

Centralized Admissions for Engineering Colleges in India*

Surender Baswana[†] Partha Pratim Chakrabarti[‡] Sharat Chandran[§]
Yashodhan Kanoria[¶] Utkarsh Patange[¶]

Abstract

We designed and implemented a new joint seat allocation process for undergraduate admissions to over 500 programs spread across 80 technical universities in India, including the prestigious Indian Institutes of Technology (IITs). Our process is based on the well-known Deferred Acceptance algorithm; however, complex affirmative action seat reservations led us to make a number of algorithmic innovations, including (1) a carefully constructed heuristic for incorporating nonnested common quotas that span multiple programs, (2) a method to utilize unfilled reserved seats with no modifications to the core software, and (3) a robust approach to reducing variability in the number of reserved-category candidates admitted, while retaining fairness. Our new seat allocation process went into production in 2015; based on its success, including a significant and provable reduction in vacancies, it has since remained in successful use and has been improved continually.

Keywords: stable matching; college admission; deferred acceptance; affirmative action; algorithm; implementation; market design

1 Introduction

Among the most select universities in the world, the prestigious Indian Institutes of Technology (IITs) are considered the Ivy League of India. The schools have an admission rate of less than 1% for the 1.2 million annual applicants who, in many cases, have spent a significant amount of money on specialized coaching to gain admission. Despite the IITs being so coveted, the puzzling and frustrating reality until recently was that about 6% of the available seats at the IITs were consistently unfilled. Over nearly five years, we have worked to correct this problem via innovative changes to the seat allocation process.

One key reason for seats remaining vacant was as follows. From the 1960s to 2014, admissions to the IITs were conducted under one umbrella. Only slightly less sought after than the IITs are the non-IIT Centrally Funded Technical Institutes (CFTIs). The admissions to the non-IIT

*Published in *INFORMS Journal of Applied Analytics* Vol 49, No. 5. (Special issue for the Daniel H. Wagner Prize for Excellence in Operations Research Practice 2018 finalists) pp. 338-354.

[†]Indian Institute of Technology Kanpur, Kanpur, India. Email: sbaswana@cse.iitk.ac.in

[‡]Indian Institute of Technology Kharagpur, Kharagpur, India. Email: ppchak@cse.iitkgp.ernet.in

[§]Indian Institute of Technology Bombay, Mumbai, India. Email: sharat@cse.iitb.ac.in

[¶]Columbia Business School, New York, New York, USA. Email: ykanoria@columbia.edu, up2139@columbia.edu

CFTIs (henceforth referred to as the non-IITs) were conducted under a separate umbrella, after completion of the IIT admissions. Each candidate was eligible to apply for a seat in each of the two sets of institutes, and several hundred candidates would indeed receive two offers, one at an IIT, and later, another one at a non-IIT. Each such candidate could use at most one of the seats, leaving a vacancy in the other seat; this would not be noticed until much later, in many cases after classes had begun. Such seats would either remain vacant or would be reallocated in an unregulated, decentralized manner after classes began, leading to inefficiencies in seat allocation (e.g., unnecessary vacancies and/or unfair seat allocation). For example, a particular non-IIT seat could be offered to a candidate B, despite having earlier denied the seat to a candidate A, who ranks higher but who had meanwhile taken an IIT seat and was no longer participating in the non-IIT process.

In 2015, we designed and implemented a new combined seat allocation process based on the Deferred Acceptance (DA) algorithm (Gale and Shapley 1962). The process brings all the over 80 CFTIs (IITs and non-IITs) under one umbrella for admissions, with approximately 34,000 available seats and over 1.3 million applicants. Each candidate submits a single preference list over all available programs and can be awarded no more than a single seat from within the system, based on her submitted preferences and rank in each relevant merit list, a ranking of candidates constructed from their performance in the nationwide Joint Entrance Examination (JEE); see our detailed report (Baswana et al. 2019, Appendix 10.8) for more information. We use the DA algorithm to allot the seats because it produces a stable matching. In our context, stability corresponds to a natural notion of *fairness*—a candidate should not be denied admittance to a given program if another lower-ranked candidate, according to the merit list pertaining to that program, was offered a seat in the program.

Despite the theoretical benefit of a combined process in terms of allocating each candidate only one seat, merging these seat allocation systems introduces several challenges. The key challenges are: (1) there is no longer a single ranking of candidates, (2) the process must incorporate complex rules regarding multiple types of seat reservations for affirmative action (more than half of the seats are reserved for affirmative action), (3) speculatively admitting more students than the capacity permits in anticipation of attrition is not allowed, (4) despite its complexities, the process must be completely transparent (unlike many other college admissions mechanisms worldwide). As of this writing, our new joint seat allocation process, which addresses all these challenges, has now been running successfully for four years (2015 to 2018) and has *provably* reduced vacancies at the IITs by nearly 75%. We have continued to improve the process over the years; for example, we have additionally reduced inefficiencies resulting from several thousand students who surrendered seats they had previously accepted.

