2 are distinct, and j_1 and j_2 rank above i at c because i has the lowest priority there, as does every other student.

It remains to select the scarcity sets, and the three groups from which they are drawn all contain the same students. Write $F = \{i, j_1, j_2, k_1, k_2\}$ and let $n = |I|$. Every student outside F ranks above i at c, because i has the lowest priority there; above k_1 at d_1, because only j_1 ranks below k_1 there and $j_1 \in F$; and above k_2 at d_2, because j_2 and k_2 are the two lowest-ranked students at d_2 among those outside $\{i, j_1, k_1\}$, a set that contains every student outside F. Any pairwise disjoint subsets of $I \setminus F$ of sizes $q_c - 2$, $q_{d_1} - 1$ and $q_{d_2} - 1$ therefore serve, and such subsets exist because those sizes add up to $q_c + q_{d_1} + q_{d_2} - 4$, which the hypothesis bounds by $n - 5 = |I \setminus F|$. Students of $I \setminus F$ left over simply belong to none of the three sets. Definition 8 is satisfied with both routes guarded, and Theorem 3 applies. □

Beyond the requirement on the number of students, the hypothesis is a statement about how different the priority orders of three schools are, and it fails only when the student ranked last at one of them is also ranked among the last four at one of the other two. When priorities are drawn independently across schools, that is unlikely, and the following is immediate.

Corollary 4. *Let c be a school with $q_c \geq 2$, let $d_1, d_2 \in C \setminus \{c\}$ be distinct, and suppose that $|I| \geq q_c + q_{d_1} + q_{d_2} + 1$. If $\succ_c$, $\succ_{d_1}$ and $\succ_{d_2}$ are drawn independently, each uniformly on the orders of I, then, whatever the priority orders of the other schools, the probability that Γ admits a separating double rotation is at least $(1 - 4/|I|)^2$.*

Proof. The student i with the lowest priority at c is determined by $\succ_c$, which is independent of $\succ_{d_1}$ and $\succ_{d_2}$. The rank of i under $\succ_{d_m}$ is therefore uniform on $\{1, \ldots, |I|\}$, and at least four students rank below i at d_m unless that rank exceeds $|I| - 4$, an event of probability $4/|I|$. The two events are independent, so the hypothesis of Proposition 6 holds with probability at least $(1 - 4/|I|)^2$. □

Note that the bound depends on the number of students alone. It holds for any number of schools and for any capacities under which the three schools leave at least one student without a seat, so the probability tends to one along every sequence of problems in which the number of students grows, including sequences in which the capacities grow with it, possibly proportionally. Thus, under independent uniform priorities the obstruction is common rather than exceptional, which is worth knowing before reading Theorem 3 as a curiosity about a five-student example. Two caveats limit the statement to what the proofs actually establish. The first is that the corollary counts problems and not preference profiles. It says that with high probability the priority structure admits some pair of

reports at which the conflict is forced, not that such reports are often submitted. The second is that neither the proposition nor the corollary characterizes compatibility. A priority structure whose orders nearly coincide need not be safe, since adding students who find no school acceptable at the bottom of every order leaves Γ^1 intact while making the orders agree everywhere else, and correlated priority structures fall entirely outside the scope of the corollary. This section gives a bound on each side, together with the statement that the region between them is not reached by accident. The only compatibility result it supports is Theorem 2, and Γ^2 shows the kind of structure that result covers, with the heterogeneity confined to schools that a correction never reaches.

6 Efficiency adjustments and consent

The efficiency adjustment that this literature has developed furthest is the Efficiency-Adjusted Deferred Acceptance (EADA) mechanism of Kesten (2010), which repairs DA by asking students to waive priorities whose violation cannot affect their own assignment. Cerrone, Hermstrüwer, and Kesten (2024) study consent under it in the laboratory, and report that the Flemish Ministry of Education took the first steps towards implementing it in 2019. This section reads our results through that mechanism. The question is not whether the repair can be made bossy, which Theorem 3 already answers, but how little consent it takes and whether it matters who gives it.

6.1 Efficiency-adjusted deferred acceptance

We state the mechanism in the simplified formulation of full-consent EADA due to Tang and Yu (2014), which produces the same matching as Kesten's mechanism when every student consents to waive the priorities that do not determine her own assignment.

Algorithm 2 Full-consent efficiency-adjusted deferred acceptance, EADA(P)

1: Run DA in the current market.
2: A school is *underdemanded* at the resulting matching if no student strictly prefers that school to her current assignment. Permanently settle every underdemanded school together with the students assigned to it, and remove those schools and students from the market. Students left unassigned are settled at the outside option, which is underdemanded by definition.
3: If students remain, rerun DA in the reduced market and repeat Step 2. Stop once all students and schools are settled.

This algorithm is well defined, and the matching it produces weakly Pareto dominates

DA and is Pareto efficient (Kesten, 2010; Tang & Yu, 2014), so full-consent EADA is DA-improving and Theorem 3 applies to it.

Corollary 5. *In every problem with a separating double rotation, full-consent EADA is locally bossy. In particular it is locally bossy in* Γ^1*, and in every problem with a twin cycle.*

In Γ^1 the outcome can also be seen directly. At P the only underdemanded school at DA(P) is c_4, so settling i_4 there and rerunning DA among $\{i_1, i_2, i_3, i_5\}$ over $\{c_1, c_2, c_3\}$ yields ν. At P' the only underdemanded school is again c_4, so settling i_5 there and rerunning DA among $\{i_1, i_2, i_3, i_4\}$ yields ν'. Student i_1 receives c_1 at both profiles and her classmate changes from i_3 to i_2.

6.2 Two waivers are enough

Full consent is a strong hypothesis. In Kesten's original formulation each student decides whether to waive priorities she cannot use, the mechanism only removes the interruptions caused by students who consent, and the outcome ranges from DA, when nobody consents, to the efficient correction, when everybody does. Write EADA^W for the mechanism obtained when the set of consenting students is $W \subseteq I$.

Algorithm 3 Efficiency-adjusted deferred acceptance with consent set W, EADA$^W(P)$

1: Run DA with the current lists. If no interrupting pair (i, c) with $i \in W$ remains, stop.
2: Otherwise, let t be the last round at which such a pair is rejected, and delete c from the list of i for every interrupting pair (i, c) with $i \in W$ rejected at t.
3: Return to Step 1.

Here (i, c) is an *interrupting pair* in a run of DA if c holds i tentatively from some round a onwards, c rejects some student other than i at a round $t \geq a$, and c rejects i at a later round. With $W = \emptyset$ the mechanism is DA, and with $W = I$ it produces the matching of Algorithm 2 (Kesten, 2010; Tang & Yu, 2014).

The deletion in Step 2 undoes the rejection that the consenting student caused. She never applies to that school, so the student she displaced there keeps the seat, and the chain of rejections that followed does not occur. Since consent is the instrument a district controls, the question is how much of it the conflict needs. The answer is two students, and we can name them.

Proposition 7. *Let* Γ *have a separating double rotation and let* (P, P') *be the pair of*

profiles constructed in the proof of Theorem 3. Let

$$W = \{k_1\} \cup \begin{cases} \{k_2\} & \text{if route 2 is guarded,} \\ \{g\} & \text{if route 2 is direct,} \end{cases}$$

where k_m is the intermediate student of route m and g is the classmate crowded out of c in Step 3 of that proof. Then W has two members, EADA^W selects at both profiles the unique element of the efficient frontier, and these two matchings assign i to c with different classmates, so EADA^W is locally bossy. The same holds for every $W' \supseteq W$.

Proof. Consider the run at P described in Step 1 of the proof of Theorem 3. Student k_1 is tentatively accepted at d_1 in the first round, school d_1 rejects j_1 in that same round, and k_1 is rejected by d_1 in the third round, when i applies. Hence (k_1, d_1) is an interrupting pair. Since $k_1 \in W$, the mechanism removes d_1 from her list and reruns. In the new run d_1 receives $N_{d_1} \cup \{j_1\}$, that is q_{d_1} applicants, and rejects nobody, so j_1 keeps d_1 and never applies to c. School c then receives $N_c \cup \{i\}$, together with j_2 when route 2 is guarded, which is exactly q_c applicants, and accepts them all. The outcome is the matching ν of Step 2, the unique element of $\mathcal{E}(P)$.

Consider now the run at P'. If route 2 is guarded, the same argument applies with (k_2, d_2) in place of (k_1, d_1) and yields ν'. If route 2 is direct, no student interrupts at d_2, because d_2 rejects nobody before i applies to it; the interruption occurs at c instead. Student g is tentatively accepted at c in the first round, school c rejects i in the second round, and g is rejected by c in the fourth round, when j_2 applies after being displaced from d_2. Hence (g, c) is an interrupting pair, and since $g \in W$ the mechanism removes c from her list. In the new run c receives $(N_c \setminus \{g\}) \cup \{j_1, i\}$, that is q_c applicants, and accepts them all, so i is never rejected by c and the outcome is the matching ν' of Step 4. In both cases EADA^W selects the unique efficient correction, and Step 5 of that proof shows that the two corrections assign i to c with different classmates.

For the last claim, note that at each of the two profiles the run of DA contains exactly one interrupting pair, namely the one used above. At P the only school that rejects two students is d_1. It holds $N_{d_1} \cup \{k_1\}$ when it rejects j_1, and of those students only k_1 is rejected by d_1 later, since the members of N_{d_1} keep their seats until the end. At P' the corresponding school is d_2 when route 2 is guarded and c when it is direct, and the same argument leaves k_2 and g as the only candidates. Enlarging W therefore does not change which pair is deleted. In the rerun that follows the deletion there is at most one rejection in the whole run. The two paragraphs above list the applications: every school receives at most its capacity, except the origin of the rival who remains at c when her route is

guarded, which receives one applicant too many and rejects that rival in the first round. No school therefore rejects a student it was already holding when it rejected somebody else, so no interrupting pair is left, the mechanism stops, and its outcome is the matching of Step 2, respectively Step 4. □

Only one of the two waivers does the work at each profile, the guard's at P and the other student's at P'. The proposition says that a single pair of students, fixed in advance, is enough to make the mechanism locally bossy at that pair of profiles. It does not say that two waivers are needed in every problem, nor that any two would do. The two students named have one thing in common, which is that neither of them can be made better off by withholding her consent. Student k_1 waives a priority at a school that admits her only until i applies, and g waives a priority at a school that admits her only until j_2 is pushed back into it. In both cases the seat was never theirs to keep. This is the sense in which Kesten (2010) shows that consenting is costless, and it concerns the school a student obtains. Whether a district can expect the waivers to be given is a separate question about information and about how students decide, which is one of the things the experimental evidence cited above studies. What the proposition says is that the waiver that costs its owner nothing is what allows the reporting student to change her own classmates while keeping her school. The waiver does not hand over a gain belonging to the consenting student, since she had none to give; rather, it frees up benefits for others, thereby transferring control over who sits in the contested school.

