

Ortega, Josué; Arribillaga, R. Pablo

Working Paper
Justifiable Priority Violations

QBS Working Paper, No. 2026/04

Provided in Cooperation with:
Queen's University Belfast, Queen's Business School

Suggested Citation: Ortega, Josué; Arribillaga, R. Pablo (2026) : Justifiable Priority Violations, QBS Working Paper, No. 2026/04, Queen's University Belfast, Queen's Business School, Belfast

This Version is available at:
<https://hdl.handle.net/10419/339987>

Standard-Nutzungsbedingungen:

Die Dokumente auf EconStor dürfen zu eigenen wissenschaftlichen Zwecken und zum Privatgebrauch gespeichert und kopiert werden.

Sie dürfen die Dokumente nicht für öffentliche oder kommerzielle Zwecke vervielfältigen, öffentlich ausstellen, öffentlich zugänglich machen, vertreiben oder anderweitig nutzen.

Sofern die Verfasser die Dokumente unter Open-Content-Lizenzen (insbesondere CC-Lizenzen) zur Verfügung gestellt haben sollten, gelten abweichend von diesen Nutzungsbedingungen die in der dort genannten Lizenz gewährten Nutzungsrechte.

Terms of use:

Documents in EconStor may be saved and copied for your personal and scholarly purposes.

You are not to copy documents for public or commercial purposes, to exhibit the documents publicly, to make them publicly available on the internet, or to distribute or otherwise use the documents in public.

If the documents have been made available under an Open Content Licence (especially Creative Commons Licences), you may exercise further usage rights as specified in the indicated licence.



**QUEEN'S
UNIVERSITY
BELFAST**

**QUEEN'S
BUSINESS
SCHOOL**

Working Paper Series - QBS Working Paper 2026/04

Justifiable Priority Violations

Josué Ortega

Queen's University Belfast

R. Pablo Arribillaga

Universidad Nacional de San Luis

08 April 2026

Series edited by Philip T. Fliers.

To submit forward your paper to qbs.rps@qub.ac.uk.

Justifiable Priority Violations

JOSUÉ ORTEGA, Queen's University Belfast, UK

R. PABLO ARRIBILLAGA, Instituto de Matemática Aplicada San Luis, Universidad Nacional de San Luis, and CONICET, Argentina

Addressing the large inefficiencies generated by the Deferred Acceptance (DA) mechanism requires priority violations, but which ones are justifiable? The leading approach is to ask individuals if they consent to waive their priority ex-ante. We develop an alternative question-free solution, in which a priority violation is justifiable whenever the affected student either (i) directly benefits from the improvement, or (ii) is unimprovable under any assignment that Pareto-dominates DA. This endogenous justifiability criterion permits improvements unattainable by the leading consent-based mechanism under any consent structure. We provide a “*just below cutoffs*” mechanism that always finds a strongly justifiable matching whenever DA's outcome is inefficient, and build on it to construct a polynomial-time algorithm that expands justifiable improvements iteratively, converging to a DA improvement that cannot be Pareto-improved by any justifiable matching without strictly expanding the beneficiary set. Finally, we prove theoretically that both the ex-ante consent and the endogenous justifiability frameworks have important limitations in reaching Pareto-efficient outcomes, and use simulations to quantify how binding these constraints are in practice.

Keywords: school choice, consent-based mechanisms, justifiable priority violations.

1 Two Paths Beyond Deferred Acceptance

A school choice mechanism must balance two often conflicting objectives: respecting legally defined priorities and allocating scarce school places efficiently. The student-proposing Deferred Acceptance algorithm (DA) has become the dominant benchmark because it produces the student-optimal stable matching: an allocation that respects priorities that is not Pareto-dominated by any other stable matching. Nonetheless, DA can generate sizeable welfare losses, so it is natural for an Education Authority to consider alternative mechanisms that Pareto-improve upon the DA outcome.¹ Any such improvement necessarily entails priority violations, but which violations are justifiable? Our contribution is to propose a new concept of controlled priority violations called *justifiable* that allows us to improve upon DA whenever its outcome is Pareto-inefficient.

In our framework, students whose priorities are violated fall into three categories: (1) beneficiaries, who improve their DA placement; (2) unimprovable students, whom no DA-improvement could help; and (3) improvable non-beneficiaries, who could have benefited under some DA improvement but who do not under the chosen one. In our view, the first two have no grounds for complaint: they either gained or could never have gained. The third category is different, as such a student could object: *why was my priority overridden if I did not gain, and I could have?*

This potential objection motivates our justifiability criterion: a DA improvement is *justifiable* if it violates priorities only for beneficiaries and unimprovable students, but never for improvable non-beneficiaries. The criterion is endogenous: whether a violation is justifiable depends on what the improvement accomplishes for the affected student relative to what was feasible, not on permissions granted in advance.

We show that justifiable improvements always exist when DA is inefficient (Theorem 1) and characterize a benchmark obtained by the *Just Below Cutoffs* mechanism (JBC). JBC assigns to each school that rejected an improvable student the highest-priority among those students below the cutoff (the lowest-priority student assigned to each school in DA), so no improvable student's priority is ever

¹The Education Authority of Flanders is one such example; see [Cerrone et al. \[2024\]](#).

violated by its trades, not even those of students who happen to benefit. We call this type of justifiable improvement *strongly justifiable*. Any matching satisfying this stronger notion corresponds to executing disjoint improvement cycles. These strongly justifiable matchings are essentially unique: if a student improves in two strongly justifiable outcomes, she gets assigned to the same school in both; furthermore, JBC selects the largest collection of such cycles.

Despite its normative appeal, justifiability is generally incompatible with full Pareto-efficiency (Proposition 1): achieving full efficiency typically requires violating improvable non-beneficiaries' priorities, precisely the violations our criterion excludes. To achieve a constrained efficiency property while preserving justifiability, we construct a polynomial-time algorithm, SJBC+, that builds on JBC via a sequential expansion.

Starting from the JBC matching, each iteration searches for augmenting paths in a bipartite graph on improvable students: these paths reroute existing trading cycles to admit additional students, with priority violations permitted only against current beneficiaries. Each successful rerouting enlarges the beneficiary set, whose members then justify further violations in the next round. In a sense, the algorithm builds "consent for priority violations" endogenously: starting from unimprovable students and iterating until no augmenting path can expand the beneficiary set further. The resulting matching is justifiable and cannot be Pareto-improved by any justifiable matching without strictly expanding the set of beneficiaries (Theorem 2).

Justifiability differs fundamentally from the dominant alternative: school choice with ex-ante consent. In that approach, the mechanism designer asks students in advance if their priority could be overridden in case doing so benefits others at no cost for themselves. The set of consenting students is fed into a consent-based mechanism, such as Efficiency-Adjusted DA (EADA), which always returns a DA improvement that respects the priorities of non-consenting students. We show that justifiability allows for improvements unattainable by EADA under any consent set, just as EADA's outcomes need not be justifiable (Theorem 3; our simulations in Section 6 show that this occurs very frequently). While a number of other weaker stability notions also allow for controlled priority violations, such as legality and

priority-neutrality, they converge to EADA with full consent as the student-optimal matching within these restrictions; thus, justifiability departs from these normative frameworks too.

More importantly, despite EADA's full efficiency guarantee when all students consent, we show that EADA and any other consent-based mechanism also face important limitations when only a fraction of students consent. In particular, any consent-based mechanism that incentivizes consent (i.e. does not give a worse placement to any student who consents) and is constrained efficient (i.e. its outcome is not Pareto-dominated by another matching which respects the priorities of non-consenting students), fails to be *efficient-whenever-possible*: there are efficient matchings that respect the priorities of non-consenters, yet the mechanism returns an inefficient allocation (Theorem 4). Taken together, our results show the complexity of guaranteeing full efficiency when improving upon DA on either the justifiability or the consent-based frameworks.

To understand how frequently our theoretical results bite, we use simulations to evaluate SJBC+ against EADA with full and observed consent. SJBC+ generates more beneficiaries than EADA even when all students consent, and achieves full Pareto-efficiency in over 60% of the cases, highlighting that even though SJBC+'s outcomes can be dominated by another justifiable matching, more often than not its outcomes are not dominated by any matching, justifiable or not.

In summary, our paper proposes a new normative notion that allows for defensible priority violations for Education Authorities that would like to obtain a matching that respects most priorities but that allows for controlled priority violations in order to reach more efficient outcomes. Our results show the limits of this framework with regard to how much efficiency it can recover in theory and in practice. Importantly, the concept of justifiability allows for improvements that cannot be obtained by existing approaches, and thus provides a new direction for the school choice literature.

Outline. Section 2 discusses the literature. Section 3 introduces the model. Section 4 develops the theoretical results on justifiable DA improvements. Section 5 discusses the relationship between justifiability and school choice with consent.

Section 6 analyzes the performance of our proposed algorithm using simulations. Section 7 concludes.

2 Literature

The student-proposing Deferred Acceptance mechanism is often Pareto-inefficient for students [[Abdulkadiroğlu and Sönmez, 2003](#)], yet any Pareto improvement over it must necessarily violate some student's priority. This tension has motivated a large literature on which priority violations are acceptable when improving upon DA.

A striking feature of this literature is that several conceptually different normative frameworks end up coinciding on the same foundational adjustment, Efficiency-Adjusted Deferred Acceptance [EADA; [Kesten, 2010](#)]. Kesten's original justification is ex-ante consent: a priority violation is acceptable when the affected student agrees in advance to waive her priority. [Reny \[2022\]](#) proposes priority neutrality, requiring (roughly) that a displaced student can only lose priority when the violation is in a precise sense unavoidable for helping others; in that framework, EADA delivers the unique student-optimal priority-efficient outcome. [Ehlers and Morrill \[2020\]](#) propose legality, a set-valued stability concept in which blocking is admissible only if the student can be assigned to the blocking school in some assignment under consideration. They show there is a unique student-optimal legal assignment which coincides with EADA's outcome.² Our notion of justifiability departs from this literature: it yields justifiable matchings that lie outside the range of EADA under any consent structure, and thus need not be priority-neutral nor legal.

Our mechanisms build on the improvement-cycle tradition, but target a different constraint. The Just-Below Cutoffs (JBC) algorithm is related to the Top Priority algorithm [[Dur et al., 2019](#)], which iteratively resolves improvement cycles while allowing violations against any student who consents ex-ante; in contrast, JBC uses a single-step construction that violates only unimprovable students' priorities. JBC

²EADA's properties have been extensively studied; see [[Bando, 2014](#), [Chen and Möller, 2023](#), [Dur et al., 2019](#), [Shirakawa, 2025](#), [Tang and Yu, 2014](#), [Tang and Zhang, 2021](#), [Trojan et al., 2020](#), [Trojan and Morrill, 2020](#)].

is also closely related to stable improvement cycles [Erdil and Ergin, 2008], which address inefficiency driven by arbitrary tie-breaking under weak priorities: starting from a stable outcome under a fixed tie-break, they implement Pareto-improving cycle exchanges that remain stable with respect to the underlying weak priorities. JBC targets a different source of inefficiency: it improves upon DA even under strict priorities, but only by violating the priorities of unimprovable students.

Several other strands are complementary to ours. Kitahara and Okumura [2021] analyze how to improve upon DA under partial priority structures. Other work studies when DA can be improved while preserving strategy-proofness [Kesten and Kurino, 2019], and develops refinements based on consistency requirements [Doğan and Yenmez, 2020]. Recent work also shows that DA produces an inefficient matching with high probability, and that almost all students are improvable [Ortega et al., 2024], and documents limitations that no Pareto improvement over DA can overcome [Ortega et al., 2026]. These theoretical results, together with the rural-hospital property satisfied by every matching that Pareto-dominates DA [Alva and Manjunath, 2019], help to explain EADA’s empirical performance in counterfactuals using real-life data [Ortega and Klein, 2023]. A complementary literature examines how to improve upon DA while minimizing the number of blocking pairs or triplets generated [Afacan et al., 2025, Doğan and Ehlers, 2021, Kwon and Shorrer, 2020]. Finally, our work is also related to Decerf et al. [2024], who study incontestable assignments that can be defended against appeals; however, in their framework every improvement upon DA is incontestable.