1.1 Analytical and Algorithmic Innovations

Our complex problem required a number of innovations that might be useful elsewhere. Below, we highlight three novel algorithms we developed; we describe them in detail in Section 3.

- **A practical heuristic for nonnested common quotas.** The defense services (DS) category reservation reserves a few seats at each IIT for eligible candidates (the DS category applies to candidates who are children of military personnel killed or disabled in service); each institute has multiple programs and the reservation applies to the cumulative number of DS candidates across programs. Such a situation is typical in admissions settings; for example, each institute may have multiple programs, but it can award a limited number of scholarships across all programs. We provide a novel practical heuristic for incorporating such a “nonnested common quota” when the size of this quota is small relative to the total number of seats. Theoretically, it is NP-complete to even determine if there is a stable matching when there is such a quota. The idea is to *first run the DA algorithm with temporary “phantom” spots for the nonnested common quota, and to then eliminate these phantom spots by running rejection chains one by one*. Below, in the last part of Section 3.2, we explain why the heuristic should fail rarely. We found no failures in practice and only one failure in 50 synthetically generated test cases. Failure is handled gracefully by creating an extra seat.

- **Dereservation with no software modification.** We were asked to make unfilled reserved seats available to all candidates, but without adding algorithmic complexity to minimize the risk of errors. We show how to implement such “dereservation” without making any change to a software package that lacks this dereservation feature by *iteratively rerunning the software with adjusted seat capacities in the input*. Our approach allows us to seamlessly incorporate such process modifications. We argue that the number of software reruns needed is small; only four iterations suffice in practice.

- **Making the number of reserved-category admissions predictable.** In 2018, we were asked to design a transparent method for female reservations to ensure a target increased fraction of female admissions in each program at the IITs and several non-IITs, and to also ensure (1) the number of seats available to nonfemales remained nearly unchanged, and (2) female candidates did not face a higher bar than nonfemales in any program (i.e., fairness). We provide an innovative algorithmic approach to achieve the demanded objectives. The idea is *to divide seats into reserved and unreserved seats (and not two sets of reserved seats, to retain fairness) and to consider a reserved-category student first for a reserved seat and then for an unreserved seat (and not vice versa, to minimize variability)*. The algorithm we designed is superior to half a dozen algorithmic alternatives. Our careful design and prior simulation experiments had given us confidence to proceed with the female reservations and indeed we found that our expectations were closely met in practice in 2018; by design, our process achieved transparency, fairness, and the target minimum number of female admissions in each program. Moreover, we observed the number of seats allotted to nonfemales was within 0.1% of the number during the previous year at the IITs, far outperforming the next-best algorithmic alternative.

It would be interesting to theoretically formalize our insights corresponding to the first and third bullet above; however, such exercises are outside the scope of the current project.

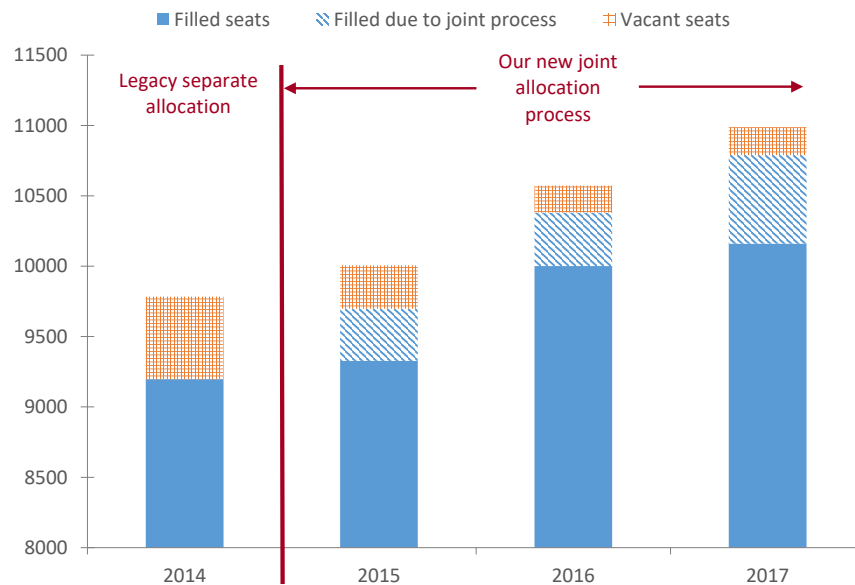
Although some of us had prior expertise in the design of matching markets, many of our practical learnings on this project were learned with difficulty. Throughout the paper, we highlight

our possibly generalizable learnings about designing centralized admissions processes as “Design Insights,” which we accompany with supporting analysis or facts as applicable. We note Design Insights 2 to 4 are novel algorithmic insights corresponding to the list above, and Design Insights 1 and 7 are novel process innovations. Design Insight 1 describes our novel multiround seat allocation process that efficiently fills seats although we were not allowed to speculatively admit more students than the capacity permitted, and thousands of students rejected allotted seats. Design Insight 7 describes our “mock seat allocation,” which allowed candidates to learn how to properly fill their preferences.