The two ends of the consent range are worth putting side by side. With no consent the mechanism is DA, which is locally non-bossy at every priority structure by Proposition 1. With the two waivers of Proposition 7, and so also with full consent, the protection is gone wherever Definition 8 is met. Within this mechanism the protection and the repair are therefore more than merely in tension, and the proposition adds that very little consent is needed to trade one for the other.

Remark 4. Partial consent does not restore the property by making the mechanism constrained efficient rather than efficient. Theorem 3 applies to DA-improving mechanisms, and EADA^W need not be one: with $W = \emptyset$ it is DA, which is not DA-improving in any problem in which DA is inefficient at some profile. Proposition 7 does not go through the theorem. It shows directly that EADA^W selects ν and ν' at the tight rigid pair constructed in the proof of Theorem 3. What the argument uses is only that the mechanism implements the available improvement at the two profiles, which is what any consent-based adjustment is designed to do.

The remark can be made precise, because what the proposition uses about EADA is

not the algorithm but two features of the profiles. At each of them DA admits a single improvement, and that improvement violates the priority of one student only, who belongs to W. We say that a matching μ *respects the priority* of student h at P if there is no school a such that $a\, P_h\, \mu_h$ and $h \succ_a h'$ for some $h' \in \mu_a$. Ortega and Arribillaga (2026) evaluate consent-based mechanisms by whether they exhaust the gains that consent makes available. Given a consent set W, let $\mathcal{M}(P, W)$ be the set of Pareto-efficient matchings that Pareto dominate $\mathrm{DA}(P)$ and respect the priority of every student outside W. A mechanism that uses the consent set W is *efficient whenever possible* if it selects an element of $\mathcal{M}(P, W)$ at every profile P at which this set is non-empty (Ortega & Arribillaga, 2026, Definition 6). They show that EADA^W does not have this property in general (their Proposition 2), so the next result is not a consequence of Proposition 7.

Proposition 8. *Let Γ have a separating double rotation, and let (P, P') and W be as in Proposition 7. For every $W' \supseteq W$, the sets $\mathcal{M}(P, W')$ and $\mathcal{M}(P', W')$ have a single element each, namely the unique elements of $\mathcal{E}(P)$ and $\mathcal{E}(P')$. Consequently every mechanism that uses the consent set W' and is efficient whenever possible is locally bossy.*

Proof. We use the notation of the proof of Theorem 3, with route 1 guarded, and recall that $\mathcal{D}(P) = \{\mathrm{DA}(P), \nu\}$ and $\mathcal{D}(P') = \{\mathrm{DA}(P'), \nu'\}$. A matching in $\mathcal{M}(P, W')$ is Pareto efficient and Pareto dominates $\mathrm{DA}(P)$, so it belongs to $\mathcal{E}(P) = \{\nu\}$, and similarly at P'. It remains to show that ν respects the priority of every student outside W at P, and ν' at P'; since $W \subseteq W'$, this gives $\nu \in \mathcal{M}(P, W')$ and $\nu' \in \mathcal{M}(P', W')$.

A student who finds no school acceptable prefers no school to her assignment. At P, students i and j_1 hold their first choices c and d_1 under ν, and every member of $N_c \cup N_{d_1} \cup N_{d_2}$, together with k_2 when route 2 is guarded, holds the only school she finds acceptable. Student j_2 holds her first choice d_2 when route 2 is direct. When it is guarded, she holds c, and the only school she prefers is d_2, whose members $N_{d_2} \cup \{k_2\}$ all rank above her there. The only remaining student is k_1, who is unassigned and outranks $j_1 \in \nu_{d_1}$ at d_1. Hence ν violates the priority of k_1 only, and $k_1 \in W$.

At P', students i and j_2 hold c and d_2, which are their first choices, and j_1 holds c and prefers only d_1, which is full with $N_{d_1} \cup \{k_1\}$, all ranked above her there. Every member of $N_{d_1} \cup N_{d_2}$, every member of N_c other than g when route 2 is direct, and k_1 hold the only school they find acceptable. The only remaining student is the one displaced in Step 3 of that proof, who is unassigned: she is k_2, who outranks $j_2 \in \nu'_{d_2}$ at d_2, when route 2 is guarded, and g, who outranks $i \in \nu'_c$ at c, when route 2 is direct. In both cases she is the second member of W. A mechanism that uses the consent set W' and is efficient whenever possible therefore selects ν at P and ν' at P', and these assign i to c with different classmates. □

6.3 Consent that the district does not choose

Proposition 7 names the two students whose waivers break the protection. A district does not choose who consents, and the next result shows that the waivers are found, with high probability, in any consent set that is not too small and is unrelated to priorities.[9] Recall from Definition 9 that a twin cycle is a separating double rotation in which both routes are guarded, so that the set W of Proposition 7 consists of its two guards.

Proposition 9. *Let c, d_1, d_2 be distinct schools with $q_c \geq 2$, let $n = |I|$ and suppose that $q_c + q_{d_1} + q_{d_2} \leq n - 1$. Let $W \subseteq I$ have w members, and draw $\succ_c$, $\succ_{d_1}$ and $\succ_{d_2}$ independently, each uniformly on the orders of I. Then, for every integer $m \geq 1$ with $2m + 3 \leq n - q_{d_1} - q_{d_2}$, with probability at least*

$$1 - \frac{4m+4}{n} - 2\Big(1 - \frac{w}{n}\Big)^m$$

the school choice problem Γ has a twin cycle at c whose two guards belong to W. In that event, EADA^W is locally bossy, and so is every mechanism that uses the consent set W and is efficient whenever possible.

The proof is in the Appendix. It looks for the guards among the students ranked just above the bottom of d_1 and of d_2, in windows of m students, and uses the remaining students to fill the scarcity sets of Definition 8. This is why the windows must leave enough students outside them. When W is random and independent of the three priority orders, the bound holds conditionally on W. The asymptotic statement follows by letting the windows grow slowly.

Corollary 6. *Consider a sequence of problems with $n \to \infty$ students and three schools c, d_1, d_2 with $q_c \geq 2$ and $n - q_c - q_{d_1} - q_{d_2} \to \infty$, with the priority orders of the three schools drawn as in Proposition 9, and let W_n be fixed in advance, or drawn independently of those orders, with $|W_n| \geq \delta n$ for some $\delta > 0$. Then, as $n \to \infty$, the probability that EADA^{W_n} is locally bossy tends to one, and so does the probability that every mechanism that uses the consent set W_n and is efficient whenever possible is locally bossy.*

Proof. For all n large enough, $q_c + q_{d_1} + q_{d_2} \leq n - 1$. Let $m_n = \min\{\lfloor (n - q_{d_1} - q_{d_2} - 3)/2 \rfloor, \lfloor \sqrt{n} \rfloor\}$. Since $n - q_{d_1} - q_{d_2} \geq n - q_c - q_{d_1} - q_{d_2} + 2$, we have $m_n \to \infty$, so $m_n \geq 1$ for all n large enough, and $2m_n + 3 \leq n - q_{d_1} - q_{d_2}$. The bound of Proposition 9 is then at least $1 - (4\sqrt{n} + 4)/n - 2(1 - \delta)^{m_n}$, which tends to one. □

[9] I thank Josué Ortega for suggesting this result and the idea of its proof.

Three caveats apply to this analysis. The first is that consent must be unrelated to priorities. The argument finds consenting students among those ranked just above the bottom of two schools, and a population in which exactly those students withhold consent escapes it. The second is the one made after Corollary 4: the statement counts problems and not profiles. The third is that the fraction δ is more than the proof needs. Since $(1 - w/n)^m \leq e^{-mw/n}$, the bound still tends to one when $|W_n|$ grows more slowly than n, as long as $m_n|W_n|/n \to \infty$ for some sequence m_n with $2m_n + 3 \leq n - q_{d_1} - q_{d_2}$ and $m_n/n \to 0$.

6.4 A margin the incentive criteria do not measure

Since no efficient DA-dominating mechanism is strategy-proof (Abdulkadiroğlu, Pathak, & Roth, 2009; Kesten, 2010), the literature evaluating efficiency adjustments has replaced strategy-proofness by weaker requirements. EADA is not obviously manipulable (Troyan & Morrill, 2020), truth-telling under it is maxmin optimal (Reny, 2022), it is immune to simple misreports (Shirakawa, 2025), and it is upper-manipulation-proof (Afacan et al., 2026). Chen and Möller (2024) show that the version of the mechanism in which students also report consent decisions is regret-free truth-telling under incomplete information, and that no efficient rule dominating a stable rule is regret-free truth-telling once those decisions are absent. Each of these criteria is stated for the school a student is assigned to.

The deviation of this paper lies outside what they measure, and not because it is small. At a rigid pair every efficient DA-dominating mechanism, and at a tight one every DA-improving mechanism, gives the deviating student the same school under both reports, so her gain in the ranking of schools is exactly zero under these two reports, and a criterion that compares the school she obtains will correctly record no manipulation while the reported list has changed the cohort of the school she attends. This does not say that those criteria fail, nor that they could not be restated once preferences include classmates. We have not studied their status on such a domain. It says that their guarantees, stated for the school a student receives, leave the present margin unaddressed, and that this margin lies in the dimension the literature on peer effects and on segregation treats as the outcome of interest.

6.5 Preferences over classmates

Duque-Rosas, Pereyra, and Torres-Martínez (2026) consider *school-lexicographic preferences over colleagues*. Students first compare matchings by the school they receive. Conditional on receiving the same school, they may care about that school's composition, but they are

indifferent between two matchings that give them both the same school and the same set of schoolmates. Let $\mathcal{D}^c$ denote the corresponding domain of students' preference profiles over matchings, and let $P(\trianglerighteq)$ be the strict ranking of schools induced by a profile $\trianglerighteq \in \mathcal{D}^c$. A mechanism φ defined on school rankings induces a mechanism on this richer domain by

$$\bar{\varphi}(\trianglerighteq) = \varphi(P(\trianglerighteq)).$$

Duque-Rosas, Pereyra, and Torres-Martínez (2026) show that the mechanism induced by DA is stable and strategy-proof on $\mathcal{D}^c$. The two possible motives for a lie are handled separately. A student may lie to change her school, which ordinary strategy-proofness prevents, or she may lie while keeping her school in order to change her classmates, which local non-bossiness prevents.