3 Model

A many-to-one school choice problem P consists of two finite sets of students I and schools S . Each student i has a strict preference \succ_i over $S \cup \{s_\emptyset\}$, where s_\emptyset denotes being unassigned. Each school s has quota $q_s \in \mathbb{N}$ and a strict priority order \triangleright_s over students (except for s_\emptyset , with $q_{s_\emptyset} = \infty$).³

A matching μ assigns each student to one school, $\mu : I \rightarrow S \cup \{s_\emptyset\}$, such that $|\mu^{-1}(s)| \leq q_s$ for all $s \in S$. A matching μ is non-wasteful if, whenever a school $s \in S$

³All of our results hold for arbitrary quotas, although our examples use unit capacities for clarity.

has an empty seat under μ , every student prefers her assignment to s . Throughout the paper, we restrict attention to non-wasteful matchings.

Let $\text{rk}_i(s) \in \{1, \dots, |S| + 1\}$ be the rank of school s under \succ_i (lower is better), and let $\text{rk}_s(i)$ be the priority rank of student i at school s (lower is better priority). The matching μ weakly Pareto-dominates ν if $\text{rk}_i(\mu_i) \leq \text{rk}_i(\nu_i)$ for all i ; it Pareto-dominates if strict for some i . A matching is Pareto-efficient if it is not Pareto-dominated by any other matching.

We say student j violates student i 's priority under μ if (i) i prefers s to μ_i , (ii) $\mu_j = s$, and (iii) $\text{rk}_s(i) < \text{rk}_s(j)$. A matching is stable if it does not violate the priority of any student.

Deferred Acceptance and its Envy Digraph. A mechanism maps a problem P to a matching. The student-proposing Deferred Acceptance (DA) is a well-known example.⁴ Let $\text{DA}(P)$ denote the matching produced by Deferred Acceptance in problem P , and let $\text{DA}_i(P)$ denote student i 's assignment under $\text{DA}(P)$.

The envy digraph $G^{\text{DA}(P)}$ has vertex set I and a directed edge $i \rightarrow j$ if student i prefers $\text{DA}_j(P)$ to $\text{DA}_i(P)$. A cycle is a sequence (i_1, \dots, i_k) of distinct students with edges $i_\ell \rightarrow i_{\ell+1}$ for $\ell = 1, \dots, k-1$ and $i_k \rightarrow i_1$. A cycle packing Π is a collection of disjoint cycles in $G^{\text{DA}(P)}$. Let $V(\Pi) \subseteq I$ be the set of students covered by cycles in Π . For $j \in I$, let $N^-(j) := \{h \in I : h \rightarrow j \text{ in } G^{\text{DA}(P)}\}$ denote the in-neighborhood of j in the envy digraph, i.e. the students who envy j in DA. It is well-known that every Pareto cycle-packing in $G^{\text{DA}(P)}$ corresponds to a sequence of trades that lead to a matching that Pareto dominates $\text{DA}(P)$.

Unimprovable students. Given a problem P , let $\mathcal{M}(P)$ denote the set of matchings that Pareto-dominate $\text{DA}(P)$. A student $i \in I$ is unimprovable if $\mu_i = \text{DA}_i(P)$ for every $\mu \in \mathcal{M}(P)$. Let $\mathcal{U}(P)$ denote the set of unimprovable students, and let $\mathcal{I}(P) := I \setminus \mathcal{U}(P)$ denote the set of improvable students.

The well-known observation below provides a useful connection between unimprovability and the structure of the DA envy digraph.

⁴We postpone its description to Appendix A.

LEMMA 1 ([DUR ET AL., 2019, ORTEGA ET AL., 2024, TANG AND YU, 2014]). *The following three statements are equivalent:*

- (1) *A student is improvable.*
- (2) *The corresponding node lies on a cycle in the DA envy digraph $G^{\text{DA}(P)}$.*
- (3) *The corresponding node belongs to a non-trivial strongly connected component of $G^{\text{DA}(P)}$.*

Justifiability. We now formalize which priority violations can be justified when implementing efficiency improvements over $\text{DA}(P)$. The idea behind justifiability is that, while any Pareto improvement over $\text{DA}(P)$ must violate some priorities, not all violations are equally defensible. We distinguish between violations affecting students who can never benefit from any improvement (unimprovable students) and those affecting students who benefit from the improvement itself.

For any matching $\mu \in \mathcal{M}(P)$, define the set of students who benefit from μ relative to DA .⁵

DEFINITION 1. *Given $\mu \in \mathcal{M}(P)$, the set of beneficiaries is*

$$\mathcal{B}(\mu) := \{i \in I : \text{rk}_i(\mu_i) < \text{rk}_i(\text{DA}_i(P))\}.$$

Unimprovable students are unaffected by the choice of efficiency improvement: regardless of which Pareto improvement is selected, they remain assigned to their DA school. By contrast, non-beneficiaries are students who could benefit under some Pareto improvement but do not benefit under the particular improvement μ .

We now introduce justifiability, our main conceptual contribution: a restriction on which priority violations can be justified when improving upon DA.

DEFINITION 2. *Given $\mu \in \mathcal{M}(P)$, a priority violation against student i under μ is justifiable if*

$$i \in \mathcal{U}(P) \cup \mathcal{B}(\mu).$$

A priority violation against an improvable non-beneficiary is unjustifiable.

⁵For simplicity, we use the notation $\mathcal{B}(\mu)$ instead of the more cumbersome $\mathcal{B}(\mu, P)$.

The criterion is endogenous: whether a violation is justifiable depends on what the improvement accomplishes for the affected student, not on permissions granted in advance.

DEFINITION 3. *A matching $\mu \in \mathcal{M}(P)$ is justifiable if every priority violation in μ is justifiable.*

Labelled Envy Digraph. The envy digraph $G^{\text{DA}(P)}$ records who envies whose assignment under $\text{DA}(P)$ and tells us which improvements can occur. To analyze the priority violations that such exchanges would generate, we enrich $G^{\text{DA}(P)}$ with labels. To each edge $i \rightarrow j$ we associate the set of improvable students who desire school $\text{DA}_j(P)$ (i.e., point to j in $G^{\text{DA}(P)}$) and whose priority at $\text{DA}_j(P)$ would be violated if i were assigned to $\text{DA}_j(P)$:

$$l(i \rightarrow j) := \{h \in \mathcal{I}(P) \cap N^-(j) : \text{rk}_{\text{DA}_j(P)}(h) < \text{rk}_{\text{DA}_j(P)}(i)\}.$$

For a cycle packing Π , define $l(\Pi) := \bigcup_{(i \rightarrow j) \in \Pi} l(i \rightarrow j)$. We remark that the labels tell us the priorities that an exchange would violate fixing the allocation for every other student to be DA , without verifying the actual cycle implemented: $i \xrightarrow{k} j$ tell us that an improvement cycle where i takes j 's place would *potentially* violate k 's priority, without verifying whether k 's priority is actually violated in the actual matching generated by the exchange (which may not occur if k is assigned to a school she prefers over $\text{DA}_j(P)$).

We now connect justifiability to the labelled DA envy digraph. Recall that a cycle packing Π consists of disjoint cycles in $G^{\text{DA}(P)}$ and induces a Pareto improvement over $\text{DA}(P)$ by assigning each student in a cycle the DA school of her successor. The set $l(\Pi)$ collects the improvable students whose priorities may be violated by executing the trades in Π . This allows us to express justifiability as a simple condition on labels.

LEMMA 2. *Fix a problem P and a cycle packing Π with induced matching μ^Π . The matching $\mu^\Pi \in \mathcal{M}(P)$ is justifiable if and only if:*

$$l(\Pi) \subseteq \mathcal{B}(\mu^\Pi).$$

PROOF. Suppose first that μ^Π is justifiable, and let $h \in l(\Pi)$. Then $h \in l(i \rightarrow j)$ for some edge $i \rightarrow j$ in Π . Thus h is improvable, prefers $DA_j(P)$ to $DA_h(P)$, and has higher priority at $DA_j(P)$ than i . If $h \notin \mathcal{B}(\mu^\Pi)$, then h does not belong to any cycle in Π , so $\mu_h^\Pi = DA_h(P)$. Hence assigning i to $DA_j(P)$ creates a priority violation against h under μ^Π . Since h is improvable and not a beneficiary, this violation is unjustifiable, contradicting that μ^Π is justifiable. Therefore $h \in \mathcal{B}(\mu^\Pi)$, and so $l(\Pi) \subseteq \mathcal{B}(\mu^\Pi)$.

Conversely, suppose that $l(\Pi) \subseteq \mathcal{B}(\mu^\Pi)$. Consider μ^Π against some student h at a school $DA_j(P)$ assigned to student i via an edge $i \rightarrow j$ in Π . If $h \in \mathcal{U}(P)$ the violation is justifiable by definition, so suppose $h \in \mathcal{I}(P)$. Then h prefers $DA_j(P)$ to μ_h^Π , and has higher priority at $DA_j(P)$ than i . Since μ^Π Pareto-dominates $DA(P)$, we have $\mu_h^\Pi \succeq_h DA_h(P)$, so in particular h prefers $DA_j(P)$ to $DA_h(P)$. Thus $h \in l(i \rightarrow j) \subseteq l(\Pi)$. By assumption, $h \in \mathcal{B}(\mu^\Pi)$. Hence every priority violation under μ^Π is justifiable, and therefore μ^Π is justifiable.

□

In particular, note that a matching μ^Π is justifiable if the following local condition on edges holds: $l(\Pi) = \emptyset$. In such case, we will say that the corresponding matching is *strongly justifiable*. Strongly justifiable matchings have minimal priority violations, but are of interest for a technical reason too. To verify that a matching is justifiable, one must verify each label in a cycle against every node who is part of the cycle packing. On the other hand, this is much easier in a strongly justifiable matching: we only need to check that each label in a cycle is \emptyset , and thus the justifiability property can be verified locally rather than globally.

4 Results

4.1 Existence and Structure

We first show that our notion of justifiability is not vacuous, and gives us at least one Pareto-improvement over DA whenever DA's matching is inefficient. To do so, we describe a simple mechanism that constructs such an improvement using the

set of schools who reject at least one improvable student.

We will use the following notation. Let $S^*(P) \subseteq S$ denote the set of schools that reject at least one improvable student during the execution of DA. For each school $s \in S^*(P)$, the *cutoff student* at s , denoted by $\underline{i}(s)$, is the lowest-priority student assigned to s in DA. Similarly, for each $s \in S^*(P)$, the *below cutoff set* of s is

$$A(s) := \{i \in \mathcal{I}(P) : s \succ_i \text{DA}_i(P) \text{ and } \underline{i}(s) \succ_s i\},$$

i.e., the improvable students who prefer s to their DA assignment but have lower priority at s than the cutoff. Note that, since $s \in S^*(P)$, the set $A(s)$ is non-empty. Thus, for any school $s \in S^*(P)$, we can find the *just below cutoff student* at s , denoted i_s^* , which is the highest-priority student in $A(s)$ according to \succ_s .

The Just-Below Cutoffs (JBC) mechanism.

- (1) Construct a directed graph on $S^*(P)$ by adding, for each $s \in S^*(P)$, a directed edge $s \rightarrow \text{DA}_{i_s^*}(P)$ (recall i_s^* is the just below cutoff student).
- (2) Since every node has out-degree exactly one, the digraph contains at least one cycle. For every such cycle $(s_1 \rightarrow s_2 \rightarrow \dots \rightarrow s_k \rightarrow s_1)$, execute the corresponding trades simultaneously: each student $i_{s_\ell}^*$ moves from $s_{\ell+1}$ to s_ℓ (indices mod k).

We use $\text{JBC}(P)$ to denote the matching produced by JBC. We show below that JBC always produces a strongly justifiable improvement whenever DA is Pareto-inefficient.

THEOREM 1. *Every problem with an inefficient DA outcome admits a strongly justifiable improvement.*

Moreover, the set of strongly justifiable matchings coincides with the set of matchings obtained by executing subsets of the disjoint cycles found by JBC.

PROOF. Assume $\text{DA}(P)$ is Pareto-inefficient. Then $\mathcal{I}(P) \neq \emptyset$, and therefore $S^*(P) \neq \emptyset$. For each $s \in S^*(P)$, the set $A(s)$ is non-empty by definition, so i_s^* is well-defined.

We first show that the school graph used by JBC is well-defined. Fix $s \in S^*(P)$. Since $i_s^* \in \mathcal{I}(P)$, Lemma 1 implies that i_s^* belongs to a cycle in $G^{\text{DA}(P)}$. Let h be the

predecessor of i_s^* on such a cycle. Then h prefers $DA_{i_s^*}(P)$ to $DA_h(P)$, so during the execution of DA student h must have applied to school $DA_{i_s^*}(P)$ and been rejected there. Since h is improvable, it follows that $DA_{i_s^*}(P) \in S^*(P)$.