1.2 Provable Impact

We analyze the impact of the new seat allocation process in Section 5. The introduction of a combined process in 2015 resulted in a *reduction in vacancies* when classes began: these reductions were over 50% at the IITs (the baseline was 587 vacancies in 9,784 seats in 2014, as Table 2 shows) and nearly 8% at the non-IITs (the baseline was 5,596 vacancies in about 21,285 seats in 2014, as Table 3 shows). Additional significant reductions in vacancies (by over 70% at the IITs) followed in 2016 when candidates who had previously accepted a seat but no longer wanted it were—as per our recommendation—permitted to withdraw, allowing such seats to be assigned to other candidates before classes began. Figure 1 shows the reduction in vacancies at the IITs.

Figure 1. The graph shows vacancies in IITs before and after the implementation of our joint seat allocation process in 2015



Notes: Seats that would not have been filled under the legacy process (estimated via a counterfactual experiment) are shown separately for 2015 onwards. The option to Withdraw after accepting a seat was introduced in 2016, leading to additional reductions in vacancies. In this graph, the Y-axis begins at 8,000.

How do we know our process was the cause of the reduction in vacancies? We rigorously

quantified the benefits of the new process by conducting a careful counterfactual experiment, which we describe in Section 4. Based on the preferences filled by candidates in 2015, we simulate the disjoint allocation process that was in place up until 2014. We first allot candidates to the IITs, and then to the non-IITs. Candidates who receive a better non-IIT seat vacate their IIT seat in the counterfactual. Figure 1 displays the counterfactual-based estimated reduction in the IIT vacancies in each year, resulting from the new joint seat allocation process. For example, in 2017, the IITs had only 198 vacancies under our process; in contrast, they would have had 629 additional vacancies (827 in total) under the legacy process. The new process also allowed many candidates (3,672 candidates in 2017) to be admitted to more preferred programs than the legacy process would have.

An additional benefit of the combined process was simplification of logistics for both colleges and students. Previously, the IIT admissions ran until late July and admissions to the non-IITs would happen only after that, delaying the start of classes and continuing even after classes began. Now the non-IIT admissions are conducted simultaneously with those of the IITs, allowing the non-IITs to begin classes in late July. We observe that, in addition to improving the allocative efficiency, centralization can also greatly reduce the logistical burden of participation on both sides of the market.

Related Work

As of now, many matching markets have been successfully centralized with a clearinghouse that collects preferences and determines matches. Most of these clearinghouses use versions of the DA algorithm proposed by Gale and Shapley (1962), which produces a stable matching. In fact, stability has been found to be essential to the success of such clearinghouses (Kagel and Roth 2000).

One of the earliest documentations of this type of clearinghouse design was for the National Residency Matching Program (NRMP) (Roth and Peranson 1999). A number of cities in the United States use such a DA-based clearinghouse for admissions to public schools; examples include New York (Abdulkadiroglu et al. 2005), Boston, Denver, Washington DC, New Orleans and Chicago. School admissions in Hungary (Biró 2008) and in some other countries, are done similarly using DA. Needless to say, there are hundreds of such marketplaces in addition to the ones we mention above. Our setting differs from the NRMP (and is more similar to many school admission systems) in that the preferences of programs over candidates are determined entirely by exam scores; therefore, there is no strategic behavior on the part of the programs. However, indifferences in program rankings of students are rare by design and do not play a significant role in our setting (in contrast to many school admissions settings).

Following the seminal paper by Gale and Shapley (Gale and Shapley 1962), a vast amount of theoretical literature has been developed on the topic of stable matching. Instead of attempting to survey this literature, we highlight a small subset of papers that are relevant to our project. Dubins and Freedman (1981) showed that candidate-proposing DA cannot be manipulated by candidates, in addition to being candidate-optimal, making it natural for us to appeal to candidate-proposing

DA. However, consistent with previous empirical evidence starting with the NRMP, as well as theoretical work (most recently Ashlagi et al. 2015), we find the set of stable matchings is small; we find that the candidate-optimal stable match is identical to the program-optimal stable match in all the data sets we checked; thus, we would achieve the same allocation if we instead used program-proposing DA. Finally, a key issue we had to address was that of complex business rules including a variety of quotas. Quotas that are nested preserve the existence of stable matchings and allow a stable matching to be computed via a modification of DA. This helped us with incorporating most of the key quotas. However, a problematic DS quota was *not nested*; in general this can lead to nonexistence as well as computational difficulties (Biró et al. 2010). We were able to construct a heuristic approach that almost always finds a stable matching despite this quota; however, it made our algorithm much more complicated. We remark here that many of the serious issues we faced went beyond algorithm design and involved understanding the marketplace and then engineering all its aspects to make it work correctly (Roth 2002).

1.3 Paper Outline

We provide some history and background regarding the CFTIs and their admissions processes in Baswana et al. (2019, Appendix 10.8). Next, we describe the business rules (Section 2), followed by our algorithm design (Section 3). Section 4 describes our implementation, and Section 5 the impact of our new joint seat allocation process. In Section 6, we propose concrete and simple solutions for the large number of vacancies at the non-IITs. We conclude with a discussion in Section 7.