Corollary 7. *Let (P, P') be a rigid pair for a student i and let φ be efficient and DA-dominating. Then $\bar{\varphi}$ is not strategy-proof on $\mathcal{D}^c$, and the profitable deviation can be made by i, whose school is the same under both reports for every mechanism in this class.*

Proof. Let x be the school that every admissible correction assigns to i at both profiles, and let $\nu = \varphi(P)$ and $\nu' = \varphi(P')$, so that $\nu_i = \nu'_i = x$ and $\nu_x \neq \nu'_x$. Choose $\trianglerighteq_i \in \mathcal{D}^c$ inducing P_i and such that, among matchings assigning i to x, student i strictly prefers the cohort ν'_x to the cohort ν_x, which is possible because the two cohorts differ and school-lexicographic preferences leave the comparison of cohorts at a given school unrestricted. Complete the profile with any $\trianglerighteq_{-i}$ inducing P_{-i}. Then $\bar{\varphi}(\trianglerighteq) = \nu$, while reporting any $\trianglerighteq'_i \in \mathcal{D}^c$ inducing P'_i gives ν', and $\nu' \rhd_i \nu$ by construction. In Γ^1 the deviation is the one of Theorem 1, made by i_1 at c_1, whose classmate changes from i_3 to i_2; by Theorem 3 the same argument applies in every problem with a separating double rotation. □

The proof uses only that φ selects an element of $\mathcal{E}(P)$ at P and of $\mathcal{E}(P')$ at P'. When the rigid pair is tight, as in Theorems 1 and 3, every DA-improving mechanism does so by Corollary 2, and the conclusion extends to all of them.

It is also worth recording what a welfare correction itself does on this domain, since that comparison is not the one the rest of the paper makes.

Remark 5. Fix P, let $\mu = \mathrm{DA}(P)$ and let ν weakly Pareto dominate μ in the rankings of schools. Then ν weakly Pareto dominates μ under every extension of P in $\mathcal{D}^c$ if and only if every student with $\nu_i = \mu_i$ also has $\nu_{\mu_i} = \mu_{\mu_i}$. One direction is immediate: a student who changes school is strictly better off in the ranking of schools and the lexicographic criterion settles her comparison, while a student who keeps her school and her set of schoolmates is indifferent, because preferences in $\mathcal{D}^c$ do not separate two matchings that give her both.

For the converse, a student who keeps her school and sees that set change can be given an extension in $\mathcal{D}^c$ that ranks the old set above the new one. Unassigned students raise no issue, since $\nu_\emptyset = \mu_\emptyset$ by Lemma 1(i).

The correction that Theorem 1 forces is of the second kind. Passing from DA(P) to ν in Γ^1 makes i_1 and i_2 strictly better off in their rankings of schools, and leaves i_3 at c_1 with i_1 in place of i_2 as her classmate. If i_3 ranks the cohort $\{i_2, i_3\}$ above $\{i_1, i_3\}$, the repair of DA's inefficiency makes her worse off. The efficiency notion of Definition 1 is stated in the rankings of schools, which is the domain on which the mechanisms of this literature are defined, and the remark says what is being left out. Students who move to a better school gain under every extension in $\mathcal{D}^c$. For a student who keeps her school, a gain that holds under every such extension requires her set of schoolmates to be unchanged. The invariance this asks for between two matchings at one profile is the same one that local non-bossiness asks for between two profiles, although the two comparisons are different exercises.

The deviation is as small as Proposition 3 allows, since student i_1 moves one school from just below to just above the school DA assigns her and changes nothing else. It leaves her own school unchanged under every mechanism in the class, so it is not a manipulation that the criteria of Section 6.4 can register, and on the domain where composition enters preferences it is strictly profitable. What it moves is the identity of a classmate. A family that cares only about the proportions of students of each observable type in its school gains from it only when the two classmates it can choose between differ in that type, and Section 7 shows when that happens.

6.6 Three properties that fail only together

The impossibility is joint, and dropping any one of the three requirements resolves it. Dropping efficiency leaves DA, which is DA-dominating and locally non-bossy, and DA is also what remains when the weaker requirement of DA-improvement is dropped. Dropping DA-dominance leaves serial dictatorship, or any other efficient and group strategy-proof mechanism, which is non-bossy and therefore locally non-bossy. Dropping local non-bossiness leaves the whole class of efficient DA-dominating mechanisms, which is non-empty at every problem by Lemma 2. The paper shows that the three cannot be had together, and Section 5 bounds the problems where this happens from both sides.

The comparison with DA is the one to keep in mind. DA satisfies two of the three requirements and is locally non-bossy at every priority structure. Repairing its inefficiency is what breaks the property, and it breaks it along the dimension the repair does not touch,

since the manipulator's own school is the same under both reports and what her report moves is the set of students she enters with.

7 Composition by type

The deviation behind our results changes the identity of a classmate. Districts usually follow the composition of their schools through observable characteristics, such as income, ethnicity or prior achievement, and the literature on segregation measures composition by the number of students of each type that a school admits (Escobar & Huerta, 2026; Ortega et al., 2026). Whether a report that leaves its author at the same school can change those numbers depends on how types are related to priorities, and the two extreme cases give opposite answers.

A *type function* is a map $\tau : I \to \{1, \dots, K\}$, and $T_k = \{h \in I : \tau(h) = k\}$ is the set of students of type k. We write $U_k = T_1 \cup \cdots \cup T_k$, with $U_0 = \emptyset$. The priority profile is *stratified by* τ if every school ranks every student of a type with a lower index above every student of a type with a higher index, that is, if $\tau(h) < \tau(h')$ implies $h \succ_c h'$ at every school c. Within a type, schools may rank students in any way. This is the tiered priority structure of Ortega et al. (2026), who give as examples the precedence of citizens over permanent residents and international students in Singapore's primary-school admissions, and the precedence of looked-after children under the School Admissions Code in England.

7.1 Stratified priorities

Ortega et al. (2026) show that under stratified priorities every rule that dominates DA admits the same number of students of each type to each school as DA does. We use the result in the following form, and include a short proof because the argument behind the next proposition builds on it.

Lemma 6 (Ortega et al. (2026), Theorem 2)**.** *Let the priority profile be stratified by τ, let μ be stable at P, and let ν weakly Pareto dominate μ at P. Then $|\nu_x \cap T_k| = |\mu_x \cap T_k|$ for every $x \in C \cup \{\emptyset\}$ and every type k.*

Proof. Fix k and a school x. Suppose first that some student $h \in U_k$ satisfies $\nu_h = x \neq \mu_h$. Then $x \, P_h \, \mu_h$ by Lemma 1(iii), and since (h, x) does not block μ, school x is full at μ and every member of μ_x has higher priority than h at x. Under stratification a student with higher priority than h has a type with index at most $\tau(h) \leq k$, so $\mu_x \subseteq U_k$, and Lemma 1(i) gives $|\nu_x \cap U_k| \leq |\nu_x| = |\mu_x| = |\mu_x \cap U_k|$. If instead no student of U_k moves into x, then $\nu_x \cap U_k \subseteq \mu_x \cap U_k$. In both cases $|\nu_x \cap U_k| \leq |\mu_x \cap U_k|$. Since $\nu_\emptyset = \mu_\emptyset$ by

Lemma 1(i), the number of students of U_k assigned to some school is the same at ν and at μ, so the inequalities, summed over $x \in C$, add up to an equality, and each of them holds with equality. Taking differences in k gives the claim for every school, and for the outside option it follows from $\nu_\emptyset = \mu_\emptyset$. □

Their statement is for $\mu = \mathrm{DA}(P)$, which is the case used below. The second ingredient is that, under stratified priorities, DA treats the students of the first k types as if the others were absent. This is the sequential description that Ortega et al. (2026) give of the rules they study, stated here for DA; the proof is in the Appendix.

Lemma 7. *Let the priority profile be stratified by τ, fix k, and let P and P' be preference profiles with $P_h = P'_h$ for every $h \in U_k$. Then $\mathrm{DA}_h(P) = \mathrm{DA}_h(P')$ for every $h \in U_k$.*

The two lemmas extend the invariance of Ortega et al. (2026) from the matchings available at one profile to the reports of one student.

Proposition 10. *Let the priority profile be stratified by τ, and let φ be DA-dominating. If $\varphi_i(P) = \varphi_i(P'_i, P_{-i}) = x$ for some student i, some profile P and some report P'_i, then $\varphi_x(P)$ and $\varphi_x(P'_i, P_{-i})$ contain the same number of students of each type.*

Proof. Write $P' = (P'_i, P_{-i})$, $\mu = \mathrm{DA}(P)$, $\mu' = \mathrm{DA}(P')$, $\nu = \varphi(P)$ and $\nu' = \varphi(P')$. By Lemma 6, ν_x and μ_x contain the same number of students of each type, and so do ν'_x and μ'_x. It is therefore enough to show that μ_x and μ'_x do.

If $x = \emptyset$, Lemma 1(i) gives $\mu_\emptyset = \nu_\emptyset$ and $\mu'_\emptyset = \nu'_\emptyset$, so i is unassigned at μ and at μ', and Proposition 1 gives $\mu_\emptyset = \mu'_\emptyset$. Let $x \in C$. If $\mu_i = \mu'_i = x$, Proposition 1 gives $\mu_x = \mu'_x$. If $\mu_i = x \neq \mu'_i$, then $\nu'_i = x \neq \mu'_i$, and Lemma 1(iii) applied at P' gives $x \, P'_i \, \mu'_i$; a student with preferences P'_i would then obtain x instead of μ'_i by reporting P_i, which contradicts the strategy-proofness of DA. The case $\mu'_i = x \neq \mu_i$ is symmetric. It remains to consider $\mu_i \neq x$ and $\mu'_i \neq x$. Then $x \, P_i \, \mu_i$ by Lemma 1(iii), and since (i, x) does not block μ, school x is full at μ and every member of μ_x has higher priority than i at x, hence a type with index at most $t = \tau(i)$. The same holds at μ'. Since only i changes her report, P and P' coincide on U_{t-1}, and Lemma 7 gives $\mu_x \cap U_{t-1} = \mu'_x \cap U_{t-1}$. The two cohorts therefore have the same members of types with index below t, both have q_x members, and neither has a member of a type with index above t, so they also have the same number of students of type t. □

Under stratified priorities every DA-dominating mechanism is therefore locally non-bossy at the level of types, whether it is efficient or not and whatever consent it relies on, since EADA^W weakly Pareto dominates DA for every consent set (Kesten, 2010).