Thus the map

$$f : S^*(P) \rightarrow S^*(P), \quad f(s) := DA_{i_s^*}(P)$$

is well-defined. Because $S^*(P)$ is finite and each school in $S^*(P)$ points to exactly one school under f , the graph generated by f contains disjoint cycles. JBC executes all such cycles simultaneously: for each cycle

$$(s_1 \rightarrow s_2 \rightarrow \cdots \rightarrow s_k \rightarrow s_1),$$

student $i_{s_\ell}^*$ moves from $s_{\ell+1}$ to s_ℓ for every $\ell = 1, \dots, k$ (indices modulo k). This new matching is feasible because, along each cycle, every school loses exactly one student and gains exactly one student, and different cycles do not share any school. It is also Pareto-improving, since $i_{s_\ell}^* \in A(s_\ell)$ implies that $i_{s_\ell}^*$ strictly prefers s_ℓ to her DA assignment $s_{\ell+1} = DA_{i_{s_\ell}^*}(P)$, while all other students keep their DA assignments.

We next show that the matching produced by JBC is strongly justifiable. Fix $s \in S^*(P)$, and let $j = \underline{l}(s)$ be the cutoff student at s , i.e. the lowest-priority student assigned to s under DA. In the student envy digraph, student i_s^* envies j (since i_s^* prefers s to her DA assignment and $DA_j(P) = s$), so the JBC trade at school s corresponds to the edge $i_s^* \rightarrow j$ in $G^{\text{DA}(P)}$. The label of this edge is

$$l(i_s^* \rightarrow j) = \{h \in \mathcal{I}(P) \cap N^-(j) : \text{rk}_s(h) < \text{rk}_s(i_s^*)\}.$$

We claim that $\mathcal{I}(P) \cap N^-(j) = A(s)$. Every student in $A(s)$ is improvable, prefers $s = DA_j(P)$ to her DA assignment, and has lower priority at s than j , so she envies j and belongs to $\mathcal{I}(P) \cap N^-(j)$. Conversely, take any $h \in \mathcal{I}(P) \cap N^-(j)$. Then h prefers $s = DA_j(P)$ to $DA_h(P)$. By stability of $DA(P)$, school s is full and every student assigned to s has higher priority at s than h ; in particular $\text{rk}_s(h) > \text{rk}_s(\underline{l}(s))$, so $h \in A(s)$. Hence

$$l(i_s^* \rightarrow j) = \{h \in A(s) : \text{rk}_s(h) < \text{rk}_s(i_s^*)\}.$$

Since i_s^* is the highest-priority student in $A(s)$, this set is empty. So every edge used by JBC is empty-labelled, and the matching produced by JBC is strongly justifiable. This proves the first statement of the theorem.

For the second statement, let μ^Π be any strongly justifiable matching induced by a cycle packing Π in $G^{\text{DA}(P)}$. Since μ^Π is strongly justifiable, every edge in Π is empty-labelled. Take any school s that is entered by some student under μ^Π , and let i be the student who enters s . Since μ^Π is induced by a cycle packing in $G^{\text{DA}(P)}$, there exists some student j assigned to s under DA such that $i \rightarrow j$ is an edge of Π .

We claim that $i = i_s^*$. First, $i \in A(s)$: because $i \rightarrow j$ is an edge of $G^{\text{DA}(P)}$, student i is improvable and prefers $s = \text{DA}_j(P)$ to $\text{DA}_i(P)$. Moreover, since j is assigned to s under DA and DA is stable, every student assigned to s under DA has higher priority at s than i ; in particular $\underline{i}(s) \succ_s i$. Hence $i \in A(s)$.

Next, let $h \in A(s)$. Then $h \in \mathcal{I}(P)$, h prefers s to $\text{DA}_h(P)$, and $\underline{i}(s) \succ_s h$. Because j is assigned to s under DA, we also have $j \succ_s \underline{i}(s)$, and therefore $j \succ_s h$. Thus h envies j , so $h \in \mathcal{I}(P) \cap N^-(j)$. We have shown that

$$A(s) \subseteq \mathcal{I}(P) \cap N^-(j).$$

Since the edge $i \rightarrow j$ is empty-labelled,

$$l(i \rightarrow j) = \{h \in \mathcal{I}(P) \cap N^-(j) : \text{rk}_s(h) < \text{rk}_s(i)\} = \emptyset.$$

Therefore there is no student in $A(s)$ with higher priority than i at s . Since $i \in A(s)$, it follows that i is the highest-priority student in $A(s)$, namely i_s^* . So in any strongly justifiable matching, the student who enters school s must be i_s^* .

Now suppose i_s^* enters s under μ^Π . Then i_s^* leaves her DA school $\text{DA}_{i_s^*}(P) = f(s)$. Since μ^Π is obtained from a cycle packing, the seat vacated at $f(s)$ must in turn be filled by some student who enters $f(s)$ along an empty-labelled edge. By the previous paragraph, that student must be $i_{f(s)}^*$. Repeating the same argument, whenever a school s is used in a strongly justifiable cycle packing, the school $f(s)$ must also be used. Since only finitely many schools are involved, this process must run along a cycle of the graph generated by f . It follows that every strongly justifiable cycle packing is obtained by selecting some of the cycles found by JBC and executing the corresponding trades. Conversely, any subset of the cycles found

Table 1. Cycle Packings in Example 1

Cycle Packing	SJ?	J?	PE?
$(i_1 \xrightarrow{\emptyset} i_2 \xrightarrow{i_5} i_1)$	No	No	No
$(i_1 \xrightarrow{i_4} i_5 \xrightarrow{\emptyset} i_1)$	No	No	No
$(i_4 \xrightarrow{\emptyset} i_5 \xrightarrow{i_1, i_6} i_4)$	No	No	No
$(i_1 \xrightarrow{i_3, i_5} i_6 \xrightarrow{i_1} i_4 \xrightarrow{\emptyset} i_5 \xrightarrow{\emptyset} i_1)$	No	No	Yes
$(i_3 \xrightarrow{\emptyset} i_6 \xrightarrow{i_1} i_4 \xrightarrow{\emptyset} i_5 \xrightarrow{\emptyset} i_3)$	No	No	No
$(i_1 \xrightarrow{\emptyset} i_4 \xrightarrow{\emptyset} i_5 \xrightarrow{\emptyset} i_1)$	Yes	Yes	No
$(i_1 \xrightarrow{i_5} i_3 \xrightarrow{\emptyset} i_6 \xrightarrow{i_1} i_4 \xrightarrow{\emptyset} i_5 \xrightarrow{\emptyset} i_1)$	No	Yes	No
$\{(i_1 \xrightarrow{\emptyset} i_2 \xrightarrow{i_5} i_1), (i_3 \xrightarrow{\emptyset} i_6 \xrightarrow{i_1} i_4 \xrightarrow{\emptyset} i_5 \xrightarrow{\emptyset} i_3)\}$	No	Yes	Yes

SJ: strongly justifiable; J: justifiable; PE: Pareto-efficient.

Example 1 illustrates that JBC selects a strongly justifiable improvement but the resulting matching need not be Pareto-efficient. At the same time, justifiability allows for further efficiency gains once we allow for edges that are not necessarily empty-labelled (Table 1).⁶ In particular, in our Example we can find a cycle packing that is both justifiable and Pareto-efficient. However, this fortunate coincidence does not hold in general, as we show below.


PROPOSITION 1. *There exists a problem P that admits no matching that is both justifiable and Pareto-efficient.*

PROOF. Consider the school choice problem in the Example 2 below, with six students and six schools with unit capacity. The DA outcome assigns each student to the school shown in bold.

⁶By this statement we do not mean that every JBC beneficiary improves upon her DA placement in every justifiable matching; see Example 4 in Appendix D for a counterexample. This is why the set of justifiable matchings does not have any nice lattice structure. What we mean is that our main algorithm will build on JBC to generate justifiable matchings.

EXAMPLE 2 (PREFERENCES AND PRIORITIES (LEFT); ENVY DIGRAPH (RIGHT)).

i_1	i_2	i_3	i_4	i_5	i_6	s_1	s_2	s_3	s_4	s_5	s_6
s_5	s_6	s_1	s_2	s_4	s_2	i_1	i_2	\cdot	i_4	i_5	i_6
s_4	s_1	s_2	s_4	s_5	s_4	i_3	i_4	\cdot	i_3	i_1	i_2
s_1	s_2	s_6			s_6	i_2	i_3	\cdot	i_1		
		s_4					i_6		i_6		
		s_3							i_5		



Step 1 [Only justified improvement]. The labelled DA envy digraph contains exactly four cycles, and all four contain student i_2 . Hence no two cycles are disjoint, so the only non-empty cycle packings are the singleton packings generated by these four cycles. Among them, the only cycle whose traded edges are justifiable is

$$(i_1 \xrightarrow{\emptyset} i_4 \xrightarrow{\emptyset} i_2 \xrightarrow{\emptyset} i_1).$$

Indeed, each traded edge on this cycle has empty label, so executing it violates no improvable student’s priority, and therefore the induced improvement is (strongly) justifiable.

By contrast, each of the other three cycles contains at least one traded edge whose label includes an improvable student who is not a beneficiary of that cycle, and thus fails justifiability:

- The cycle $(i_2 \xrightarrow{\emptyset} i_6 \xrightarrow{i_4} i_2)$ is not justifiable because the edge $i_6 \rightarrow i_2$ has label $\{i_4\}$, so executing the cycle violates the priority of i_4 , who is not improved by that cycle.
- The cycle $(i_2 \xrightarrow{\emptyset} i_6 \xrightarrow{i_1} i_4 \xrightarrow{\emptyset} i_2)$ is not justifiable because the edge $i_6 \rightarrow i_4$ has label $\{i_1\}$, so it violates the priority of i_1 , who is not improved by that cycle.
- The cycle $(i_2 \xrightarrow{\emptyset} i_1 \xrightarrow{\emptyset} i_5 \xrightarrow{i_1, i_6} i_4 \xrightarrow{\emptyset} i_2)$ is not justifiable because the edge $i_5 \rightarrow i_4$ has label $\{i_6\}$ (in particular), so executing the cycle violates the priority of i_6 , who is not improved by that cycle.

Consequently, the unique justifiable Pareto improvement over $DA(P)$ in Example 2 is the one induced by the 3-cycle above; denote the resulting matching by μ^J .

Step 2 [The justifiable improvement is not efficient]. We claim that μ^J is not Pareto-efficient. After executing the 3-cycle, students i_1 and i_5 both strictly prefer to exchange their assigned schools (as is immediate from the preferences in Example 2), so there exists a further Pareto improvement over μ^J that assigns i_1 to $\mu^J(i_5)$ and i_5 to $\mu^J(i_1)$, leaving all other students unchanged. However, any matching that implements this exchange violates the priority of student i_6 at school s_1 (the school assigned to one of $\{i_1, i_5\}$ under μ^J), and i_6 is an improvable non-beneficiary of the exchange. Hence this further improvement cannot be justifiable.

Since μ^J is the unique justifiable improvement over $\text{DA}(P)$ in Example 2, and μ^J is not Pareto-efficient, it follows that Example 2 admits no matching that is both justifiable and Pareto-efficient. \square

Proposition 1 shows that justifiability and full Pareto-efficiency are generally incompatible. When restricted to justifiable improvements, therefore, we must settle for constrained-efficiency. In the next subsection, we introduce a JBC-related algorithm, SJBC+, that always returns a justifiable matching that satisfies one such notion.

4.2 The Sequential Just Below Cutoffs Algorithm

JBC allows only empty-labelled edges, so only unimprovable students' priorities may be violated. But justifiability allows more: a priority violation against an improvable student is legitimate whenever that student benefits from the realized improvement. This opens the door to a richer class of improvements through sequential iteration. JBC produces an initial set of beneficiaries B^1 . In the next round, every edge whose label is contained in B^1 becomes admissible. Executing such edges may create new beneficiaries, which enlarges the admissible set further. Iterating this procedure yields an expansion phase that stops when no additional beneficiary can be created.