2 Business Rules

Prior to 2014, because of the complexity of the Indian affirmative action program and the goal of a completely fair and transparent process, the 16 (now 23) IITs had an intricate set of business rules for the allocation of their approximately 10,000 seats based on the rankings in the merit lists constructed using the candidates' performance in the JEE Advanced examination. Independently, the non-IITs, including the 30 (now 31) National Institutes of Technology (NITs) and the 12 (now 23) Indian Institutes of Information Technology (IIITs), had their own intricate business rules for allocation of their over 20,000 seats, based on a distinct set of merit lists, which were constructed using the JEE Main and high school graduation examination results. Our task of organizing the joint seat allocation process in 2015 necessitated extensive coordination, starting with the business rules. Below, we provide a summary of key aspects of the final joint seat allocation authority (JoSAA) business rules in 2015 and key changes we made since then; the business rules for the current year are available online (JoSAA 2018)

1. *Merit lists.* Distinct merit lists are constructed for (i) admissions to the IITs using the JEE Advanced scores, and (ii) admissions to the non-IITs using the JEE Main-based scores.
2. *No overbooking.* Admitting more candidates than the program capacity allows in anticipation of some offers being rejected is not permitted.

3. *Fairness.* The seat allocation produced must satisfy the property that if a candidate is denied admission to a particular program, then no other candidate with an inferior rank in the relevant merit list should be admitted to that program. That is, the allocation must be consistent with a *cutoff* rank for each program in the relevant merit list. These cutoff ranks are publicly announced.

4. *Reservation of seats for different affirmative action categories.*

- (i) By law, in each program, 15% of the seats are reserved for the scheduled castes (SCs), 7.5% for the scheduled tribes (STs), and 27% for other backward castes (OBCs). The remaining 50.5% of seats are in the open category; that is, they are available for all.
- (ii) In each category, 3% (now 5%) of the seats are reserved for persons with disabilities (PwD).
- (iii) Two seats are allocated preferentially in each IIT under the DS category.
- (iv) Beginning in 2018, 14% of admissions to the IITs and the NITs should be of female candidates; if necessary, this can be accomplished via creation of additional “supernumerary” seats, so that nonfemale candidates do not have to compete for a smaller number of seats. This number will gradually increase to 20% in the coming years. At present, only 8 to 9% of IIT undergraduate students are female, reflecting (a) a gender skew at the top of the merit lists for admission, and (b) that only one of three female candidates who is offered an IIT seat accepts it, whereas two of three male candidates who are offered an IIT seat accept it. However, once they are in attendance at the IITs, women do better academically than men on average.
- (v) Half of the seats at each NIT are reserved for candidates from the corresponding state.
- (vi) A candidate should be considered for admission to a program through all categories in which she is eligible, in a particular order described in the business rules.

5. *Dereservation of unfilled reserved seats.* Unfilled OBC seats in a program must be dereserved (i.e., made available to open candidates). SC/ST seats cannot be dereserved to the open category. However, unfilled SC-PwD seats must be made available to SC candidates; similarly for other seats reserved for PwD candidates (e.g., OBC-PwD dereserved to OBC).

6. *Multiple rounds of allocation.* Candidates fill and lock their preferences over programs at the start of the allocation process, after which the “first round” allocation is produced. To fill allotted seats that were rejected (either actively or via a no-show by the candidate at the reporting center), additional new allocations (candidates who are allotted for the first time in a round) and allocation “upgrades” (candidates who received a more preferred seat compared to their previous round allocation) are executed over multiple rounds. Candidates who accept a seat are provided the following options: Float, Slide, and Freeze. Float indicates the candidate wants to be upgraded as high as possible on her preference list. Slide indicates the candidate wants to remain in the institute to which she has currently been allotted but wants the most desirable program available at that institute. Freeze indicates she wants to remain assigned to the specific program to which she has been allotted, even if other options become available subsequently. Candidates are not permitted to change their preference list midway through the process. Candidates who accept a seat at an IIT are not permitted to apply to an IIT in the future. However, accepting a seat in either an IIT or a non-IIT does not prevent a candidate from applying to a non-IIT in the future.

Our Float/Slide/Freeze mechanism may be useful in other multi-round allocation settings. We observed that the candidates valued this flexibility. For example, after the first round in 2017, 14.2% of the candidates chose to freeze their allotted program, 8.8% chose the slide option, and 59.6% chose the float option; the proportion of candidates who chose the freeze and slide options increased in later rounds. The remainder (17.4%) decided to exit the system by rejecting their seats.

Design Insight 1 *In a multi-round seat allocation process, a candidate who is allotted a seat may be given the following options: Float (i.e., the candidate is open to upgrades in future rounds), Slide (i.e., the candidate is open to upgrades but only within the same institute), Freeze (i.e., the candidate elects to keep the current allocation), or Reject (i.e. the candidate rejects the seat offered).*

One additional difficult business rule was that if the category of a candidate changes because that candidate reported incorrect information, the candidate must be penalized by being allocated a seat only from those remaining unfilled after the previous round.