A student can still change who her classmates are, and Theorem 3 applies to stratified priorities as to any others, but she cannot change how many of them are of each type. The separating double rotation shows why. Its requirements $j_m \succ_c i$ and $i \succ_{d_m} j_m$ force j_m to have the type of i, a guard lies between i and j_m at d_m, and the classmate g of a direct route lies between j_1 and i at c. Under stratified priorities the manipulator, the two rivals, the guards and the classmate g therefore all have the same type, and the classmate her report removes is replaced by another student of that type.

7.2 Types unrelated to priorities

Stratification is a strong assumption, as Ortega et al. (2026) acknowledge. In most systems the types a district monitors are not priority classes: priorities come from siblings, distance and lotteries, and where disadvantage enters the assignment it may do so through reserved seats rather than through absolute precedence, as in the Chilean system studied by Escobar and Huerta (2026). Without stratification the conclusion of Proposition 10 fails, and the construction behind Theorem 3 shows exactly how.

Corollary 8. *Let Γ have a separating double rotation, labelled so that route 1 is guarded, and let (P, P'), ν and ν' be as in the proof of Theorem 3. If both routes are guarded, then $\nu_c \setminus \nu'_c = \{j_2\}$ and $\nu'_c \setminus \nu_c = \{j_1\}$. If route 2 is direct, then $\nu_c \setminus \nu'_c = \{g\}$ and $\nu'_c \setminus \nu_c = \{j_1\}$. Consequently, for every type function τ, the two cohorts of c contain different numbers of students of some type if and only if $\tau(j_1) \neq \tau(j_2)$ in the first case, and $\tau(j_1) \neq \tau(g)$ in the second.*

Proof. By Steps 2, 4 and 5 of the proof of Theorem 3, $\nu_c = N_c \cup \{i, j_2\}$ and $\nu'_c = N_c \cup \{i, j_1\}$ if both routes are guarded, while $\nu_c = N_c \cup \{i\}$, which contains g, and $\nu'_c = (N_c \setminus \{g\}) \cup \{i, j_1\}$ if route 2 is direct. Since i, j_1 and j_2 do not belong to N_c, the two cohorts share all their other members and differ in one student each, and their numbers of students of each type differ exactly when those two students have different types. □

Whether the two students differ in type depends on the problem. The next result shows that, with uniform priorities and types fixed in advance, problems in which they do are the rule rather than the exception. Types play no role in the priorities here, and the only requirement on them is that no type contain almost every student.

Proposition 11. *Let c, d_1, d_2 be distinct schools with $q_c \geq 2$ and $q_{d_2} \leq q_{d_1}$, let $n = |I|$ and suppose that $q_c + q_{d_1} + q_{d_2} \leq n - 1$. Let τ be a type function in which no type has more than $(1-\varepsilon)n$ members, for some $\varepsilon \in (0, 1)$, and draw $\succ_c$, $\succ_{d_1}$ and $\succ_{d_2}$ independently, each*

uniformly on the orders of I. Then, for every integer m with $1 \leq m \leq (n-5)/2$, with probability at least

$$1 - \frac{m+6}{n} - \left(1 - \varepsilon + \frac{1}{n}\right)^{m+1}$$

the problem Γ admits a tight rigid pair for some student i at which every DA-improving mechanism assigns i to c under both reports and the two cohorts of c contain different numbers of students of some type.

The proof is in the Appendix. Taking $m = \lfloor \sqrt{n} \rfloor$, the bound tends to one as n grows, for any fixed ε, and it does not depend on the number of schools or on their capacities beyond the requirement that the three schools leave a student without a seat. The labelling $q_{d_2} \leq q_{d_1}$ is without loss of generality, since the roles of d_1 and d_2 are symmetric.

The two results describe the two ends of the relation between types and priorities. When the types a district monitors are themselves priority classes, a reform that improves upon DA preserves the number of students of each type at every school, across the matchings available at one profile, as Ortega et al. (2026) show, and across the reports of any single student, as Proposition 10 shows. Families can then reshuffle identities but not proportions. When types and priorities are unrelated, the same reforms hand a family that keeps its school a choice between classmates of different types in almost every large problem. Priority structures that favour a type at some schools only, or reserves of the kind studied by Escobar and Huerta (2026), lie between the two cases and are covered by neither result.

8 Concluding remarks

The literature on improvements over Deferred Acceptance has asked what such improvements must give up in stability, in incentives, in consistency or in the size of the gains they deliver. This paper adds a requirement of a different kind, about who a student is assigned with rather than where she is assigned. It is a requirement that DA itself satisfies at every priority structure, while full non-bossiness costs acyclicity.

The protection does not survive the repair, and it is lost for every mechanism that implements an improvement over DA whenever one exists, whether or not it exhausts the available gains. Theorem 1 shows this in a balanced problem with five students, four schools and complete preference lists, where the impossibility owes nothing to outside options or to unassigned students. Section 5 places the problems in which the protection is lost between two conditions on priorities and capacities. The conflict requires first that priorities let a welfare correction reach a school with several seats at all, that is, a

chain of denied claims that closes on such a school with each of its links filled by students with a stronger claim. When there is no such rotation, every improvement over DA reproduces the DA cohort at every school with several seats. When the conflict comes from a rigid pair, as every conflict exhibited here does, it requires second that the chain be available to the manipulator from two different schools, because at a rigid pair DA denies her the contested school under both reports and sends her to two different schools. What produces the conflict is a double rotation whose two routes are separating, meaning that each of them is closed while the manipulator travels the other, and priorities can close a route either through an intermediate student or through the manipulator herself. Proposition 6 and Corollary 4 show that this configuration is common when priorities are drawn independently and uniformly, also when capacities grow with the market, so the impossibility is not an artefact of the examples.

Two further features bear on how a district might read the result. Consent does not contain it. Under EADA two waivers that cost their owners nothing are enough, and so, with probability approaching one, is the consent of a fixed fraction of students unrelated to priorities. The same holds for every consent-based mechanism that is efficient whenever possible. Furthermore, the margin is not always confined to identities. When the types a district monitors are priority classes, no improvement over DA lets a family change how many students of each type its school admits, but when types and priorities are unrelated, the family can choose between classmates of different types in almost every large problem. Which of the two descriptions fits a given district is an empirical question about its priority rules, and the answer decides whether the bossiness that efficiency adjustments introduce reaches the composition measures that segregation policy targets.

References

Abdulkadiroğlu, A., Pathak, P. A., & Roth, A. E. (2009). Strategy-proofness versus efficiency in matching with indifferences: Redesigning the NYC high school match. *American Economic Review*, *99*(5), 1954–1978.

Abdulkadiroğlu, A., Pathak, P. A., Schellenberg, J., & Walters, C. R. (2020). Do parents value school effectiveness? *American Economic Review*, *110*(5), 1502–1539.

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

Afacan, M. O., Dur, U., Gitmez, A. A., & Yılmaz, Ö. (2026). Improving the deferred acceptance with minimal compromise. *Games and Economic Behavior*, *156*, 64–81.

Afacan, M. O., & Dur, U. M. (2017). When preference misreporting is harm[less]ful? *Journal of Mathematical Economics*, *72*, 16–24.

Alva, S., & Manjunath, V. (2019). Strategy-proof Pareto-improvement. *Journal of Economic Theory*, *181*, 121–142.

Balinski, M., & Sönmez, T. (1999). A tale of two mechanisms: Student placement. *Journal of Economic Theory*, *84*(1), 73–94.

Beuermann, D. W., Jackson, C. K., Navarro-Sola, L., & Pardo, F. (2023). What is a good school, and can parents tell? Evidence on the multidimensionality of school output. *The Review of Economic Studies*, *90*(1), 65–101.

Burgess, S., Greaves, E., Vignoles, A., & Wilson, D. (2015). What parents want: School preferences and school choice. *The Economic Journal*, *125*(587), 1262–1289.

Cerrone, C., Hermstrüwer, Y., & Kesten, O. (2024). School choice with consent: An experiment. *The Economic Journal*, *134*(661), 1760–1805.

Chen, Y., & Möller, M. (2024). Regret-free truth-telling in school choice with consent. *Theoretical Economics*, *19*(2), 635–666.

Correa, J., Epstein, N., Epstein, R., Escobar, J., Ríos, I., Aramayo, N., Bahamondes, B., Bonet, C., Castillo, M., Cristi, A., Epstein, B., & Subiabre, F. (2022). School choice in Chile. *Operations Research*, *70*(2), 1066–1087.

Cox, N., Fonseca, R., Pakzad-Hurson, B., & Pecenco, M. (2025). *Peer preferences in centralized school choice markets: Theory and evidence. Working paper.*

Doğan, B., & Ehlers, L. (2021). Minimally unstable Pareto improvements over deferred acceptance. *Theoretical Economics*, *16*(4), 1249–1279.

Doğan, B., & Yenmez, M. B. (2020). Consistent Pareto improvement over the student-optimal stable mechanism. *Economic Theory Bulletin*, *8*(1), 125–137.

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

Duque-Rosas, E., Pereyra, J. S., & Torres-Martínez, J. P. (2026). Local non-bossiness. *Journal of Economic Theory*, *236*, 106230.

Erdil, A. (2014). Strategy-proof stochastic assignment. *Journal of Economic Theory*, *151*, 146–162.

Ergin, H. I. (2002). Efficient resource allocation on the basis of priorities. *Econometrica*, *70*(6), 2489–2497.

Escobar, J. F., & Huerta, L. (2026). Segregation and affirmative action in school choice. *Journal of Political Economy: Microeconomics*, *4*(2), 410–460.

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

Kesten, O. (2010). School choice with consent. *Quarterly Journal of Economics*, *125*(3), 1297–1348.

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

Knipe, T., & Ortega, J. (2025). *The trade-off between minimal instability and larger improvements over deferred acceptance. Working paper*, arXiv 2504.12871.

Kojima, F. (2010). Impossibility of stable and nonbossy matching mechanisms. *Economics Letters*, *107*(1), 69–70.

Moreno-García, E., & Torres-Martínez, J. P. (2026). Coalitional manipulability in school choice. *Theoretical Economics*, *21*(3), 669–690.