To find one of the largest cycle packings at each iteration, we represent the problem as a bipartite graph H^t . Both sides contain one node per improvable student. For each edge $i \rightarrow j$ in $G^{\text{DA}(P)}$ with $l(i \rightarrow j) \subseteq B^t$, we add an edge from i (left) to j (right). We also add a self-loop $i \rightarrow i$ for every improvable student,

representing remaining at the DA assignment. A perfect matching in H^t then corresponds to a cycle packing: matched pairs $i \rightarrow j$ with $i \neq j$ form trading cycles, while self-loops represent uncovered students. Hence the relevant objective is not the cardinality of a perfect matching in H^t —all perfect matchings have the same size—but rather the number of non-loop edges it contains, which is exactly the number of beneficiaries created at that iteration.

Equivalently, let \widehat{H}^t denote the bipartite graph obtained from H^t by deleting all self-loops. A matching in \widehat{H}^t specifies the students who trade at step t , and every such matching can be completed uniquely to a perfect matching in H^t by assigning each unmatched student to herself. Therefore maximizing the number of non-loop edges in a perfect matching of H^t is equivalent to computing a maximum-cardinality matching in \widehat{H}^t .⁷

Our proposed algorithm iterates this maximum-cardinality matching step on \widehat{H}^t , and then adds self-loops for any uncovered students.

The Sequential Just Below Cutoffs (SjBC+) algorithm.

- (1) Set $\mu^0 = \text{JBC}(P)$ and $B^1 = \mathcal{B}(\mu^0)$.

Build the bipartite graph \widehat{H}^0 whose left and right vertex sets are both copies of $\mathcal{I}(P)$, and whose edges are all pairs $i \rightarrow j$ in $G^{\text{DA}(P)}$ with $l(i \rightarrow j) = \emptyset$. Let M^0 be the unique maximum-cardinality matching in \widehat{H}^0 corresponding to μ^0 .⁸ Complete M^0 to a perfect matching in the augmented graph H^0 by adding a self-loop $i \rightarrow i$ for every uncovered student $i \in \mathcal{I}(P)$.

- (2) For each $t = 1, 2, \dots$, given B^t , build the bipartite graph \widehat{H}^t whose left and right vertex sets are both copies of $\mathcal{I}(P)$, and whose edges are all pairs $i \rightarrow j$ in $G^{\text{DA}(P)}$ satisfying $l(i \rightarrow j) \subseteq B^t$.

Start from the non-loop matching induced by the previous iteration. Using augmenting paths in \widehat{H}^t , compute a maximum-cardinality matching M^t in \widehat{H}^t . Complete M^t to a perfect matching in the augmented graph H^t by adding a self-loop $i \rightarrow i$ for every uncovered student $i \in \mathcal{I}(P)$.

⁷This is a standard bipartite matching problem, solvable for example by the Hopcroft–Karp algorithm; see [Cormen et al. \[2009, Chapter 26\]](#).

⁸Uniqueness in this first iteration follows from [Theorem 1](#).

Let μ^t be the matching induced by this perfect matching, and set

$$B^{t+1} = \mathcal{B}(\mu^t).$$

If $B^{t+1} = B^t$, stop the expansion phase and let $\mu^* = \mu^t$. Otherwise continue to iteration $t + 1$.

- (3) Starting from μ^* , run the refinement phase. At each refinement iteration, call an edge $i \rightarrow j$ in the envy digraph at the current matching *admissible* if every improvable student $h \neq j$ who prefers the school currently assigned to j to $DA_h(P)$ and has higher priority than i at that school belongs to B^* . Here $B^* = \mathcal{B}(\mu^*)$. Repeatedly find any directed cycle consisting only of admissible edges and execute it. Continue until no such cycle remains.
- (4) Output the final matching μ^+ .

We call any matching produced by this procedure an SJBC+ outcome.⁹ The outcome need not be unique: different selections of maximum matchings in the expansion phase or different collections of cycles in the refinement phase may yield different matchings. In the expansion phase we compute a maximum-cardinality matching in the bipartite graph of admissible non-loop edges at each iteration, and then complete uncovered students with self-loops. Equivalently, in the augmented graph with self-loops, we select a perfect matching that minimizes the number of self-loops.

We illustrate the execution of SJBC+ on Example 1 (for a graphical version see Figure 2). Step 1 yields the JBC cycle $(i_1 \xrightarrow{\emptyset} i_4 \xrightarrow{\emptyset} i_5 \xrightarrow{\emptyset} i_1)$ and $B^1 = \{i_1, i_4, i_5\}$, with self-loops on i_2, i_3 and i_6 . At $t = 1$, the newly added edges with labels contained in B^1 are: $i_1 \xrightarrow{i_5} i_3$, $i_1 \xrightarrow{i_4} i_5$, $i_2 \xrightarrow{i_5} i_1$, and $i_6 \xrightarrow{i_1} i_4$. An augmenting path, which alternates between edges not in the current matching and edges in the current matching, starts at an unmatched vertex on the left-hand side of \widehat{H}^t (equivalently, a student currently assigned to a self-loop in the augmented graph H^t), namely i_2^L (the superscript L denotes the LHS), who is now allowed to traverse to i_1^R . The augmenting path alternates:

- (1) $i_2^L \rightarrow i_1^R$ (unmatched, newly available),

⁹See Appendix C for an example showing why the refinement step is necessary.

- (2) $i_1^R \rightarrow i_5^L$ (matched),
- (3) $i_5^L \rightarrow i_3^R$ (unmatched),
- (4) $i_3^R \rightarrow i_3^L$ (matched, self-loop),
- (5) $i_3^L \rightarrow i_6^R$ (unmatched),
- (6) $i_6^R \rightarrow i_6^L$ (matched, self-loop),
- (7) $i_6^L \rightarrow i_4^R$ (unmatched, newly available),
- (8) $i_4^R \rightarrow i_1^L$ (matched),
- (9) $i_1^L \rightarrow i_2^R$ (unmatched),
- (10) $i_2^R \rightarrow i_2^L$ (matched, self-loop).

The path returns to i_2 , where it started. Then we flip, so that we preserve only new edges, and execute the corresponding cycle:

$$(i_1 \xrightarrow{\emptyset} i_2 \xrightarrow{i_5} i_1), (i_3 \xrightarrow{\emptyset} i_6 \xrightarrow{i_1} i_4 \xrightarrow{\emptyset} i_5 \xrightarrow{\emptyset} i_3)$$

covering all six improvable students. Thus $B^* = \{i_1, i_2, i_3, i_4, i_5, i_6\}$, the unique justifiable and Pareto-efficient matching from Table 1.

Any SJBC+ outcome satisfies several important desiderata.

THEOREM 2. *If DA(P) is inefficient, any SJBC+ outcome μ^+ :*

- (1) *Pareto-dominates DA(P);*
- (2) *is justifiable;*
- (3) *is not Pareto-dominated by any justifiable matching μ' unless $\mathcal{B}(\mu^+) \subsetneq \mathcal{B}(\mu')$;*
- (4) *satisfies $\mathcal{B}(\text{JBC}(P)) \subseteq \mathcal{B}(\mu^+)$*

Furthermore, any SJBC+ outcome can be computed in time polynomial in $|I|$ and $|S|$.

PROOF. [Part 1] By Theorem 1, JBC(P) Pareto-dominates DA(P). The expansion phase maintains $B^t \subseteq B^{t+1}$ and ensures every student in B^t strictly prefers μ^t to DA(P). At termination, every student in $B^* \supseteq \mathcal{B}(\text{JBC}(P))$ strictly prefers μ^* to DA. The refinement step permutes schools only among B^* via cycles in which each participant improves over her current assignment, hence strictly improves over DA. Students outside B^* retain their DA assignments throughout. Thus μ^+ Pareto-dominates DA(P).

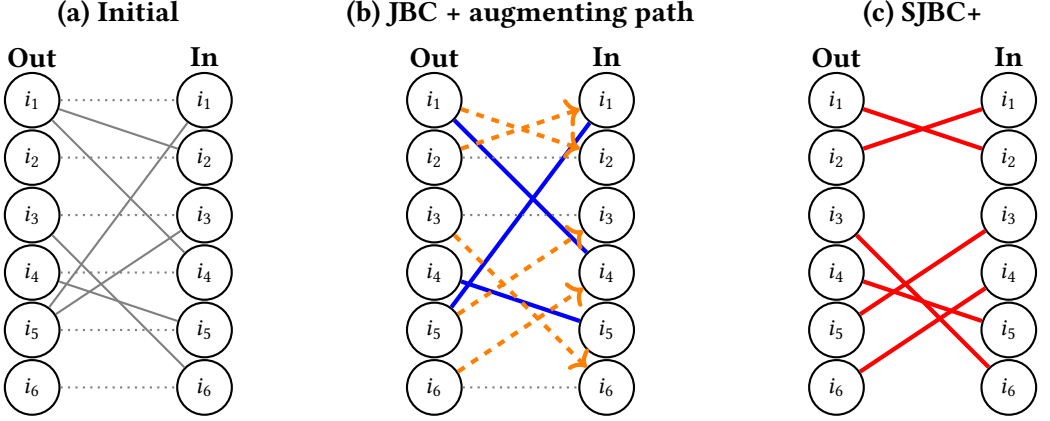


Fig. 2. Expansion phase in Example 1. Panel (a): after JBC, the JBC cycle (solid blue) covers i_1, i_4, i_5 ; remaining students are on self-loops (dotted). Panel (b): an augmenting path (dashed orange) reroutes through the existing matching, incorporating i_2, i_3, i_6 . Panel (c): after flipping, all six students are on non-trivial cycles (red).

[Part 2] Every edge implemented during the expansion phase satisfies $l(i \rightarrow j) \subseteq B^t \subseteq B^*$ at the step in which it is used. Hence every priority violation created during the expansion phase is against a student in $B^* = \mathcal{B}(\mu^*)$. By Lemma 2, μ^* is therefore justifiable. Since the refinement permutes assignments only among B^* and each refinement-cycle participant strictly improves over her current school, $\mathcal{B}(\mu^+) = B^*$. It suffices to show that no improvable non-beneficiary suffers a priority violation under μ^+ . Suppose for contradiction that $h \in \mathcal{I}(P) \setminus B^*$ prefers school $s = \mu_i^+$ to $\mu_h^+ = DA_h(P)$ and $\text{rk}_s(h) < \text{rk}_s(i)$. If $\mu_i^* = s$, the same violation exists under μ^* , contradicting its justifiability. Otherwise i received s during the refinement phase: specifically, i took the school of some student j along a cycle of admissible edges. Because h is improvable, $h \neq j$ (since $j \in B^*$), h prefers s to $DA_h(P)$, and $\text{rk}_s(h) < \text{rk}_s(i)$, admissibility of the edge $i \rightarrow j$ requires $h \in B^*$, a contradiction.

[Part 3] Suppose for contradiction that there exists a justifiable matching μ' that Pareto-dominates μ^+ but does not satisfy $\mathcal{B}(\mu^+) \subseteq \mathcal{B}(\mu')$.

Because μ' Pareto-dominates μ^+ , every student in $\mathcal{B}(\mu^+)$ weakly prefers μ' to μ^+ , and since each such student strictly prefers μ^+ to $\text{DA}(P)$, she also strictly prefers μ' to $\text{DA}(P)$. Hence $\mathcal{B}(\mu^+) \subseteq \mathcal{B}(\mu')$. By assumption this inclusion is not strict, so $\mathcal{B}(\mu') = \mathcal{B}(\mu^+)$. By Parts 1 and 4, $\mathcal{B}(\mu^+) = B^*$, and therefore $\mathcal{B}(\mu') = B^*$ as well.

Since students outside B^* receive their DA assignments under both μ^+ and μ' , the two matchings differ only in how they assign the schools held by students in B^* . Equivalently, μ' is obtained from μ^+ by permuting the assignments of students in B^* . This permutation decomposes into disjoint cycles. Along each such cycle, every student receives under μ' the school held under μ^+ by her successor. Since μ' Pareto-dominates μ^+ , every student on every cycle weakly prefers her successor's school to her own under μ^+ , and on at least one nontrivial cycle at least one student strictly prefers it.

Fix such a cycle, and consider any edge $i \rightarrow j$ on it, so that under μ' student i receives school μ_j^+ . Take any $h \in \mathcal{I}(P) \setminus \{j\}$ who prefers μ_j^+ to $\text{DA}_h(P)$ and satisfies $\text{rk}_{\mu_j^+}(h) < \text{rk}_{\mu_j^+}(i)$. If $h \notin B^*$, then $\mathcal{B}(\mu') = B^*$ implies $\mu'_h = \text{DA}_h(P)$. But then h prefers $\mu'_i = \mu_j^+$ to μ'_h and has higher priority than i at that school, so μ' creates an actual priority violation against an improvable non-beneficiary, contradicting justifiability. Hence every such h belongs to B^* , and therefore every edge on this cycle is admissible at μ^+ .