The complexities of these rules necessitated a number of algorithmic and process-related innovations. We describe our algorithmic innovations in Section 3, and then describe our process implementation in Section 4.

3 Algorithm Design

Early in the creation of a joint seat allocation process, the authorities suggested on several occasions that we should keep the processes for the IITs and the non-IITs separate but should require candidates who received two offers to reject one and iterate. Even if implemented in the best possible way, such an approach is identical to program-proposing DA, but with iterations occurring in the real world instead of on a computer. Our simulations showed as many as six iterations would have been needed to obtain convergence. We ultimately convinced all concerned to instead collect integrated preferences from candidates and then run the algorithm on a computer.

Our software, which we describe in Section 4.1 below, was created by the National Informatics Centre (NIC). NIC initially suggested that we should collect candidate preferences as a single list over all IIT and non-IIT programs and then run an algorithm similar to program-proposing DA on the computer. We convinced them to use candidate-proposing DA on the preferences instead, since it produces a candidate-optimal allocation, and is incentive compatible for candidates (i.e., candidates have nothing to gain by misreporting their preferences). Incidentally, it turned out that the two versions of DA produced identical allocations on the preferences collected from 2015 to 2017. This is consistent with the finding in the literature that in typical matching markets, there is an essentially unique stable matching (e.g., see Ashlagi et al. (2015)).

The complex business rules necessitated a handcrafted variant of candidate-proposing DA. Our full algorithm is specified in an 75-page document that includes pseudocode, explanations, and examples (Baswana et al. 2019). In the remainder of this section, we summarize how our algorithm incorporated, in turn, (1) affirmative action reservations, (2) nonnested common quota reservations

(for DS candidates), (3) dereservation of unfilled seats, and (4) the target of 14% female candidates in each program (in 2018), all while satisfying stringent practical requirements. Although our solution for affirmative action reservations is straightforward, we constructed novel approaches to incorporate the other listed requirements. Because these types of requirements arise naturally in planning school and college admissions, our algorithmic design may help other market designers faced with similar problems choose the best practical solution.

3.1 Affirmative Action Reservations

Birth-category affirmative action reservations (Rules 4(i) and 4(ii) in Section 2) are implemented using their nested structure by dividing each program into multiple “virtual” programs, one for each category. Their capacity is set to the number of reserved seats for that category in that program. The preference lists of candidates are modified to use virtual programs instead of academic programs. The sequence in which a candidate applies to virtual programs is based on business rules. For example, a candidate who appears in the Open, OBC, Open-PwD and OBC-PwD merit lists applies to virtual programs in the order $\text{Open} \rightarrow \text{Open-PwD} \rightarrow \text{OBC} \rightarrow \text{OBC-PwD}$.

3.2 Nonnested Common Quota Reservations

One of the most difficult business rules to implement was the nonnested common quota for DS candidates, Rule 4(iii): at most two DS candidates are admitted in total across all programs in an institute. This quota is nonnested; as such, a stable matching may not exist and checking whether a stable matching exists is computationally hard (Biró et al. 2010). We now present our practical solution to the challenge of incorporating a small nonnested common quota; such a quota may arise naturally in admissions settings; for example, each institute may have multiple programs, and a limited number of scholarships it can award across all programs. In principle, to solve this problem, one could appeal to the integer programming method (Biró and McBride 2014), which finds a stable outcome when it exists. However, such an approach was untenable in practice because of its complexity, relative opaqueness, and the likelihood of an unreasonably long run time on our large problem.

Our Heuristic Algorithm

To implement the DS reservation, a new virtual DS program with two seats is added *per institute* (not per program) (e.g., IITK-DS for IIT Kanpur). An equivalent way to view our setting is to think of one DS virtual program for each program, such that all DS virtual programs in a particular institute have a common quota of two seats, and that individual DS virtual programs have no additional capacity constraints of their own. Under this view, the DS reservation is a nonnested common quota (Biró et al. 2010). However, for purposes of describing our algorithm, we define a single DS virtual program for the entire institute. Only DS candidates are eligible for these virtual programs. Furthermore, the preference list of each DS candidate is modified as follows. If the preference list of a candidate is $\langle p_1, p_2, p_3 \rangle$, then his preference list is first augmented as $\langle p_1,$

Institute(p_1)-DS, p_2 , Institute(p_2)-DS, ...). Then p_1, p_2, \dots are each replaced by multiple virtual programs as we described in Section 3.1. Note that any seat allotted based on the DS status of a candidate must only be from the open category.