Ortega, J., & Arribillaga, R. P. (2026). *Justifiable priority violations. Working paper*, arXiv 2604.06396.

Ortega, J., Ziegler, G., Arribillaga, R. P., & Zhao, G. (2026). What Pareto-efficiency adjustments cannot fix. *Journal of Political Economy: Microeconomics*. Forthcoming.

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

Raghavan, M. (2020). Influence in private-goods allocation. *Journal of Mathematical Economics*, *89*, 14–28.

Reny, P. J. (2022). Efficient matching in the school choice problem. *American Economic Review*, *112*(6), 2025–2043.

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

Rothstein, J. (2006). Good principals or good peers? Parental valuation of school characteristics, Tiebout equilibrium, and the incentive effects of competition among jurisdictions. *American Economic Review*, *96*(4), 1333–1350.

Satterthwaite, M. A., & Sonnenschein, H. (1981). Strategy-proof allocation mechanisms at differentiable points. *Review of Economic Studies*, *48*(4), 587–597.

Shirakawa, R. (2025). Simple manipulations in school choice mechanisms. *American Economic Journal: Microeconomics*, *17*(4), 202–233.

Tang, Q., & Yu, J. (2014). A new perspective on Kesten's school choice with consent idea. *Journal of Economic Theory*, *154*, 543–561.

Tang, Q., & Zhang, Y. (2021). Weak stability and Pareto efficiency in school choice. *Economic Theory*, *71*(2), 533–552.

Thomson, W. (2016). Non-bossiness. *Social Choice and Welfare*, *47*(3), 665–696.

Troyan, P., Delacrétaz, D., & Kloosterman, A. (2020). Essentially stable matchings. *Games and Economic Behavior*, *120*, 370–390.

Troyan, P., & Morrill, T. (2020). Obvious manipulations. *Journal of Economic Theory*, *185*, 104970.

Appendix

This Appendix contains the material omitted from Section 5, in the order in which the main text calls for it. Lemma 8 verifies that the problem Γ^2, which Theorem 2 places outside the impossibility, admits no rotation through its three-seat school. Theorem 3 is then proved. The rest presents the problems Γ^4 and Γ^5, which separate the conditions of Figure 4 from one another and show what an exact description of the conflict would have to accommodate; Proposition 12 and Proposition 13 record the impossibility in each of them, and Lemma 9 shows that the first is outside the reach of the guarded case. The Appendix closes with the proofs of Proposition 9, Lemma 7 and Proposition 11.

We begin with the problem Γ^2 of Section 5.1, whose three-seat school no welfare correction can reach.

Lemma 8. *The problem Γ^2 of Section 5.1 admits no rotation through c_1.*

Proof. Suppose a rotation through c_1 exists, with schools $a_1 = c_1, \dots, a_L$ and students $k_1, \dots, k_L$ as in Definition 6. The requirement at a_1 provides $|N_{c_1}| = q_{c_1} - 1 = 2$ students ranked above k_L at c_1 and distinct from k_1, together with k_1 itself, so at least three students rank above k_L at c_1 and therefore $k_L \in \{i_4, i_5\}$.

If $L = 2$ and $a_2 = c_2$, the requirements $k_1 \succ_{c_1} k_2$ and $k_2 \succ_{c_2} k_1$ are incompatible because $\succ_{c_1}$ and $\succ_{c_2}$ coincide. If $L = 2$ and $a_2 = c_3$, then $k_2 = k_L \in \{i_4, i_5\}$ and $k_2 \succ_{c_3} k_1$; no student ranks below i_5 at c_3, so $k_2 = i_4$ and $k_1 = i_5$, contradicting $k_1 \succ_{c_1} k_2$.

If $L = 3$, then $\{a_2, a_3\} = \{c_2, c_3\}$ and $k_3 = k_L \in \{i_4, i_5\}$. Suppose $(a_2, a_3) = (c_2, c_3)$. Since $\succ_{c_1}$ and $\succ_{c_2}$ coincide, the requirements $k_2 \succ_{c_2} k_1$ and $k_1 \succ_{c_1} k_3$ place k_2, k_1 and k_3 in this order in the common ranking, and we also need $k_3 \succ_{c_3} k_2$. If $k_3 = i_5$ this fails because no student ranks below i_5 at c_3. If $k_3 = i_4$ then $k_2 \in \{i_1, i_2\}$, and neither satisfies $i_4 \succ_{c_3} k_2$. Suppose instead $(a_2, a_3) = (c_3, c_2)$. Then $k_3 \succ_{c_2} k_2$ and $k_1 \succ_{c_1} k_3$ place k_1, k_3 and k_2 in this order in the common ranking, and we also need $k_2 \succ_{c_3} k_1$. If $k_3 = i_5$ there is no student below i_5 available for k_2; if $k_3 = i_4$ then $k_2 = i_5$, and no student ranks below i_5 at c_3. In every case a rotation through c_1 is impossible. □

We turn to the sufficiency result. The construction is the one sketched in Section 5.2: the manipulator declares acceptable only the contested school and one of her two origins, and every other student involved declares only the schools her role requires. Throughout, $\beta \in \{0, 1\}$ counts the direct routes.

Proof of Theorem 3. Fix a separating double rotation with the notation of Definition 8, relabelling the routes so that route 1 is guarded. Consider the preference profile P in which

$$P_i:\ c,\, d_1, \qquad P_{j_1}:\ d_1,\, c, \qquad P_{j_2}:\ d_2,\, c, \qquad P_{k_m}:\ d_m \ \text{ for each guarded route } m,$$

every member of N_c finds only c acceptable, every member of N_{d_m} finds only d_m acceptable, and every remaining student finds no school acceptable. Let P' be obtained from P by replacing the report of i with $P'_i:\ c,\, d_2$. The two profiles differ only in the second school on the list of i.

Throughout, call a student *fixed* if she finds at most one school acceptable. Every member of $N_c \cup N_{d_1} \cup N_{d_2}$, every k_m and every student outside $\{i, j_1, j_2, k_1, k_2\}$ is fixed, and each of them is assigned to her unique acceptable school at any profile at which she is not rejected from it.

Step 1: the DA outcome at P. In the first round school c receives i and the members of N_c, that is $q_c - 1 + \beta \leq q_c$ applicants, and rejects nobody. School d_1 receives j_1, k_1 and the members of N_{d_1}, that is $q_{d_1} + 1$ applicants, and rejects the applicant with the lowest priority, which is j_1 because every member of N_{d_1} ranks above k_1 and $k_1 \succ_{d_1} j_1$. School d_2 receives j_2, the members of N_{d_2} and, if route 2 is guarded, also k_2. If route 2 is guarded it has $q_{d_2} + 1$ applicants and rejects j_2, by the same argument. If route 2 is direct it has exactly q_{d_2} applicants and rejects nobody, so j_2 keeps d_2, her most preferred school, and never applies to c.

In the second round j_1 applies to c, and so does j_2 if route 2 is guarded. School c now has $q_c + 1$ applicants, namely $N_c \cup \{i, j_1\}$ together with j_2 when route 2 is guarded, and it rejects i, since every other applicant ranks above her. In the third round i applies to d_1, which is full with $N_{d_1} \cup \{k_1\}$, and displaces k_1 because $i \succ_{d_1} k_1$ and every member of N_{d_1} ranks above k_1. Student k_1 has no further acceptable school. Hence

$$\mathrm{DA}(P):\quad c = N_c \cup \{j_1\} \cup \{j_2 :\ \text{route 2 guarded}\},\quad d_1 = N_{d_1} \cup \{i\},\quad d_2 = N_{d_2} \cup \{k_2 \text{ or } j_2\},$$

with k_1 and all remaining students unassigned.

Step 2: the efficient frontier at P. Every fixed student other than k_1 holds her most preferred option in $C \cup \{\emptyset\}$, so she is fixed in every element of $\mathcal{D}(P)$, and by Lemma 1(i) the unassigned student k_1 stays unassigned. If route 2 is guarded, d_2 is full with fixed students, so j_2 cannot move and remains at c; if route 2 is direct, j_2 already holds her most preferred school. In either case the only students who can move are i and j_1, and by Lemma 1(i) every element of $\mathcal{D}(P)$ assigns q_c students to c and q_{d_1} students to d_1.

One seat at c and one seat at d_1 are contested, and each of them must go to i or to j_1. Assigning d_1 to i returns $\mathrm{DA}(P)$; assigning c to i and d_1 to j_1 gives a matching ν that makes both of them strictly better off. Hence $\mathcal{D}(P) = \{\mathrm{DA}(P), \nu\}$ and $\mathcal{E}(P) = \{\nu\}$, with

$$\nu_i = c, \qquad \nu_c = N_c \cup \{i\} \cup \{j_2 : \text{ route 2 is guarded}\}.$$

Step 3: the DA outcome at P'. The first two rounds are exactly as in Step 1, since i applies to c in both and has not yet reached the school she names second: in the second round school c again has $q_c + 1$ applicants and rejects i, so that c is left with $N_c \cup \{j_1\}$, together with j_2 when route 2 is guarded. In the third round i applies to d_2.