Therefore this directed cycle would be available at termination, contradicting the stopping rule of the refinement phase. Hence no such μ' exists.

[Part 4] The expansion phase initializes $B^1 = \mathcal{B}(\text{JBC}(P))$ and maintains $B^t \subseteq B^{t+1}$. Since flipping along augmenting paths only extends the matching, every student in B^1 remains a beneficiary under μ^t for all $t \geq 1$, and in particular under μ^* . The refinement step implements trades only among students in B^* , so students in $\mathcal{B}(\text{JBC}(P)) \subseteq B^*$ remain beneficiaries under μ^+ . Hence $\mathcal{B}(\text{JBC}(P)) \subseteq \mathcal{B}(\mu^+)$.

[Part 5] DA runs in $O(|I| \cdot |S|)$ time, since each student applies to each school at most once, so there are at most $|I| \cdot |S|$ proposals in total, and each proposal is processed once.

The labelled envy digraph has at most $|I|(|I| - 1)$ edges. For each edge $i \rightarrow j$, its label is obtained by scanning the set of improvable students who prefer $\text{DA}_j(P)$ to

their DA assignment and checking which of them have higher priority than i at $DA_j(P)$. Thus all labels can be computed in polynomial time.

JBC constructs the auxiliary digraph on $S^*(P)$ by assigning to each school $s \in S^*(P)$ the unique outgoing edge to $DA_{i_s^*}(P)$. Since this digraph has at most $|S|$ nodes and out-degree one at every node, its directed cycles can be identified in $O(|S|)$ time.

The expansion phase runs for at most $|I|$ iterations, since B^t strictly increases at each successful iteration. At each iteration, the expansion phase computes a maximum-cardinality matching in the bipartite graph \widehat{H}^t of admissible non-loop edges, where the two sides are copies of the improvable students. Hence $|V(\widehat{H}^t)| \leq 2|I(P)|$ and $|E(\widehat{H}^t)| \leq |I(P)|^2$, so each iteration can be solved in $O(|E(\widehat{H}^t)|\sqrt{|V(\widehat{H}^t)|})$ time by the Hopcroft-Karp algorithm; see [Cormen et al. \[2009, Chapter 26\]](#). Completing the resulting matching with self-loops is linear-time. Therefore the expansion phase runs in polynomial time.

For the refinement phase, fix an iteration r . The restricted envy digraph on B^* has at most $|B^*|(|B^*| - 1)$ edges. Admissibility of each edge $i \rightarrow j$ is checked by scanning the students in $I(P) \setminus \{j\}$ and verifying preference and priority at school $\mu_{\text{ref},j}^r$. Hence the admissible subgraph can be computed in polynomial time.

A maximal collection of vertex-disjoint directed cycles in the admissible subgraph can also be found in polynomial time: repeatedly find any directed cycle in the current graph by depth-first search, record it, delete its vertices, and continue until no directed cycle remains. Since each search and deletion step is polynomial and at most $|B^*|$ vertices are deleted overall, this procedure is polynomial.

At each refinement iteration, every student who moves along a selected cycle strictly prefers her new assignment to her current one. Therefore the sum of the ranks of the current assignments of students in B^* strictly decreases at every successful refinement iteration. Since each such rank lies between 1 and $|S|$, this sum can decrease at most $|B^*|(|S| - 1)$ times. Hence the refinement phase performs at most $|B^*|(|S| - 1)$ successful iterations, and each iteration is polynomial-time.

Therefore SJBC+ runs in time polynomial in $|I|$ and $|S|$. □

We make an observation regarding the efficiency guarantees of SJBC+. Part 3 of Theorem 2 guarantees that no SJBC+ outcome is Pareto-dominated by any justifiable matching unless that matching improves every SJBC+ beneficiary and more. A stronger property would be that no SJBC+ outcome is dominated by any justifiable matching, which we cannot guarantee in general (see Appendix E for an example). We can attain this stronger guarantee at the cost of losing the polynomial-time algorithm. To see this, after the sequential expansion converges, add all remaining non-loop edges to the bipartite graph and compute a maximum-cardinality matching: if the result is justifiable, it covers weakly more beneficiaries than any justifiable matching; if not, one can enumerate subsets of $\mathcal{I}(P) \setminus B^*$ to find the largest justifiable improvement, at the cost of losing polynomiality (with the refinement step still needed in the end).¹⁰ In practice, however, this stronger guarantee is rarely needed: SJBC+ already covers the vast majority of improvable students, and as we show in Section 6, it frequently achieves the even stronger property of full Pareto-efficiency: it is not dominated by any matching, justifiable or not.

Before describing SJBC+ performance in simulations, we address two questions. First, how our justifiability framework is inherently different from existing approaches. Second, how consent-based mechanisms also face limitations in their ability to restore efficiency if consent is not self-harming.

5 EADA and Ex-ante Consent

So far, we have treated consent as an endogenous constraint: whether a priority violation is permissible depends on who ultimately benefits in the realized improvement. An alternative perspective treats consent as an ex-ante input rather than an ex-post justification. Under this view, a subset of students agrees in advance to waive priority, and the designer is asked to improve upon the DA outcome while continuing to respect the priorities of all non-consenting students.

We can formalize this perspective through the notion of a consent-based mechanism. A consent-based mechanism takes as input a school choice problem P and

¹⁰It remains an open question whether there exists a polynomial time algorithm that finds a justifiable improvement upon DA that is not Pareto dominated by any other justifiable matching.

a consent set $W \subseteq I$, interpreted as the set of students whose priorities may be violated, and returns a matching that respects the priorities of $I \setminus W$, i.e. there do not exist $i \in I \setminus W$, $j \in I$, and $s \in S$ such that (i) i prefers s to μ_i , (ii) $\mu_j = s$, and (iii) $\text{rk}_s(i) < \text{rk}_s(j)$.

The celebrated Efficiency-Adjusted Deferred Acceptance mechanism (EADA; [Kesten, 2010](#)) is the leading example of a consent-based mechanism. For a consent set $W \subseteq I$, let $\text{EADA}(P, W)$ denote the matching produced by EADA when the set of consenting students is W ; we write $\text{EADA}_i(P, W)$ for student i 's assignment under $\text{EADA}(P, W)$.¹¹ It is easy to see that, when only a subset of unimprovable students consent, EADA generates a justifiable matching, since by design EADA only violates priorities of consenting students. This observation may lead the reader to conclude that our justifiability framework may be strongly linked to EADA's outcomes. This is not so, as [Theorem 3](#) shows.

THEOREM 3. *There exists a problem P such that $\text{EADA}(P, I)$ is not justifiable, and the unique justifiable and Pareto-efficient matching is not equal to $\text{EADA}(P, W)$ for any $W \subseteq I$.*

PROOF. [Part 1] In [Example 1](#), the EADA outcome under full consent implements the following trades

$$(i_1 \xrightarrow{i_3, i_5} i_6 \xrightarrow{i_1} i_4 \xrightarrow{\emptyset} i_5 \xrightarrow{\emptyset} i_1)$$

Note that the priority of i_3 is violated but she is an improvable student who does not benefit from this exchange. Thus, the resulting matching is not justifiable.

[Part 2] For the same example, in [Appendix B](#) we consider all possible last interrupters during EADA's execution, and construct a consent tree that shows each possible outcome. The unique justifiable and Pareto-efficient matching is never produced. □

[Figure 3](#) below presents the improvement executed by EADA with full consent, as well as the unique justifiable and Pareto-efficient matching in [Example 1](#).¹²

¹¹EADA is well-known in the literature and thus we postpone its description to [Appendix A](#).

¹²In this example, the unique justifiable and Pareto-efficient matching improves more students and generates fewer blocking pairs than EADA with full consent; this double-dominance of the EADA outcome is explored by [Knipe and Ortega \[2026\]](#).




Fig. 3. EADA with full consent (left) and the justifiable and efficient matching (right) in Example 1.

An immediate consequence of Theorem 3, together with the fact that under full consent EADA selects the unique student-optimal priority-efficient assignment and the unique student-optimal legal assignment, is that justifiability is also at odds with other prominent frameworks for controlled priority violations.¹³

COROLLARY 1. *The student-optimal priority-efficient assignment and the student-optimal legal assignment need not be justifiable. Conversely, there exists a problem P with a unique Pareto-efficient and justifiable matching μ that is neither legal nor priority-neutral (and hence not priority-efficient).*

A similar statement can be made regarding essential stability, a weakening of stability that is always satisfied by EADA with full consent [Trojan et al., 2020].¹⁴ It is a direct implication of Theorem 3 and the aforementioned property that essentially stable matchings need not be justifiable. One can straightforwardly verify (as we do in Appendix F) that the unique justifiable and Pareto-efficient matching in Example 1 admits a non-vacuous reassignment chain, showing that a justifiable and Pareto-efficient improvement upon DA need not be essentially stable either.

¹³A matching μ is priority neutral if no matching ν can make any student whose priority is violated by μ better off unless ν violates the priority of some student who is made worse off [Reny, 2022]. A matching ν blocks μ if there exists a student who prefers her assignment under ν to her assignment under μ and who could be admitted to that school ahead of at least one student assigned there under μ . A set L of matchings is legal if (i) every matching outside L is blocked by some matching in L , and (ii) no matching in L blocks another matching in L . Ehlers and Morrill [2020] show that the legal set is unique and contains a student-optimal element.

¹⁴A matching is essentially stable if any priority-based claim initiates a chain of reassignments that results in the initial claimant losing the object.

5.1 The Limitations of Ex-Ante Consent

We now articulate three properties that one might reasonably want from a consent-based mechanism. First and foremost, consent should not be self-harming.

DEFINITION 4. *A consent-based mechanism M incentivizes consent if for all P , all $W \subseteq I$, and all $i \in I$,*

$$\text{rk}_i(M_i(P, W \cup \{i\})) \leq \text{rk}_i(M_i(P, W)).$$

Said differently, every student should be weakly incentivized to consent to waive her priorities. Second, given a fixed consent set, the mechanism should not leave obvious efficiency gains unexploited.

DEFINITION 5. *A consent-based mechanism M is constrained efficient if for all P and $W \subseteq I$, there is no matching μ such that*

- (1) μ respects the priorities of $I \setminus W$, and
- (2) μ Pareto-dominates $M(P, W)$.

EADA satisfies both incentivized consent and constrained efficiency, and these guarantees are central to its appeal to both academics and policymakers.¹⁵ Yet, a different efficiency desideratum is that a consent-based mechanism should return a Pareto-efficient allocation whenever efficiency is compatible with the announced waivers.

DEFINITION 6. *Given P and $W \subseteq I$, let $\mathcal{M}(P, W)$ denote the set of matchings that*

- (1) respect the priorities of $I \setminus W$,
- (2) are Pareto-efficient (among all feasible matchings), and
- (3) Pareto-dominate $DA(P)$.

A consent-based mechanism M is efficient whenever possible if whenever $\mathcal{M}(P, W) \neq \emptyset$, it selects such a matching, i.e., $M(P, W) \in \mathcal{M}(P, W)$.

¹⁵In fact, EADA satisfies a property stronger than incentivizing consent: giving consent does not harm or benefit any consenter [Tang and Yu, 2014].

This benchmark strengthens a well-known property of EADA: when all students consent, EADA restores Pareto-efficiency. Efficiency whenever possible asks for more. It requires that even under partial consent, the mechanism exploit all feasible efficiency gains whenever an efficient, priority-respecting improvement over DA exists. Unfortunately, EADA fails this stronger requirement.

PROPOSITION 2. *For partial consent sets $W \subsetneq I$, $EADA(P, W)$ is not efficient whenever possible: it may fail to be Pareto-efficient even though $\mathcal{M}(P, W) \neq \emptyset$.*

PROOF. Consider Example 1 and let $W = \{i_1, i_5, i_7\}$. Run Kesten's EADA (see a detailed execution in Appendix B). In the first DA run on P , the last interrupter in W is i_7 at s_4 . Since $i_7 \in W$, EADA deletes s_4 from i_7 's preference list and reruns DA. In the resulting DA run, the last interrupter is i_3 at s_6 , and since $i_3 \notin W$ the procedure stops. Hence $EADA(P, W)$ implements the exchange $(i_1 \rightarrow i_4 \rightarrow i_5 \rightarrow i_1)$, so in particular $EADA_{i_1}(P, W) = s_4$ and $EADA_{i_6}(P, W) = s_6$.