We start by running the DA algorithm to completion, while artificially allowing the institute DS virtual programs to admit up to two candidates, over and above the capacity of any other virtual program. Upon completion of the DA algorithm, if each candidate allotted to a DS virtual program is given a seat in the respective open-category virtual program she had requested, we may have artificially increased the capacity of some (open-category) programs by up to two seats per institute. To prevent this overage, we sequentially process candidates in DS virtual programs as follows. Suppose Candidate c , allotted the IITB-DS virtual program, asks for a seat in, for example, the IITB-EE program. Let x be the candidate with the worst rank among those currently in the virtual program IITB-EE-Open. Then, the IITB-EE-Open program rejects Candidate x . We run the candidate-proposing DA starting from the current allocation and with Candidate x applying to her next most preferred virtual program, thus, triggering a *rejection chain*. This rejection chain could possibly loop back to cause the rejection of Candidate c who started it; in this case, we would fail to obtain a fair allocation. In any such situation, our algorithm rolls back the rejection chain and allocates that DS candidate a supernumerary seat (i.e., a seat in excess of program capacity). We provide an example of such a “failure” in the appendix. Unsurprisingly, however, we did not observe a single failure in practice to date, and in our 50 synthetically generated test cases, we encountered only one instance of failure. For a more detailed discussion of this algorithm, we refer an interested reader to Baswana et al. (2019, Appendix 10.7).

Why Failures Are Rare

The example of failure in our handling of nonnested common quota, which we describe in the appendix, not only demonstrates a case of failure, but also throws light on why such failures are rare as long as the number of nonnested common quota (DS) seats and the number of candidates eligible for that quota are small (there are tens or fewer of each in our setting). For a failure to occur, one of the rejection chains initiated by a candidate occupying a seat in a DS virtual program must displace a DS candidate from an open virtual program; that is, at least one step of the rejection chain must involve an open virtual program such that its worst-ranked open candidate is eligible for the DS quota. In our setting, which has fairly sizeable open virtual programs, few DS candidates, and relatively short rejection chains, this is already unlikely to happen. In addition, if our heuristic wisely chooses the order in which to process the DS candidates, the rejection chain needs to loop back to the DS virtual program from which it started. This occurs in all likelihood because the displaced DS candidate herself applies to that program, and this is also unlikely given the sizeable number of DS virtual programs (one per institute, with many institutes). In practice, we did not use smart (re)ordering of DS candidates, because we prioritized algorithmic simplicity over the small risk of creating an extra seat due to failure. Still, no extra seats were created in practice.

It would be an interesting theoretical exercise to formalize this intuition in the future, along the lines of the analysis of the heuristic for accommodating couples in the NRMP; see Kojima et al. (2013) and Ashlagi et al. (2014). The notion of influence tree (roughly, the programs and candidates that may be affected by the rejection chain of a DS candidate) in the latter paper may be especially useful.

Design Insight 2 *Nonnested common quotas that include a relatively small number of seats can be accommodated using our simple heuristic, while creating very few (or no) extra seats. The idea of our heuristic is to first run DA with temporary extra spots for the nonnested common quota, and then to eliminate these extra spots by sequentially running rejection chains.*

3.3 Dereservation of Unfilled Seats

Business rule 5 required unfilled OBC seats to be made available to open-category candidates. The approach we initially suggested involved construction of augmented merit lists, thus making open-category candidates eligible for OBC seats but at a lower priority than all OBC candidates, and modification of virtual preference lists so that general candidates now apply for both the open and the OBC virtual programs. We showed that running our algorithm on these modified inputs would produce the candidate-optimal allocation satisfying the business rules. However, the authorities feared this approach could cause issues with correctly computing the closing ranks (i.e., the rank of the worst ranking candidate from each merit list that was allotted to a program; these have to be publicly declared after each round as per our transparency mandate) (see Design Insight 6) or might have some other hidden problem. An authority running centralized college or school admissions is typically loathe to modify, add complexity to, or replace software that has been tried and tested; we also see this in Chilean college admissions (Rios et al. 2019) and NYC public school admissions (Abdulkadiroglu et al. 2005). Upon reflection, we realized that if we relaxed our goal of candidate optimality, and were willing to tolerate a slightly longer computation time, we could use the existing software as a black box, and yet incorporate dereservations.

Our Algorithm

Our approach was remarkably simple: (1) Run the core algorithm with no dereservations to completion, (2) move the vacant seat capacity in each OBC virtual program to the corresponding open virtual program, (3) rerun the core algorithm, and (4) iterate until convergence.

Properties of Our Algorithm

Each successive run of our core algorithm (which is essentially candidate-proposing DA) will produce (weakly) fewer assignments to each OBC virtual program. This is true because adding more seat capacity to one or more virtual programs only makes candidates (weakly) better off (Roth and Sotomayor 1990); therefore, some candidates may be granted an upgrade *out* of an OBC virtual program, but no candidate may be granted an upgrade *to* an OBC virtual program that already had

vacancies, given that the previous allocation was stable. As a result of this monotonicity, we can conclude that only a finite number of reruns is needed. In fact, each of our iterations resembles the iterations of Manjunath and Turhan (2016), who find fast convergence when there are two parallel school systems drawing upon the same set of candidates. In our setting, convergence is even faster; we needed only about four iterations. Our intuition for the observed rapid convergence is that accommodating several open candidates in vacant OBC seats lowers the bar for open-category candidates; however, this allows only a few OBC candidates to upgrade because most of the OBC candidates are still given their allocation via OBC seats. As such, very few additional vacancies are generated in OBC seats.