If route 2 is guarded, d_2 is full with $N_{d_2} \cup \{k_2\}$ and i displaces k_2, because $i \succ_{d_2} k_2$ and every member of N_{d_2} ranks above k_2; student k_2 has no further acceptable school, the run ends, and c holds $N_c \cup \{j_1, j_2\}$. If route 2 is direct, d_2 is full with $N_{d_2} \cup \{j_2\}$ and i displaces j_2, because $i \succ_{d_2} j_2$ and every member of N_{d_2} ranks above j_2. Student j_2 then applies to c in the fourth round, where the applicants ranked above i are $N_c \cup \{j_1, j_2\}$, one more than the number of seats. School c rejects the applicant with the lowest priority, and by hypothesis some member of N_c ranks below both j_1 and j_2 at c, so the rejected student is a member of N_c, say g, who has no further acceptable school. Writing $N_c' = N_c$ when $\beta = 0$ and $N_c' = N_c \setminus \{g\}$ when $\beta = 1$, a set of $q_c - 2$ students in both cases,

$$\mathrm{DA}(P') : \quad c = N_c' \cup \{j_1, j_2\}, \quad d_1 = N_{d_1} \cup \{k_1\}, \quad d_2 = N_{d_2} \cup \{i\}.$$

Step 4: the efficient frontier at P'. As in Step 2, every fixed student is now either at her unique acceptable school or unassigned, and Lemma 1(i) keeps unassigned whichever of the two is displaced, namely k_2 when $\beta = 0$ and g when $\beta = 1$. School d_1 is full with fixed students, so j_1 cannot move and remains at c. The only students who can move are i and j_2, and the same counting argument gives $\mathcal{D}(P') = \{\mathrm{DA}(P'), \nu'\}$ and $\mathcal{E}(P') = \{\nu'\}$ with

$$\nu_i' = c, \qquad \nu_c' = N_c' \cup \{i, j_1\}.$$

Step 5: conclusion. Student i is assigned to c at both profiles. If $\beta = 0$ then $N_c' = N_c$ and

$$\nu_c = N_c \cup \{i, j_2\} \neq N_c \cup \{i, j_1\} = \nu_c',$$

because $j_1 \neq j_2$. If $\beta = 1$ then $\nu_c = N_c \cup \{i\}$ contains g and not j_1, while $\nu_c' = (N_c \setminus \{g\}) \cup \{i, j_1\}$ contains j_1 and not g, and $g \neq j_1$ because $g \in N_c$. In both cases $\nu_i = \nu_i' = c$ and $\nu_c \neq \nu_c'$, so (P, P') is a rigid pair for i. By Steps 2 and 4, $\mathcal{D}(P)$ and $\mathcal{D}(P')$ have two elements each, so the pair is tight, and Corollary 2 applies. □

The two problems that close the Appendix show that the direct route of Definition 8 is not an artefact of the construction. Both satisfy the definition with one direct route, and the first of them has no twin cycle, so there the impossibility is outside the reach of the guarded case, which is the configuration that Theorem 1 already covers. The second problem also displays three features that the proof of Theorem 3 never produces.

Let Γ^4 have $I = \{i_1, \ldots, i_5\}$, $C = \{c_1, c_2, c_3\}$, capacities $q_{c_1} = 2$ and $q_{c_2} = q_{c_3} = 1$, and the priorities and preferences of Table 2. Total capacity is one seat short of the number of students, so exactly one student is unassigned at every individually rational and non-wasteful matching. The two profiles differ only in the report of i_5, and only in the school she ranks second.

Priorities		**Preferences**	
$\succ_{c_1}$:	i_1, i_2, i_3, i_4, i_5	P_{i_1}:	c_2, c_1
$\succ_{c_2}$:	i_3, i_2, i_4, i_5, i_1	P_{i_2}:	c_3, c_1
$\succ_{c_3}$:	i_5, i_3, i_2, i_1, i_4	P_{i_3}:	c_3
		P_{i_4}:	c_1
		P_{i_5}:	c_1, c_2
		P'_{i_5}:	c_1, c_3

Table 2: The problem Γ^4 and the two preference profiles.

Lemma 9 below shows that Γ^4 has no twin cycle. It does have a separating double rotation: take

$$c = c_1, \quad i = i_5, \quad j_1 = i_2, \quad j_2 = i_1, \quad d_1 = c_3, \quad d_2 = c_2, \quad k_1 = i_3, \quad N_c = \{i_4\},$$

with $N_{c_2} = N_{c_3} = \emptyset$. Route 1 is guarded, since $i_5 \succ_{c_3} i_3 \succ_{c_3} i_2$, and route 2 is direct, since $i_5 \succ_{c_2} i_1$ and the two are adjacent at c_2. The classmate i_4 ranks above i_5 and below both i_1 and i_2 at c_1, as the direct route requires. The reports of Table 2 are exactly the ones the proof of Theorem 3 prescribes for this configuration, with the two profiles interchanged, so that theorem gives the following statement directly.

Proposition 12. *In Γ^4, no DA-improving mechanism is locally non-bossy, and in particular no mechanism is simultaneously efficient, DA-dominating and locally non-bossy. Hence the absence of a twin cycle does not guarantee that the three properties are compatible.*

The two runs are verified directly at the end of this subsection, without appealing to the theorem. At the two profiles the DA outcomes are

$$\mathrm{DA}(P): \quad c_1 = \{i_1, i_2\},\ c_2 = \{i_5\},\ c_3 = \{i_3\},\ \text{with } i_4 \text{ unassigned},$$

$$\mathrm{DA}(P'): \quad c_1 = \{i_2, i_4\},\ c_2 = \{i_1\},\ c_3 = \{i_5\},\ \text{with } i_3 \text{ unassigned},$$

and the unique efficient improvements are

$$\nu: \quad c_1 = \{i_2, i_5\},\ c_2 = \{i_1\},\ c_3 = \{i_3\}, \qquad \nu': \quad c_1 = \{i_4, i_5\},\ c_2 = \{i_1\},\ c_3 = \{i_2\}.$$

Student i_5 obtains c_1 under both reports and her classmate changes from i_2 to i_4. Since Proposition 4 and Proposition 12 cannot both hold with an acyclic priority profile, the priorities of Γ^4 contain an Ergin cycle. Consistently with Proposition 5, the manipulator is denied c_1 by DA at both profiles and her two DA assignments, c_2 and c_3, are the origins of the two rotations that Corollary 3 delivers.

Both forced improvements in Γ^4 are trading cycles of length two between i_5 and one incumbent of c_1: at P she exchanges c_2 for the seat that i_1 vacates, and at P' she exchanges c_3 for the seat that i_2 vacates. What a twin cycle would require and Γ^4 lacks is an intermediate student on the second route: no student lies between i_5 and i_1 at c_2, so nobody is displaced there whose departure empties the seat that i_1 moves into. Student i_5 empties it herself. The price is paid at c_1, where the arrival of i_1 pushes out i_4, and it is the presence or absence of i_4 that makes the two cohorts of i_5 differ. This is the direct route of Definition 8, and Γ^4 shows that it cannot be dispensed with.

Lemma 9. *The problem Γ^4 has no twin cycle.*

Proof. Only c_1 has capacity at least two, so a twin cycle must have $c = c_1$, and then $|N_{c_1}| = |N_{c_2}| = |N_{c_3}| = 0$, so the scarcity requirements are vacuous. For $h \in I$ let

$$B(h) = \big\{a \in \{c_2, c_3\} :\ \exists\, j, k \in I \setminus \{h\} \text{ distinct with } j \succ_{c_1} h \text{ and } h \succ_a k \succ_a j\big\},$$

and for $a \in B(h)$ let $J_a(h)$ collect the students j that qualify. A twin cycle requires two distinct schools $d_1, d_2 \in B(h)$ for some h, together with $j_1 \in J_{d_1}(h)$ and $j_2 \in J_{d_2}(h)$ that are distinct, and witnesses k_1, k_2 such that the five students are distinct.

Student i_1 has the highest priority at c_1, so no j satisfies $j \succ_{c_1} i_1$ and $B(i_1) = \emptyset$. For i_2 the only candidate is $j = i_1$; at c_2 students i_4 and i_5 lie between i_2 and i_1, so $c_2 \in B(i_2)$, while at c_3 students i_2 and i_1 are adjacent, so $c_3 \notin B(i_2)$ and $B(i_2) = \{c_2\}$. For i_4 the candidates are $j \in \{i_1, i_2, i_3\}$; at c_3 student i_4 has the lowest priority, so $c_3 \notin B(i_4)$, while at c_2 the pair $j = i_1$, $k = i_5$ works, so $B(i_4) = \{c_2\}$. For i_5, at c_2 the only student ranked below her is i_1 and the two are adjacent, so $c_2 \notin B(i_5)$, while at c_3 the pair $j = i_2$, $k = i_3$ works, so $B(i_5) = \{c_3\}$. For these four students $|B(h)| \leq 1$, so none of them can serve as the shared student of a twin cycle.

Only i_3 has $B(i_3) = \{c_2, c_3\}$, and there the two qualifying sets coincide. At c_2 student i_3 has the highest priority and the students above her at c_1 are i_1 and i_2; since i_2 is immediately below i_3 at c_2, no witness separates them, so $J_{c_2}(i_3) = \{i_1\}$. At c_3 student i_3 is second and i_2 is immediately below her, so again $J_{c_3}(i_3) = \{i_1\}$, with the unique witness $k = i_2$. Any twin cycle with shared student i_3 would therefore need $j_1 = j_2 = i_1$, which Definition 9 forbids. Hence Γ^4 has no twin cycle. □

The two runs in Γ^4.

The two DA outcomes. At P, students $i_1, \ldots, i_5$ apply to c_2, c_3, c_3, c_1 and c_1 respectively. School c_3 keeps i_3 and rejects i_2, and c_1 keeps i_4 and i_5. Student i_2 applies to c_1, which now keeps i_2 and i_4 and rejects i_5. Student i_5 applies to c_2 and displaces i_1, since $i_5 \succ_{c_2} i_1$. Student i_1 applies to c_1, which keeps i_1 and i_2 and rejects i_4, who has no further acceptable school. Hence

$$\mathrm{DA}(P): \quad c_1 = \{i_1, i_2\},\ c_2 = \{i_5\},\ c_3 = \{i_3\},\ \text{with } i_4 \text{ unassigned.}$$

At P' the algorithm is identical until i_5 is rejected by c_1. She then applies to c_3 and displaces i_3, since $i_5 \succ_{c_3} i_3$, and i_3 has no further acceptable school. Hence

$$\mathrm{DA}(P'): \quad c_1 = \{i_2, i_4\},\ c_2 = \{i_1\},\ c_3 = \{i_5\},\ \text{with } i_3 \text{ unassigned.}$$

The unique efficient improvement at P. By Lemma 1(i), every element of $\mathcal{D}(P)$ leaves i_4 unassigned, assigns two students to c_1 and one student to each of c_2 and c_3. Student i_3 is at her only acceptable school and is fixed at c_3; consequently i_2, whose only improvement over c_1 is c_3, is fixed at c_1. The remaining students i_1 and i_5 can hold c_1 or c_2, and exactly one seat of c_1 and the seat of c_2 are available to them. Hence $\mathcal{D}(P) = \{\mathrm{DA}(P), \nu\}$ with

$$\nu: \quad c_1 = \{i_2, i_5\},\ c_2 = \{i_1\},\ c_3 = \{i_3\},\ \text{with } i_4 \text{ unassigned.}$$

Both i_1 and i_5 obtain their first choice under ν, so ν Pareto dominates $\mathrm{DA}(P)$ and $\mathcal{E}(P) = \{\nu\}$.