Define ν by swapping the assignments of i_1 and i_6 in $EADA(P, W)$ and leaving everyone else unchanged. This yields a feasible matching and makes both i_1 and i_6 strictly better off (in Example 1, i_1 prefers s_6 to s_4 and i_6 prefers s_4 to s_6), so ν Pareto-dominates $EADA(P, W)$. Thus $EADA(P, W)$ is not Pareto-efficient.

On the other hand, with the same consent set W , the cycle packing $\{(i_1 \rightarrow i_2 \rightarrow i_1), (i_3 \rightarrow i_6 \rightarrow i_4 \rightarrow i_5 \rightarrow i_3)\}$ induces a matching that Pareto-dominates $DA(P)$, is Pareto-efficient, and respects the priorities of every non-consenting student in $I \setminus W$ (this can be verified directly from Example 1). Hence $\mathcal{M}(P, W) \neq \emptyset$. \square

Proposition 2 shows that fixing waivers ex-ante does not, by itself, guarantee that available efficiency gains will be realized. This raises a deeper question: is the failure specific to EADA, or does it reflect a more fundamental tension among the desiderata above? Our next Theorem shows that the aforementioned tension is structural in the ex-ante consent framework.

THEOREM 4. *There is no consent-based mechanism M that weakly Pareto-dominates DA and simultaneously satisfies all three properties: (i) efficiency whenever possible, (ii) constrained efficiency, and (iii) incentivizing consent.*

PROOF. Fix the school choice problem P from Example 1. Suppose, by contradiction, that there exists a consent-based mechanism M that weakly Pareto-dominates DA and satisfies efficiency whenever possible, constrained efficiency, and incentivizing consent. Consider the following three nested consent sets:

- (1) $W_1 = \{i_7\}$. In Example 1, the unique improvement over $DA(P)$ that can be implemented without violating the priorities of any non-consenting student in $I \setminus W_1$ is the cycle

$$(i_1 \rightarrow i_4 \rightarrow i_5 \rightarrow i_1).$$

This cycle yields a matching that Pareto-dominates $DA(P)$. By constrained efficiency, $M(P, W_1)$ must implement this improvement.

- (2) $W_2 = \{i_5, i_7\}$. Incentivizing consent implies that i_5 must get s_1 or better, and must violate the priorities only of W_2 . There is only one improvement that achieves this, i.e. exactly the one we found before:

$$(i_1 \rightarrow i_4 \rightarrow i_5 \rightarrow i_1).$$

In particular, note that if we were to allow any cycle with the edge $i_6 \xrightarrow{i_1} i_4$, i_1 's priority at s_4 would be violated. And the 4-node cycle implemented by EADA with full consent is not implementable here because it violates the priority of $i_3 \notin W_2$.

- (3) $W_3 = \{i_1, i_5, i_7\} = W_2 \cup \{i_1\}$. Incentivizing consent requires that i_1 is assigned to s_4 or better. Only two improvement cycles do this, namely:

(a) $(i_1 \xrightarrow{i_3, i_5} i_6 \xrightarrow{i_1} i_4 \xrightarrow{\emptyset} i_5 \xrightarrow{\emptyset} i_1)$

(b) $(i_1 \xrightarrow{\emptyset} i_4 \xrightarrow{\emptyset} i_5 \xrightarrow{\emptyset} i_1)$

Improvement (a) violates the priorities of $i_3 \notin W_3$, so improvement (b) is chosen. Yet, the improvement

$$\{(i_1 \xrightarrow{\emptyset} i_2 \xrightarrow{i_5} i_1), (i_3 \xrightarrow{\emptyset} i_6 \xrightarrow{i_1} i_4 \xrightarrow{\emptyset} i_5 \xrightarrow{\emptyset} i_3)\}$$

yields an efficient matching that Pareto-dominates $DA(P)$ while respecting the priorities of all non-consenting students in $I \setminus W_3$. Since M is efficient whenever possible, it must satisfy $M(P, W_3) \in \mathcal{M}(P, W_3)$, contradicting the fact that

$M(P, W_3)$ is induced by (b) and is inefficient. The contradiction completes the proof.

□

Theorem 4 identifies an inherent limitation of ex-ante consent designs. If consent is treated as an arbitrary input, then the natural objective of exploiting all feasible efficiency gains under partial consent conflicts with the very incentive requirement needed to elicit consent in the first place.

Taken together with our earlier results on endogenous justifiability, a coherent picture emerges that shows that achieving efficiency and Pareto-dominating DA comes at important costs, even when some priorities can be violated. On our endogenous framework, it requires violating the priorities of non-beneficiaries. On the ex-ante side, even when priority waivers are explicitly granted in advance, robustly achieving efficiency under partial consent runs into an equally sharp obstacle: the mechanism cannot systematically use the available waivers without sometimes undermining the incentives that make waivers available in the first place.

Having clarified these theoretical limits, both under endogenous justifiability and under ex-ante consent, we now turn to the empirical question of how our proposed SJBC+ algorithm fares in practice. The next section compares SJBC+ and EADA in simulations, shedding light on the practical relevance of the theoretical trade-offs identified above.

6 Simulations

The results so far clarify what can and cannot be guaranteed in theory. We now ask how often these impossibilities bind in random environments, and how frequently the mechanisms nevertheless deliver Pareto-efficient outcomes. To answer this, we evaluate SJBC+ in simulations and benchmark it against DA and the leading consent-based mechanism, EADA. To calibrate the ex-ante consent environment, we use observed behavior from laboratory experiments: [Cerrone et al. \[2024\]](#) document that around 50% of subjects voluntarily consent to waive priorities (52%

to be precise). We therefore take this one-half rate as our baseline for generating consent sets in EADA.¹⁶

We generate random school choice problems with $n \in \{50, 100\}$ students and n schools, each with unit capacity. For each market size, we construct 2,000 independent instances under two preference structures: (i) independent preferences (students and schools rank uniformly at random), and (ii) correlated preferences, where students exhibit positive correlation in how they rank schools, reflecting realistic settings where school quality is commonly valued.¹⁷ For each instance, we compute outcomes under DA (baseline), SJBC+, EADA with 50% consent, and EADA with full consent. We measure average student rank, the number of beneficiaries relative to DA, and the frequency with which the outcome is Pareto-efficient or justifiable.

Table 2 reveals three robust results. First, SJBC+ attains Pareto-efficiency in the majority of instances (63–85%). This indicates that, in these random environments, the incompatibility between justifiability and full Pareto-efficiency (Proposition 1) is often not binding. Second, SJBC+ compares favorably to EADA under partial consent. When only half of students consent, EADA’s Pareto-efficiency rate is less than 10% in all cases, and exactly zero in correlated markets, where SJBC+ continues to return Pareto-efficient outcomes in over 70% of instances. SJBC+ yields more beneficiaries on average than EADA with full consent, at a modest cost in average rank.

And third, SJBC+ is justifiable in every instance by construction, whereas EADA’s justifiability rate deteriorates sharply with market size and preference correlation. Under full consent, it falls from 27.3% in iid markets with $n = 50$ to 0.3% in correlated markets with $n = 100$. The comparison between the two EADA variants is also intuitive. Full consent allows more trades and therefore delivers larger welfare gains, but it also creates more opportunities to violate the priorities of improvable

¹⁶The replication code is available at <https://www.josueortega.com>. See also Freer et al. [2026] for further experimental evidence on EADA’s performance.

¹⁷Specifically, we draw a common latent value $q_s \sim \mathcal{N}(0, 1)$ for each school. Each student i ranks schools according to utilities $u_{is} = \rho q_s + \sqrt{1 - \rho^2} \varepsilon_{is}$, where $\varepsilon_{is} \sim \mathcal{N}(0, 1)$ are i.i.d. shocks. Preferences are obtained by sorting these utilities. This specification yields positively correlated rankings across students, with correlation increasing in ρ (here $\rho = 0.50$).

Table 2. SJBC+ and EADA’s performance in random markets

	iid		correlated	
	$n = 50$	$n = 100$	$n = 50$	$n = 100$
Average rank				
DA	4.2 (0.023)	4.9 (0.025)	10.4 (0.052)	18.0 (0.088)
EADA (full)	2.6 (0.007)	2.7 (0.005)	5.3 (0.018)	6.8 (0.019)
EADA (50%)	3.3 (0.016)	3.6 (0.015)	8.2 (0.046)	12.7 (0.072)
SJBC+	2.7 (0.009)	2.9 (0.008)	5.8 (0.023)	8.0 (0.029)
Beneficiaries				
EADA (full)	19.8 (0.172)	47.5 (0.275)	32.7 (0.131)	78.0 (0.165)
EADA (50%)	10.6 (0.179)	27.1 (0.320)	13.6 (0.200)	36.2 (0.371)
SJBC+	22.0 (0.240)	55.6 (0.452)	38.1 (0.217)	89.9 (0.253)
PE rate (%)				
EADA (full)	100 (0.0)	100 (0.0)	100 (0.0)	100 (0.0)
EADA (50%)	7.9 (0.6)	0.8 (0.2)	0.0 (0.0)	0.0 (0.0)
SJBC+	66.9 (1.1)	62.6 (1.1)	70.6 (1.0)	85.2 (0.8)
Justifiable rate (%)				
EADA (full)	27.3 (1.0)	3.3 (0.4)	2.2 (0.3)	0.3 (0.1)
EADA (50%)	36.1 (1.1)	10.4 (0.7)	20.9 (0.9)	3.5 (0.4)
SJBC+	100 (0.0)	100 (0.0)	100 (0.0)	100 (0.0)

Notes: Results from 2,000 replications per condition. The first two columns use iid preferences and priorities. The last two columns use correlated preferences ($\rho = 0.50$) with iid priorities. Standard errors in parentheses. Lower rank indicates higher welfare. Beneficiaries are the average number of students who strictly prefer the mechanism outcome to DA. “50%” means the consent set has size $\lfloor n/2 \rfloor$. The PE rate reports the fraction of instances in which the outcome is Pareto-efficient with respect to students’ preferences. The justifiable rate reports the fraction of instances in which the mechanism outcome is justifiable.

students who do not themselves benefit. With only 50% consent, EADA is more constrained: it improves fewer students and is much less often Pareto-efficient, but precisely for that reason it is more often justifiable. In our simulations, partial consent dampens both the gains and the harms of EADA.

7 Conclusion

We introduced justifiability as a framework for determining which improvements over DA can be implemented through priority waivers, and showed that it enables Pareto improvements that are unattainable under existing consent-based approaches. We proposed an algorithm that selects the unique maximal strongly justifiable matching, and then built on it to obtain further justifiable improvements satisfying a weaker, yet meaningful, constrained efficiency guarantee. Simulations indicate that, despite the inherent tension between justifiability and full Pareto-efficiency, our proposed algorithm frequently achieves Pareto-efficient outcomes while restricting priority violations to the justifiable class, and it improves the DA assignments of substantially more students than leading consent-based mechanisms.

An interesting open question that goes beyond the aim of our paper is the analysis of JBC and SJBC+'s performance in equilibrium. Both mechanisms are manipulable and thus their theoretical properties we documented may be affected by strategic behaviour.¹⁸ We leave a formal analysis of manipulation under JBC and SJBC+ for future work.

Acknowledgments

We are grateful to Bettina Klaus and Juan Sebastián Pereyra for their detailed feedback, and to audiences at various conferences and workshops for their comments.

¹⁸Every mechanism that Pareto-dominates DA is manipulable [Abdulkadiroğlu and Sönmez, 2003]. On the positive side, every mechanism that Pareto-dominates DA is not *obviously* manipulable, which means every potential manipulation is risky [Trojan and Morrill, 2020].