Our implementation of dereservation can be viewed as a hybrid of program-proposing DA (in the outer loop of multiple runs) and candidate-proposing DA (in the core algorithm), for the economy consisting of the virtual preferences and augmented merit lists, which we define above in the first paragraph of this subsection. Therefore, the allocation we arrive at is stable, but not necessarily candidate optimal in theory.

Design Insight 3 *Practitioners are loath to modify or replace existing software. Dereservation of unfilled seats, and conceptually similar problems, such as the integration of two separate admission systems, can be implemented practically by treating the existing software as a black box, and running it repeatedly while iteratively updating the input provided until convergence.*

3.4 A Target Minimum Fraction of Females in Each Program

In 2018, we were asked to ensure that 14% of admitted candidates in each program in the IITs and the NITs are female. In 2017 and in prior years, gender played no role in admissions, and the fraction of admitted females was only 8-9% overall in the IITs. The fraction was approximately 14% on average at the NITs, but lower in some programs and higher in others. In addition, we had to satisfy the following constraints:

1. The number of nonfemales admitted should not increase compared to the number admitted in 2017 because the institutes have resource constraints and cannot simultaneously accommodate an increase in both nonfemale and female candidates.
2. Nonfemales should not be disadvantaged while admitting the target minimum fraction of females in a program. Therefore, the number of nonfemales admitted should not decrease significantly compared to those admitted in 2017.
3. The program capacities must be frozen prior to the joint seat allocation because they are publicly announced; moreover, they should be determined in a simple, transparent, and fair way. The allocation produced should not violate the announced capacities.
4. The allocative approach should ensure that the number of females is at least the larger of (i) the typical number of females in the past, and (ii) 14% of the total capacity of the program, but not substantially more than this target.

5. For each program, the admission cutoff for female candidates should not be more stringent than that for nonfemale candidates, to ensure female candidates are not unfairly disadvantaged.

The following simple algorithm may initially appear to meet all the above constraints: Divide each virtual program into two separate virtual programs for nonfemales and females, respectively. In the seat allocation process, nonfemales compete for nonfemale virtual programs and females compete only for female virtual programs. As per Requirement 3, the capacities of these virtual programs may be fixed beforehand as follows—the capacity of the nonfemale virtual program is set equal to the number of nonfemales admitted to that virtual program in the previous year (thus satisfying Requirements 1 and 2), whereas the capacity of the female virtual program is chosen to satisfy Requirement 4 under the assumption that the applicant pool will be similar to that of the previous year. Although this algorithm meets the first four constraints, it violates the fifth constraint; that is, some female candidates may be denied a seat under this algorithm despite nonfemales with inferior ranks being admitted. A violation of the fifth constraint would defeat the purpose of this exercise, because in pursuit of guaranteeing at least 14% of the seats to females, we would deny some female candidates of their right to compete for seats on the basis of merit. This serious problem in the algorithm was not merely a theoretical possibility; our simulations on 2017 preference data revealed this algorithm would deny seats to several deserving female candidates. We now present our algorithm, which rectifies this problem and which we used in 2018.

Our Algorithm

The algorithm (and corresponding business rule) divides each virtual program p into two separate virtual programs as follows:

- Female(p): exclusively for females
- Gender-neutral(p): admits candidates based only on merit

A Nonfemale(p) virtual program is intentionally avoided, to satisfy Constraint 5.

The key feature of our algorithm is that each female candidate interested in a program first competes for a seat in the relevant Female(p) virtual program(s); only if she fails to get a seat does she compete for a seat in the Gender-neutral(p) virtual programs. This is in contrast to Business Rule 4 for other reservations. (In Boston school admissions, a similar precedence-order design for walk-zone reservations ended up with an allocation almost identical to what would have happened with no reservation because the number of reserved seats chosen was too small; see Dur et al. (2018). In our implementation, the fraction of seats allocated to females is guaranteed to increase from 9% to at least 14%, and the chosen precedence order will ensure it does not substantially exceed 14%). As a result, if the cutoff for Female(p) ends up being less stringent than the one for Gender-neutral(p), then females do not occupy any seat in Gender-neutral(p) and hence, Constraints 1 and 2 are satisfied exactly. Conversely, if the cutoff for Female(p) is the more stringent one, then notice that the algorithm allows female candidates to compete for Gender-neutral(p). This ensures the fulfillment of Constraint 5 and hence fairness to female candidates. One may be concerned

Table 1: The table summarizes the performance of our design for female reservation in 2018

	IITs	NITs
Female seats	1,852	2,947
Gender-neutral seats	10,227	15,673
#Gender-neutral seats taken by females	10	467
Average excess over target minimum #females	0.6%	15.8%
Average reduction in seats for nonfemales	0.1%	3.0%

Note. The reduction in seats for nonfemales is minimal.

this outcome is unfair to nonfemale candidates because female candidates might evict nonfemale candidates from their seats in p . However, in any scenario in which even one female occupies a seat in $\text{Gender-neutral}(p)$, we observe the allocation of all seats in p , including both $\text{Gender-neutral}(p)$ and $\text{Female}(p)$, is based purely on merit with no regard to gender. Therefore, neither females nor nonfemales are unfairly disadvantaged in our algorithm.