The unique efficient improvement at P'. By Lemma 1(i), student i_3 is unassigned in every element of $\mathcal{D}(P')$. Student i_1 is assigned her first choice c_2 and student i_4 her only acceptable school c_1, so both are fixed. The remaining students i_2 and i_5 can hold the free seat of c_1 or the seat of c_3, so $\mathcal{D}(P') = \{\mathrm{DA}(P'), \nu'\}$ with

$$\nu': \quad c_1 = \{i_4, i_5\},\ c_2 = \{i_1\},\ c_3 = \{i_2\},\ \text{with } i_3 \text{ unassigned,}$$

and since ν' gives i_2 and i_5 their first choices, $\mathcal{E}(P') = \{\nu'\}$.

Conclusion. We have $\nu_{i_5} = c_1 = \nu'_{i_5}$ while $\nu_{c_1} = \{i_2, i_5\} \neq \{i_4, i_5\} = \nu'_{c_1}$, so (P, P') is a rigid pair for i_5, as Theorem 3 predicts, and it is tight because $\mathcal{D}(P)$ and $\mathcal{D}(P')$ have two elements each. Corollary 2 applies.

Trading cycles of length two are not the general case. In the next problem the impossibility also holds, and the profiles we record display three further phenomena that an exact condition would have to accommodate.

Let Γ^5 have six students $I = \{i_1, \ldots, i_6\}$, four schools $C = \{c_1, c_2, c_3, c_4\}$, capacities $q_{c_1} = 2$ and $q_{c_2} = q_{c_3} = q_{c_4} = 1$, and the priorities and preferences of Table 3. All schools are acceptable to all students, and total capacity is one seat short of the number of students. The two profiles differ only in the report of i_1, and only in the order of the two schools she ranks last.

Priorities		**Preferences**	
$\succ_{c_1}$:	$i_2, i_6, i_1, i_4, i_3, i_5$	P_{i_1}:	c_1, c_4, c_3, c_2
$\succ_{c_2}$:	$i_1, i_3, i_6, i_2, i_5, i_4$	P'_{i_1}:	c_1, c_4, c_2, c_3
$\succ_{c_3}$:	$i_2, i_6, i_1, i_4, i_3, i_5$	P_{i_2}:	c_4, c_1, c_3, c_2
$\succ_{c_4}$:	$i_4, i_6, i_5, i_2, i_1, i_3$	P_{i_3}:	c_2, c_4, c_1, c_3
		P_{i_4}:	c_3, c_1, c_2, c_4
		P_{i_5}:	c_1, c_2, c_3, c_4
		P_{i_6}:	c_2, c_1, c_3, c_4

Table 3: The problem Γ^5 and the two preference profiles. Note that $\succ_{c_3}$ coincides with $\succ_{c_1}$.

Proposition 13. *In Γ^5, no DA-improving mechanism is locally non-bossy, and in particular no mechanism is simultaneously efficient, DA-dominating and locally non-bossy.*

The problem also has a separating double rotation, with

$$c = c_1, \quad i = i_5, \quad j_1 = i_1, \quad j_2 = i_4, \quad d_1 = c_4, \quad d_2 = c_2, \quad k_1 = i_2, \quad N_c = \{i_3\},$$

empty scarcity sets at c_4 and c_2, route 1 guarded because $i_5 \succ_{c_4} i_2 \succ_{c_4} i_1$, and route 2 direct because i_5 and i_4 are adjacent at c_2; the classmate i_3 ranks above i_5 and below both i_1 and i_4 at c_1. Theorem 3 therefore delivers the proposition, with i_5 as the manipulator. The profiles of Table 3 witness its second statement through a different student, and we record them because the improvements they force have a shape that the proof of Theorem 3 never produces; they are verified at the end of this subsection. The DA outcomes are

$$\mathrm{DA}(P): \quad c_1 = \{i_2, i_6\},\ c_2 = \{i_3\},\ c_3 = \{i_1\},\ c_4 = \{i_4\},\ \text{with } i_5 \text{ unassigned},$$

$$\mathrm{DA}(P'): \quad c_1 = \{i_2, i_6\},\ c_2 = \{i_1\},\ c_3 = \{i_4\},\ c_4 = \{i_5\},\ \text{with } i_3 \text{ unassigned},$$

and the unique efficient improvements are

$$\nu:\ c_1 = \{i_1, i_6\},\ c_2 = \{i_3\},\ c_3 = \{i_4\},\ c_4 = \{i_2\},$$

$$\nu':\ c_1 = \{i_1, i_5\},\ c_2 = \{i_6\},\ c_3 = \{i_4\},\ c_4 = \{i_2\}.$$

Three features of Γ^5 are absent from Γ^1 and Γ^4, and Lemma 1(iv) makes them precise. First, at P the improvement forced on every efficient DA-dominating mechanism is a single trading cycle of length three, in which i_1 moves from c_3 to c_1, student i_2 from c_1 to c_4 and student i_4 from c_4 to c_3. The seat that i_1 vacates is not the seat that the departing classmate occupies, so the exchange does not close within two schools, and the rotation that Corollary 3 delivers has length three. Second, at P' the improvement consists of two disjoint cycles of length two that run simultaneously through the two seats of c_1, namely i_1 and i_6 exchanging c_2 and c_1, and i_2 and i_5 exchanging c_1 and c_4. Third, the classmate who arrives at P' is i_5, who ranks *below* the manipulator at c_1: efficiency can bring in a lower-priority student once enough higher-priority students leave. None of the three appears in the construction behind Theorem 3, which only ever produces exchanges between two schools, while a double rotation accommodates all of them. This is one reason not to expect Definition 8 to be necessary as it stands: the trades it is built from are the shortest ones, and Γ^5 shows that the impossibility can travel along longer ones in a problem that satisfies the definition for other reasons.

A fourth feature is worth recording because Proposition 2(i) does not exclude it. In Γ^5 the report of i_1 changes which student DA leaves unassigned, from i_5 to i_3, and thereby changes the pool of students available for the seats of c_1. Proposition 2(i) rules out this channel only when the manipulator is herself unassigned under both reports.

The two runs in Γ^5.

The two DA outcomes. In the first round c_1 receives i_1 and i_5 and keeps both, c_2 receives i_3 and i_6 and rejects i_6, c_3 holds i_4, and c_4 holds i_2. Student i_6 then applies to c_1, which keeps i_6 and i_1 and rejects i_5. Student i_5 is rejected in turn by c_2, where i_3 ranks above her, and by c_3, where i_4 ranks above her, and is accepted at c_4, displacing i_2. Student i_2 applies to c_1, which keeps i_2 and i_6 and rejects i_1. Up to this point the two runs coincide, because i_1 has applied only to c_1.

Under P, student i_1 applies to c_4 and is rejected because $i_5 \succ_{c_4} i_1$, then applies to c_3 and displaces i_4, since $i_1 \succ_{c_3} i_4$. Student i_4 is rejected by c_1 and by c_2, where she has the lowest priority, and is accepted at c_4, displacing i_5, who has no acceptable school left. Under P', student i_1 applies to c_4, is rejected, and then applies to c_2, where she has the

highest priority, displacing i_3. Student i_3 is rejected by c_4, by c_1 and by c_3, and remains unassigned. Therefore

$$\mathrm{DA}(P): \quad c_1 = \{i_2, i_6\},\ c_2 = \{i_3\},\ c_3 = \{i_1\},\ c_4 = \{i_4\},\ \text{with } i_5 \text{ unassigned},$$

$$\mathrm{DA}(P'): \quad c_1 = \{i_2, i_6\},\ c_2 = \{i_1\},\ c_3 = \{i_4\},\ c_4 = \{i_5\},\ \text{with } i_3 \text{ unassigned}.$$

The unique efficient improvement at P. By Lemma 1(i), every element of $\mathcal{D}(P)$ assigns two students to c_1, one to each other school, and leaves i_5 unassigned. Student i_3 is at her top choice c_2 and is fixed there; consequently i_6, whose only improvement over c_1 is c_2, is fixed at c_1. The remaining students i_1, i_2 and i_4 occupy the free seat at c_1, the seat at c_3 and the seat at c_4. Their upper contour sets restrict the possibilities: student i_1 weakly improves upon c_3 only at c_1, c_4 or c_3, student i_2 weakly improves upon c_1 only at c_4 or c_1, and student i_4 weakly improves upon c_4 only at c_3, c_1, c_2 or c_4, and c_2 is unavailable. Enumerating the bijections between $\{i_1, i_2, i_4\}$ and the three seats that respect these restrictions yields exactly four matchings, namely $\mathrm{DA}(P)$ and

$$\nu:\ c_1 = \{i_1, i_6\},\ c_3 = \{i_4\},\ c_4 = \{i_2\}; \qquad \mu^1:\ c_1 = \{i_2, i_6\},\ c_3 = \{i_4\},\ c_4 = \{i_1\};$$

$$\mu^2:\ c_1 = \{i_4, i_6\},\ c_3 = \{i_1\},\ c_4 = \{i_2\},$$

each completed with $c_2 = \{i_3\}$ and i_5 unassigned. Comparing, ν Pareto dominates μ^1, because i_1 prefers c_1 to c_4 and i_2 prefers c_4 to c_1 while i_4 is unaffected; ν Pareto dominates μ^2, because i_1 prefers c_1 to c_3 and i_4 prefers c_3 to c_1 while i_2 is unaffected; and ν Pareto dominates $\mathrm{DA}(P)$. Hence $\mathcal{E}(P) = \{\nu\}$.

The unique efficient improvement at P'. Consider

$$\nu': \quad c_1 = \{i_1, i_5\},\ c_2 = \{i_6\},\ c_3 = \{i_4\},\ c_4 = \{i_2\},\ \text{with } i_3 \text{ unassigned}.$$

Each of i_1, i_2, i_4, i_5 and i_6 is assigned to her top choice, and by Lemma 1(i) student i_3 is unassigned in every element of $\mathcal{D}(P')$. Comparing ν' with $\mathrm{DA}(P')$ shows that ν' weakly Pareto dominates it, so $\nu' \in \mathcal{D}(P')$; and since no matching can improve upon a matching that gives every assigned student her top choice while leaving the same student unassigned, ν' Pareto dominates every other element of $\mathcal{D}(P')$ and $\mathcal{E}(P') = \{\nu'\}$.

Conclusion. We have $\nu_{i_1} = c_1 = \nu'_{i_1}$ while $\nu_{c_1} = \{i_1, i_6\} \neq \{i_1, i_5\} = \nu'_{c_1}$, so (P, P') is a rigid pair for i_1 and Corollary 1 applies.

We now prove the results on consent and on types. In the two probabilistic proofs, positions are counted from the bottom of an order, so that the student in position 1 at a

school has the lowest priority there.