References

- Atila Abdulkadiroğlu and Tayfun Sönmez. 2003. School choice: A mechanism design approach. *American Economic Review* 93, 3 (2003), 729–747. [4](#), [32](#)
- Mustafa Oğuz Afacan, Umut Dur, A Arda Gitmez, and Özgür Yılmaz. 2025. Improving the deferred acceptance with minimal compromise. *Games and Economic Behavior* (2025). [5](#)
- Samson Alva and Vikram Manjunath. 2019. *Stable-dominating rules*. Technical Report. Working paper, University of Ottawa. [5](#)
- Keisuke Bando. 2014. On the existence of a strictly strong Nash equilibrium under the student-optimal deferred acceptance algorithm. *Games and Economic Behavior* 87 (2014), 269–287. [4](#)
- Claudia Cerrone, Yoan Hermstrüwer, and Onur Kesten. 2024. School choice with consent: An experiment. *The Economic Journal* (01 2024). [1](#), [29](#)
- Yiqiu Chen and Markus Möller. 2023. Regret-free truth-telling in school choice with consent. *Theoretical Economics* (2023). [4](#)
- Thomas H. Cormen, Charles E. Leiserson, Ronald L. Rivest, and Clifford Stein. 2009. *Introduction to Algorithms* (3 ed.). MIT Press. [17](#), [22](#)
- Benoit Decerf, Guillaume Haeringer, and Martin Van der Linden. 2024. Incontestable Assignments. *arXiv preprint arXiv:2401.03598* (2024). [5](#)
- Battal Doğan and Lars Ehlers. 2021. Minimally unstable Pareto improvements over deferred acceptance. *Theoretical Economics* 16, 4 (2021), 1249–1279. [5](#)
- Battal Doğan and M Bumin Yenmez. 2020. Consistent Pareto improvement over the student-optimal stable mechanism. *Economic Theory Bulletin* 8 (2020), 125–137. [5](#)
- Umut Dur, A Arda Gitmez, and Özgür Yılmaz. 2019. School choice under partial fairness. *Theoretical Economics* 14, 4 (2019), 1309–1346. [4](#), [7](#)
- Lars Ehlers and Thayer Morrill. 2020. (II) legal assignments in school choice. *The Review of Economic Studies* 87, 4 (2020), 1837–1875. [4](#), [25](#)
- Aytek Erdil and Haluk Ergin. 2008. What’s the matter with tie-breaking? Improving efficiency in school choice. *American Economic Review* 98, 3 (2008), 669–89. [5](#)
- Mikhail Freer, Thilo Klein, and Josué Ortega. 2026. Experimental School Choice with Parents. *arXiv preprint arXiv:2603.24615* (2026). [30](#)
- Onur Kesten. 2010. School choice with consent. *The Quarterly Journal of Economics* 125, 3 (2010), 1297–1348. [4](#), [24](#)
- Onur Kesten and Morimitsu Kurino. 2019. Strategy-proof improvements upon deferred acceptance: A maximal domain for possibility. *Games and Economic Behavior* 117 (2019), 120–143. [5](#)
- Minoru Kitahara and Yasunori Okumura. 2021. Improving efficiency in school choice under partial priorities. *International Journal of Game Theory* 50, 4 (2021), 971–987. [5](#)

- Taylor Knipe and Josué Ortega. 2026. The Trade-off Between Minimal Instability and Larger Improvements over Deferred Acceptance. *arXiv preprint arXiv:2504.12871* (2026). [24](#)
- Hyukjun Kwon and Ran I Shorrer. 2020. Justified-envy-minimal efficient mechanisms for priority-based matching. *Available at SSRN 3495266* (2020). [5](#)
- Josué Ortega and Thilo Klein. 2023. The cost of strategy-proofness in school choice. *Games and Economic Behavior* 141 (2023), 515–528. [5](#)
- Josué Ortega, Gabriel Ziegler, R Pablo Arribillaga, and Geng Zhao. 2024. Identifying and Quantifying (Un) Improvable Students. *arXiv preprint arXiv:2407.19831* (2024). [5](#), [7](#)
- Josué Ortega, Gabriel Ziegler, R Pablo Arribillaga, and Geng Zhao. 2026. What Pareto-Efficiency Adjustments Cannot Fix. *Journal of Political Economy: Microeconomics* (2026). [5](#)
- Philip J Reny. 2022. Efficient matching in the school choice problem. *American Economic Review* 112, 6 (2022), 2025–43. [4](#), [25](#)
- Ryo Shirakawa. 2025. Simple Manipulations in School Choice Mechanisms. *American Economic Journal: Microeconomics* (2025). [4](#)
- Qianfeng Tang and Jingsheng Yu. 2014. A new perspective on Kesten’s school choice with consent idea. *Journal of Economic Theory* 154 (2014), 543–561. [4](#), [7](#), [26](#)
- Qianfeng Tang and Yongchao Zhang. 2021. Weak stability and Pareto efficiency in school choice. *Economic Theory* 71 (2021), 533–552. [4](#)
- Peter Troyan, David Delacrétaz, and Andrew Kloosterman. 2020. Essentially stable matchings. *Games and Economic Behavior* 120 (2020), 370–390. [4](#), [25](#), [46](#)
- Peter Troyan and Thayer Morrill. 2020. Obvious manipulations. *Journal of Economic Theory* 185 (2020), 104970. [4](#), [32](#)

A DA and EADA Descriptions

This appendix describes the Deferred Acceptance (DA) algorithm and the Efficiency-Adjusted Deferred Acceptance (EADA) algorithm.

A.1 Deferred Acceptance (DA) Algorithm

The student-proposing Deferred Acceptance algorithm proceeds as follows:

- (1) **Round 1:** Each student applies to her most preferred school. Each school tentatively accepts the highest-priority applicants up to its capacity and rejects the rest.
- (2) **Round $k \geq 2$:** Each rejected student applies to her next most preferred school. Each school considers its current tentative acceptances together with new applicants, tentatively accepts the highest-priority students up to its capacity, and rejects the rest.
- (3) The algorithm terminates when no student is rejected. Tentative assignments become final.

The outcome is the student-optimal stable matching $DA(P)$.

A.2 Efficiency-Adjusted Deferred Acceptance (EADA) Algorithm

The EADA algorithm improves upon the DA outcome by allowing students to consent to waive priorities that do not affect their own assignments. Let $W \subseteq I$ be the set of consenting students.

- (1) **Round 0:** Run the DA algorithm with original preferences. Let μ^0 be the outcome.
- (2) **Round $r \geq 1$:**
 - Examine the DA execution from Round $r - 1$.
 - Find the last step where a consenting student $i \in W$ is rejected from a school s for which i is an interrupter (i.e., i was tentatively placed at s earlier, and at least one other student was rejected from s while i was tentatively placed there).
 - Identify all such interrupting pairs (i, s) with $i \in W$.
 - If none exist, stop and output the current matching.
 - Otherwise, for each such (i, s) , remove school s from i 's preference list (keeping the relative order of other schools unchanged).
 - Rerun DA with the updated preferences.
- (3) The algorithm terminates in finite time. The final matching is $EADA(P, W)$.

The EADA mechanism satisfies the following properties:

- $EADA(P, W)$ weakly Pareto dominates $DA(P)$ for any W .
- No non-consenting student's priority is ever violated.
- A consenting student is never harmed by consenting: $EADA_i(P, W \cup \{i\}) \succeq_i EADA_i(P, W)$ for all i .
- If all students consent ($W = I$), the outcome is Pareto-efficient.

B Detailed Execution of EADA in Example 1

This appendix reports a detailed execution of Kesten’s EADA procedure on our running Example 1. We first compute $DA(P)$ and then illustrate the successive EADA iterations under (i) full consent and (ii) a restricted consent set.

B.1 Computing $DA(P)$ (Iteration 0)

The DA proposals evolve as follows.

Iteration 0							
round	s_1	s_2	s_3	s_4	s_5	s_6	s_7
r1	2		5	<u>6,7</u>	4	<u>1,3</u>	
r2	2		5	<u>1,7</u>	4	<u>3,6</u>	
r3	2	1	<u>3,5</u>	7	4	6	
r4	2	1	3	7	4	<u>5,6</u>	
r5	2	1	3	<u>5,7</u>	4	6	
r6	<u>2,5</u>	1	3	7	4	6	
r7	5	<u>1,2</u>	3	7	4	6	
r8	5	2	<u>1,3</u>	7	4	6	
r9	5	2	3	7	<u>1,4</u>	6	
r10	<u>1,5</u>	2	3	7	4	6	
r11	1	2	3	7	<u>4,5</u>	6	
r12	1	2	3	<u>4,7</u>	5	6	
r13	1	2	3	4	5	6	7

B.2 EADA under full consent

Fix a consent set W that contains every potential interrupter (equivalently, assume that whenever EADA queries an interrupter, the interrupter consents). Kesten’s procedure identifies the last interrupter in W at the DA outcome, asks that student to waive priority at the relevant school, removes that school from the student’s preference list, and reruns DA on the modified problem.

In Iteration 0 the last interrupter is i_7 at s_4 . Since $i_7 \in W$, we delete s_4 from i_7 ’s preference list and rerun DA.

Iteration 1

round	s_1	s_2	s_3	s_4	s_5	s_6	s_7
r1	2		5	6	4	<u>1</u> ,3	7
r2	2		5	1, <u>6</u>	4	3	7
r3	2		5	1	4	<u>3</u> ,6	7
r4	2		3, <u>5</u>	1	4	6	7
r5	2		3	1	4	<u>5</u> ,6	7
r6	2		3	1, <u>5</u>	4	6	7
r7	<u>2</u> ,5		3	1	4	6	7
r8	5	2	3	1	4	6	7

At this modified DA outcome the last interrupter is i_3 at s_6 . We therefore delete s_6 from i_3 's preference list and rerun DA.

Iteration 2

round	s_1	s_2	s_3	s_4	s_5	s_6	s_7
r1	2		3, <u>5</u>	6	4	1	7
r2	2		3	6	4	<u>1</u> ,5	7
r3	2		3	1, <u>6</u>	4	5	7
r4	2		3	1	4	<u>5</u> ,6	7
r5	2		3	1, <u>5</u>	4	6	7
r6	<u>2</u> ,5		3	1	4	6	7
r7	5	2	3	1	4	6	7

Now the last interrupter is i_5 at s_6 . Deleting s_6 from i_5 's preference list and rerunning DA yields:

Iteration 3

round	s_1	s_2	s_3	s_4	s_5	s_6	s_7
r1	2		3, <u>5</u>	6	4	1	7
r2	2		3	<u>5</u> ,6	4	1	7
r3	<u>2</u> ,5		3	6	4	1	7
r4	5	2	3	6	4	1	7

This outcome implements the cycle

$$(i_1 \rightarrow i_6 \rightarrow i_4 \rightarrow i_5 \rightarrow i_1),$$

which yields a Pareto-efficient allocation.

B.3 EADA under partial consent

We now illustrate the failure of EADA to achieve Pareto-efficiency under a restricted consent set that we use in the proof of Theorem 4. Let $W = \{i_1, i_5, i_7\}$.

EADA starts from $DA(P)$ (Iteration 0) and identifies the last interrupter at that outcome, which is again i_7 at s_4 . Since $i_7 \in W$, we delete s_4 from i_7 's preference list and rerun DA, obtaining Iteration 1 above.

At the resulting DA outcome, the (unique) last interrupter is i_3 at s_6 , but $i_3 \notin W$. By the definition of EADA with an exogenous consent set, the procedure stops at this point. The resulting allocation corresponds to the cycle

$$(i_1 \rightarrow i_4 \rightarrow i_5 \rightarrow i_1),$$

highlighted in bold in the table below.

i_1	i_2	i_3	i_4	i_5	i_6	i_7
s_6	s_1	s_6	s_5	s_3	s_4	s_4
s_4	s_2	s_3	s_4	s_6	s_6	s_7
s_2	\cdot	\cdot	\cdot	s_4	\cdot	\cdot
s_3	\cdot	\cdot	\cdot	s_1	\cdot	\cdot
s_5	\cdot	\cdot	\cdot	s_5	\cdot	\cdot
s_1	\cdot	\cdot	\cdot	\cdot	\cdot	\cdot

This allocation is not Pareto-efficient: for instance, i_1 (assigned s_4) and i_6 (assigned s_6) can exchange assignments and both strictly gain.

Nevertheless, under the same consent set $W = \{i_1, i_5, i_7\}$ there exists a Pareto-efficient matching that weakly Pareto-dominates $DA(P)$ and respects the priorities of every student in $I \setminus W$, obtained by implementing the two disjoint cycles

$$(i_1 \rightarrow i_2 \rightarrow i_1) \quad \text{and} \quad (i_3 \rightarrow i_6 \rightarrow i_4 \rightarrow i_5 \rightarrow i_3).$$

Hence $\mathcal{M}(P, W) \neq \emptyset$ while $EADA(P, W)$ is not Pareto-efficient, establishing the claim used in Proposition 2.

B.4 Proof of Theorem 3, Part 2

Finally, we show that the allocation corresponding to the cycle packing

$$\{(i_1 \rightarrow i_2 \rightarrow i_1), (i_3 \rightarrow i_6 \rightarrow i_4 \rightarrow i_5 \rightarrow i_3)\}$$

cannot be obtained by EADA under any consent set, to finalize the proof of Theorem 3. Figure 4 enumerates all possible EADA outcomes, starting from the last interrupter, who is asked if she consents to waive her priorities or not.

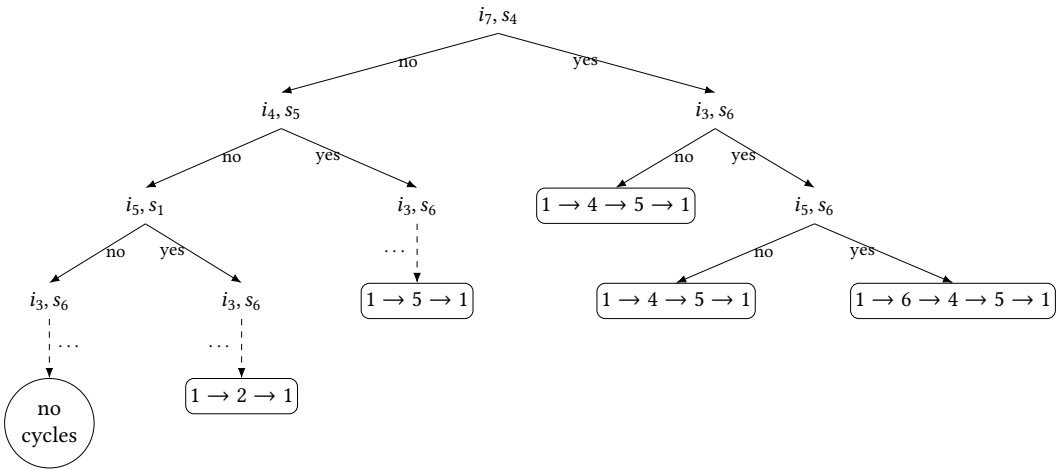


Fig. 4. EADA outcomes for all consent sets in Example 1.

C Why SJBC Requires the “+” Step

This appendix explains why the final step in SJBC+ is necessary. In particular, the naive procedure that iterates Top-Priority Cycles (JBC) sequentially—hereafter Sequential JBC—need not return a Pareto-efficient justifiable matching, even when such a matching exists. The issue is that an early locally-justifiable trade can change the residual instance in a way that prevents reaching a Pareto-efficient outcome within the class of justifiable improvements over the DA benchmark. SJBC+ addresses this by adding a final Pareto-improvement step that preserves the set of beneficiaries.

Consider the following problem with five students and five schools, each of unit capacity. Preferences are listed in the left block and priorities in the right block; the DA assignment is indicated in bold.

EXAMPLE 3 (PREFERENCES AND PRIORITIES, ONE-TO-ONE).

i_1	i_2	i_3	i_4	i_5	s_1	s_2	s_3	s_4	s_5
s_5	s_4	s_2	s_5	s_5	i_1	i_5	i_1	i_5	i_3
s_3	s_5	s_4	s_2	s_1	i_2	i_4	\cdot	i_4	i_1
\cdot	s_2	s_1	\cdot	s_4	i_5	i_3	\cdot	i_1	i_4
\cdot	s_1	s_5	\cdot	\cdot	i_4	i_2	\cdot	i_3	i_5
\cdot	\cdot	\cdot	\cdot	\cdot	i_3	i_1	\cdot	i_2	i_2

Student i_1 is unimprovable at the DA outcome (no other student envies i_1), so Sequential JBC deletes i_1 together with the school assigned to i_1 under DA (here, s_3), and proceeds with the reduced instance on $\{i_2, i_3, i_4, i_5\}$. The resulting envy digraph is depicted in Figure 5.




Fig. 5. Envy digraph for Example 3 after deleting i_1 and s_3 .

In this reduced digraph, JBC selects the \emptyset -labelled 2-cycle between i_3 and i_4 , recommending that i_3 and i_4 exchange schools. This yields the intermediate matching

$$\mu' = (i_1 \rightarrow s_3, i_2 \rightarrow s_1, i_3 \rightarrow s_2, i_4 \rightarrow s_5, i_5 \rightarrow s_4).$$

At this point, Sequential JBC continues by selecting a maximum packing of admissible cycles in the original envy digraph that is consistent with the set of beneficiaries already created (here, i_3 and i_4). One feasible packing is

$$(i_4 \xrightarrow{\emptyset} i_3 \xrightarrow{\emptyset} i_4), \quad (i_2 \xrightarrow{i_3} i_5 \xrightarrow{\emptyset} i_2),$$

which remains justifiable. The resulting allocation is Pareto-efficient; for instance, it coincides with the outcome of serial dictatorship under the order i_2, i_3, i_4, i_5, i_1 .

However, Sequential JBC can instead select the (also admissible and justifiable) cycle

$$(i_2 \xrightarrow{i_3} i_4 \xrightarrow{\emptyset} i_3 \xrightarrow{\emptyset} i_5 \xrightarrow{\emptyset} i_2),$$

which induces the corresponding allocation shown below:

i_1	i_2	i_3	i_4	i_5
s_5	s_4	s_2	s_5	s_5
s_3	s_5	s_4	s_2	s_1
\cdot	s_2	s_1	\cdot	s_4
\cdot	s_1	s_5	\cdot	\cdot
\cdot	\cdot	\cdot	\cdot	\cdot

This outcome is not Pareto-efficient: students i_2 and i_3 can exchange assignments and both strictly gain.

The failure arises from selection among admissible cycle packings: Sequential JBC can terminate at a justifiable matching that is Pareto-dominated by another justifiable matching (and, crucially, without changing the set of beneficiaries). This motivates the final “+” step in SJBC+: after Sequential JBC determines the set of beneficiaries, SJBC+ performs an additional Pareto-improvement stage (restricted to justifiable trades that preserve that beneficiary set), ensuring constrained efficiency within the class of justifiable improvements over DA.

D Justifiability Does Not Imply Preserving the JBC Beneficiary Set

This appendix provides an instance in which there exists a justifiable matching that Pareto-improves on $DA(P)$, but whose beneficiary set is not nested with that of $JBC(P)$. In particular, there is a justifiable improvement that benefits a student who does not benefit under $JBC(P)$, while $JBC(P)$ benefits students who do not benefit under that justifiable improvement. Hence, justifiability alone does not force a reform to preserve the JBC beneficiary set.

Consider the following one-to-one instance with $I = \{i_1, \dots, i_6\}$ and $S = \{s_1, \dots, s_6\}$. Columns on the left list students' preferences (top to bottom), and columns on the right list schools' priorities (top to bottom). Bold entries on the left indicate DA assignments $DA(P)$, and bold entries on the right indicate the students assigned to each school under DA.

EXAMPLE 4 (PREFERENCES AND PRIORITIES).											
i_1	i_2	i_3	i_4	i_5	i_6	s_1	s_2	s_3	s_4	s_5	s_6
s_4	s_1	s_5	s_6	s_4	s_1	i_5	i_3	i_4	i_2	i_1	i_6
s_6	s_4	s_6	s_2	s_1	s_2	i_4	i_2	\cdot	i_4	i_3	i_3
s_5	s_6	s_4	s_4	s_6	s_6	i_3	i_4	\cdot	i_5	i_2	i_1
s_3	s_2	s_2	s_5	s_5	s_5	i_2	i_5	\cdot	i_6	i_4	i_2
s_1	s_3	s_1	s_1	s_2	s_4	i_6	i_6	\cdot	i_3	i_6	i_5
s_2	s_5	s_3	s_3	s_3	s_3	i_1	i_1	\cdot	i_1	i_5	i_4

The DA outcome is

$$\mu^0 = DA(P) : (i_1 \rightarrow s_5, i_2 \rightarrow s_4, i_3 \rightarrow s_2, i_4 \rightarrow s_3, i_5 \rightarrow s_1, i_6 \rightarrow s_6).$$

Applying JBC yields

$$JBC(P) : (i_1 \rightarrow s_5, i_2 \rightarrow s_1, i_3 \rightarrow s_6, i_4 \rightarrow s_3, i_5 \rightarrow s_4, i_6 \rightarrow s_2),$$

The corresponding beneficiary set is

$$\mathcal{B}(JBC(P)) = \{i_2, i_3, i_5, i_6\}.$$

Consider the following DA-improving matching:

$$\mu : (i_1 \rightarrow s_6, i_2 \rightarrow s_4, i_3 \rightarrow s_5, i_4 \rightarrow s_3, i_5 \rightarrow s_1, i_6 \rightarrow s_2).$$

Its beneficiary set is

$$\mathcal{B}(\mu) = \{i_1, i_3, i_6\}.$$

Therefore the beneficiary sets are not nested:

$$\mathcal{B}(\mu) \setminus \mathcal{B}(\text{JBC}(P)) = \{i_1\} \neq \emptyset, \quad \mathcal{B}(\text{JBC}(P)) \setminus \mathcal{B}(\mu) = \{i_2, i_5\} \neq \emptyset.$$

In particular, μ benefits a student (i_1) who does not benefit under JBC, while JBC benefits students (i_2, i_5) who do not benefit under μ .

This example shows that justifiability alone does not force a reform to preserve the JBC beneficiary set, motivating the additional requirement of preserving that beneficiary set when evaluating constrained efficiency (as implemented by SJBC+).

E The Efficiency Guarantee of SJBC+

The constrained-efficiency property in Theorem 2 is tight: there exist problems in which a justifiable matching with a strictly larger beneficiary set Pareto-dominates an SJBC+ outcome. Consider the school choice problem in Example 5 (similar to the one used in the proof of Theorem 1).

EXAMPLE 5 (PREFERENCES AND PRIORITIES).

i_1	i_2	i_3	i_4	i_5	i_6	s_1	s_2	s_3	s_4	s_5	s_6
s_5	s_6	s_1	s_2	s_4	s_2	i_1	i_2	\cdot	i_4	i_5	i_6
s_4	s_1	s_2	s_4	s_5	s_4	i_3	i_4	\cdot	i_3	i_1	i_2
s_1	s_2	s_6			s_1	i_2	i_3	\cdot	i_1		
		s_4			s_6	i_5	i_6		i_6		
		s_3				i_6			i_5		

The DA matching appears in bold. The JBC mechanism finds a three-node cycle:

$$i_1 \xrightarrow{\emptyset} i_4 \xrightarrow{\emptyset} i_2 \xrightarrow{\emptyset} i_1$$

Note that SJBC+ cannot expand to include i_6 , which would require violating i_5 's priority at s_1 , nor expand to include i_5 , which would violate i_6 's priority at s_4 . However, consider the cycle:

$$i_1 \xrightarrow{\emptyset} i_5 \xrightarrow{i_1, i_6} i_4 \xrightarrow{\emptyset} i_2 \xrightarrow{\emptyset} i_6 \xrightarrow{i_2, i_5} i_1$$

which includes all improvable students, and thus is justifiable. Furthermore, note that it Pareto-dominates the matching generated by JBC. Although one could strengthen SJBC+ to close this gap, it remains an open question whether doing so necessarily sacrifices the polynomial-time computation of SJBC+.

F Essential Stability and Justifiability

In the main text we stated the existing result that EADA with full consent is essentially stable, which means that any claim starting from a priority violation is vacuous, i.e. ends up with the claimant losing her claim [Trojan et al., 2020]. This directly implies that essentially stable matchings need not be justifiable by Theorem 3. Now we show that justifiable and Pareto-efficient matchings need not be essentially stable by constructing a non-vacuous reassignment chain starting from a priority violation in Example 1.

Consider the unique Pareto-efficient and justifiable matching obtained with the following trading cycles:

$$\{(i_1 \xrightarrow{\emptyset} i_2 \xrightarrow{i_5} i_1), (i_3 \xrightarrow{\emptyset} i_6 \xrightarrow{i_1} i_4 \xrightarrow{\emptyset} i_5 \xrightarrow{\emptyset} i_3)\}$$

Let us start with someone who suffers a priority violation making a claim. Say i_1 at s_4 . This claim displaces i_6 , who in turn claims s_6 . This last claim displaces i_3 , who in turn claims s_3 , displacing i_5 . In turn, i_5 claims s_1 , displacing i_2 , who now can go to the empty school s_2 . We conclude that the claim is not vacuous, which in turn implies that the unique justifiable and Pareto-efficient matching need not be essentially stable.