Proof of Proposition 9. Let i be the student in position 1 at c and j_1 the student in position 1 at d_1. Let B_1 be the set of students in positions $1, \dots, m+1$ at d_1, let $S_1 = B_1 \setminus \{j_1\}$, and let k_1 be the member of $S_1 \cap W$ with the lowest priority at d_1 if this set is non-empty, and $k_1 = j_1$ otherwise. At d_2, let R be the set of students in positions 1, 2 and 3, let S_2 be the set of students in positions $4, \dots, m+3$, and let $B_2 = R \cup S_2$. Consider the events

$$E_1 : \ i \notin B_1, \qquad E_2 : \ S_1 \cap W \neq \emptyset, \qquad E_3 : \ i \notin B_2,$$
$$E_4 : \ j_1 \notin S_2 \text{ and } k_1 \notin S_2, \qquad E_5 : \ S_2 \cap W \neq \emptyset.$$

Suppose that all five occur. Let j_2 be the member of $R \setminus \{j_1, k_1\}$ with the lowest priority at d_2, which exists because R has three members, and let k_2 be any member of $S_2 \cap W$. The students i, j_1, j_2, k_1, k_2 are distinct. Indeed, $k_1 \in S_1$ differs from j_1, and i differs from both by E_1; j_2 differs from j_1 and k_1 by construction and from i by E_3; and $k_2 \in S_2$ differs from i by E_3, from j_1 and k_1 by E_4, and from $j_2 \in R$ because R and S_2 are disjoint. Since i has the lowest priority at c, we have $j_1 \succ_c i$ and $j_2 \succ_c i$. At d_1, student j_1 is in position 1, student k_1 is in a position between 2 and $m+1$, and i is above position $m+1$ by E_1, so $i \succ_{d_1} k_1 \succ_{d_1} j_1$. At d_2, student j_2 is in one of the positions 1 to 3, student k_2 is in one of the positions 4 to $m+3$, and i is above position $m+3$ by E_3, so $i \succ_{d_2} k_2 \succ_{d_2} j_2$. Both routes are therefore guarded, and both guards belong to W.

For the scarcity sets, let $Q = I \setminus (B_1 \cup B_2 \cup \{i\})$. Every member of Q ranks above every member of B_1 at d_1, and in particular above k_1, and above every member of B_2 at d_2, and in particular above k_2. The set Q contains none of i, j_1, j_2, k_1, k_2 and has at least $n - (m+1) - (m+3) - 1 = n - 2m - 5$ members, which is at least $q_{d_1} + q_{d_2} - 2$ by the hypothesis on m. Choose disjoint sets $N_{d_1}, N_{d_2} \subseteq Q$ with $q_{d_1} - 1$ and $q_{d_2} - 1$ members. The students outside $\{i, j_1, j_2, k_1, k_2\} \cup N_{d_1} \cup N_{d_2}$ number $n - 5 - (q_{d_1} + q_{d_2} - 2)$, which is at least $q_c - 2$ because $q_c + q_{d_1} + q_{d_2} \leq n - 1$; let N_c consist of $q_c - 2$ of them. Every member of N_c ranks above i at c. Definition 8 holds with both routes guarded, so Γ has a twin cycle at c whose guards belong to W.

It remains to bound the probability that some event fails. The student i is determined by $\succ_c$, which is independent of $\succ_{d_1}$ and $\succ_{d_2}$, so her position at each of these schools is uniform on $\{1, \dots, n\}$, and E_1 and E_3 fail with probabilities $(m+1)/n$ and $(m+3)/n$. The set S_1 is a uniformly distributed set of m students, so

$$\Pr(S_1 \cap W = \emptyset) = \binom{n-w}{m} \Big/ \binom{n}{m} = \prod_{t=0}^{m-1} \frac{n-w-t}{n-t} \leq \left(1 - \frac{w}{n}\right)^m.$$

Conditional on $\succ_c$ and $\succ_{d_1}$, which determine j_1 and k_1, the set S_2 is again a uniformly distributed set of m students. A given student belongs to it with probability m/n, so E_4 fails with probability at most $2m/n$, and E_5 fails with probability at most $(1 - w/n)^m$, as before. Adding the five bounds gives the stated probability. In the event that all five occur, Proposition 7, applied to the twin cycle just constructed with a consent set that contains its two guards, shows that EADA^W is locally bossy, and Proposition 8 gives the same conclusion for every mechanism that uses the consent set W and is efficient whenever possible. □

Proof of Lemma 7. Write $U = U_k$. Let Γ^U be the problem in which only the students of U are present, with the same schools and capacities and with each priority order restricted to U, and let μ^U be the outcome of DA in Γ^U at the restriction of P to U. We show that $\text{DA}_h(P) = \mu^U_h$ for every $h \in U$. Since μ^U depends on P only through the reports of the students of U, which are the same at P', the lemma follows.

Let $\mu = \text{DA}(P)$. First, the restriction of μ to U is stable in Γ^U. It is individually rational because μ is. Let $h \in U$ and let x be a school with $x \, P_h \, \mu_h$. Since (h, x) does not block μ, school x is full at μ and every member of μ_x has higher priority than h at x, hence belongs to U by stratification. In Γ^U, school x therefore holds q_x students of U, all ranked above h, and (h, x) does not block there either. Since μ^U is the student-optimal stable matching of Γ^U (Gale & Shapley, 1962), $\mu^U_h \, R_h \, \mu_h$ for every $h \in U$.

Second, extend μ^U to a matching λ of Γ by assigning the students of $I \setminus U$ according to the outcome of DA in the problem with students $I \setminus U$, priorities restricted to them, and capacities $q_x - |\mu^U_x|$. We claim that λ is stable at P. It is individually rational because both parts are. Let $x \, P_h \, \lambda_h$. If $h \in U$, stability of μ^U in Γ^U gives that x holds q_x students of U under μ^U, all ranked above h at x, so $\lambda_x = \mu^U_x$ and (h, x) does not block λ. If $h \notin U$, stability in the residual problem gives that x has no residual seat left and that every student of $I \setminus U$ assigned to x ranks above h there, while every student of U ranks above h at x by stratification. School x is then full at λ with students who all outrank h, and again (h, x) does not block. Hence λ is stable at P, and student optimality of μ gives $\mu_h \, R_h \, \lambda_h = \mu^U_h$ for every $h \in U$.

The two comparisons and strict preferences give $\mu_h = \mu^U_h$ for every $h \in U$. □

Proof of Proposition 11. Let i be the student in position 1 at c, let j_1 and k_1 be the students in positions 1 and 2 at d_1, let B be the set of students in positions $1, \ldots, m + 1$

at d_2, and let k_2 be the student in position $m+2$ at d_2. Consider the events

$$F_1:\ i \notin \{j_1, k_1\}, \qquad F_2:\ i \notin B \cup \{k_2\}, \qquad F_3:\ k_2 \notin \{j_1, k_1\},$$
$$F_4:\ \tau(h) \neq \tau(j_1) \text{ for some } h \in B \setminus \{j_1, k_1\}.$$

Suppose that all four occur, and let $j_2 \in B \setminus \{j_1, k_1\}$ have $\tau(j_2) \neq \tau(j_1)$. The students i, j_1, j_2, k_1, k_2 are distinct: i differs from j_1 and k_1 by F_1 and from j_2 and k_2 by F_2, student j_2 differs from j_1 and k_1 by construction and from k_2 because $k_2 \notin B$, and k_2 differs from j_1 and k_1 by F_3. Since i has the lowest priority at c, we have $j_1 \succ_c i$ and $j_2 \succ_c i$. At d_1, $i \succ_{d_1} k_1 \succ_{d_1} j_1$ by F_1. At d_2, student j_2 is in one of the positions 1 to $m+1$, student k_2 is in position $m+2$, and i is above it by F_2, so $i \succ_{d_2} k_2 \succ_{d_2} j_2$. Both routes are guarded.

For the scarcity sets, let $Q = I \setminus (\{i, j_1, k_1, k_2\} \cup B)$. It contains none of the five students above, it has at least $n-m-5$ members, and all of them rank above k_2 at d_2. Since $q_{d_2} \leq q_{d_1}$ and $q_{d_1} + q_{d_2} \leq n-1-q_c \leq n-3$, we have $q_{d_2} - 1 \leq (n-5)/2 \leq n-m-5$, where the last inequality uses $m \leq (n-5)/2$. Choose $N_{d_2} \subseteq Q$ with $q_{d_2} - 1$ members. Every student other than j_1 and k_1 ranks above k_1 at d_1, and every student other than i ranks above i at c. The students outside $\{i, j_1, j_2, k_1, k_2\} \cup N_{d_2}$ number $n - 5 - (q_{d_2} - 1)$, which is at least $(q_{d_1} - 1) + (q_c - 2)$ because $q_c + q_{d_1} + q_{d_2} \leq n-1$, so disjoint sets N_{d_1} and N_c with $q_{d_1} - 1$ and $q_c - 2$ members can be chosen among them. Definition 8 holds with both routes guarded. By Theorem 3, the profiles (P, P') constructed in its proof form a tight rigid pair for i at which every admissible correction assigns i to c, and by Corollary 2 every DA-improving mechanism selects ν at P and ν' at P'. By Corollary 8, ν_c and ν'_c differ only in j_2 and j_1, which have different types, so the two cohorts contain different numbers of students of type $\tau(j_1)$.

For the probability, the position of i at d_1 and at d_2 is uniform on $\{1, \ldots, n\}$, because i is determined by $\succ_c$, so F_1 and F_2 fail with probabilities $2/n$ and $(m+2)/n$. Conditional on $\succ_c$ and $\succ_{d_1}$, which determine i, j_1 and k_1, the student in position $m+2$ at d_2 is uniform on I, so F_3 fails with probability $2/n$, and B is a uniformly distributed set of $m+1$ students. Let $a = \tau(j_1)$. If F_4 fails, every member of B other than j_1 and k_1 has type a, and since j_1 itself has type a, the set B is contained in $T_a \cup \{k_1\}$, which has at most $\lfloor (1-\varepsilon)n \rfloor + 1 \leq n$ members. Hence

$$\Pr(F_4 \text{ fails}) \leq \binom{\lfloor (1-\varepsilon)n \rfloor + 1}{m+1} \Big/ \binom{n}{m+1} \leq \left(1 - \varepsilon + \frac{1}{n}\right)^{m+1}.$$

Adding the four bounds gives the stated probability. □