Note that the alternative of having females apply first to the gender-neutral virtual program would lead to more variability in the fraction of females admitted, because now the *top* female candidates would compete for gender-neutral seats and an unpredictable number would gain admission via a gender-neutral virtual program (compared to very few with our approach); see the details in the discussion in the next two paragraphs. Thus, the choice of precedence order affects not only the *number* of reserved-category candidates admitted, as noted by Dur et al. (2018), but crucially also affects the *unpredictability/variability* in the number of reserved-category admissions.

Design Insight 4 *If reserved-category candidates are considered first for reserved seats and then for unreserved seats, this choice of precedence order can reduce the variability in the total number of reserved-category candidates admitted.*

Table 1 summarizes the superior performance of our design for female reservations in 2018. Note how only 10 female candidates at the IITs and 467 female candidates at the NITs took gender-neutral seats, thus ensuring that Constraints 4 and 2 were closely met under our approach; by design, Constraints 1, 3, and 5 were already met a priori. In contrast, if female candidates had been considered first for gender-neutral virtual programs, we found that females would have captured 66 gender-neutral seats in the IITs and 847 in the NITs, unevenly and idiosyncratically distributed across programs. This would have caused fewer seats to be available for nonfemales (a violation of Constraint 2), particularly in some programs.

3.5 Additional Comments on Algorithmic Aspects

Needless to say, we had to address various additional complexities. One noteworthy challenge was how to handle the several candidates who were found to have misreported their birth category. These candidates were then supposed to be allotted spots only from unfilled seats. We refer the

interested reader to our full report (Baswana et al. 2019) for our algorithmic approach to carefully handle such candidates.

In addition, it is fairly typical to have several candidates with identical ranks in our setting, and candidates with the same rank may apply to the same program. Our approach was to create additional seats to accommodate all candidates at the cutoff rank, corresponding to the “L-stable score-limit” solution in Biró and Kiselgof (2015), although we learned of that paper only later. We employed the natural modification of DA to accommodate ties in this way.

We advocated for modification of the difficult DS and category change rules to simplify the implementation. Our advice was eventually accepted: Since 2016, the DS seats have been allotted in a supernumerary manner, decoupling DS allotments from other allotments and greatly simplifying the process. In 2018, the category misreporting penalty was also eliminated.

4 Implementation

“Market design involves a responsibility for detail, a need to deal with all of a market’s complications, not just its principal features” (Roth 2002, p. 1341).

The orders of a Delhi High Court judge under public interest court case W.P.(C) 2275/2010 catalyzed the launch of a joint seat allocation process for all CFTIs in 2015; we provide details in Baswana et al. (2019, Appendix 10.8).


The separate processes under the IITs umbrella and the non-IITs umbrella each involved multiple (three or four) rounds of seat allocation conducted in rapid succession. The process began with candidates submitting preferences over programs; a variant of serial dictatorship was then used to produce an allocation in the first round. A candidate who was allotted a seat was asked to pay a fee and accept the allotted seat. Some candidates did not accept their seats, and these seats were then allotted again in the second round, and so on. Multiple rounds were especially important to fill seats because overbooking was not permitted (Rule 2), and the yield was much less than 100% especially in the non-IITs. This multi-round structure had to be retained for joint seat allocation. However, the process in each round now had to account for different merit lists for the IIT programs and the non-IIT programs, necessitating that we construct a suitable multi-round adaptation of DA.

Design Insight 5 *Multiple rounds of centralized allocation can serve to mitigate the number of vacancies without incurring the overage risk associated with admitting more candidates than the capacity of each program permits.*

Figure 2 shows how multiple admission rounds enabled reductions in the number of vacancies using 2015 data (the scenario in 2016 and 2017 was complicated somewhat by the fact that candidates were permitted to withdraw after accepting a seat; see Section 5).

Another important feature of the legacy processes we retained was that the cutoff ranks for each program in each category were published in each round (Rule 3).

Figure 2. In the graphs below, we show vacancy progressions for the IITs and the non-IITs after each admission round in 2015



(a) Vacancies in the IITs (10,006 total seats).

(b) Vacancies in the non-IITs (24,068 total seats).

Notes. The IITs held three rounds, while the non-IITs held a fourth round and a special round in addition. Because candidates were not allowed to withdraw after accepting a seat, many vacancies were discovered after classes began, after which a special round was conducted in the non-IITs.

Design Insight 6 *Conducting admissions based on strict program preferences derived from exam scores allows for the public announcement of closing ranks/cutoffs. Such a public announcement provides (1) transparency, (2) guidance to candidates regarding their chances of admission to each program (because the cutoff ranks from the previous year are available), and (3) help to candidates in understanding they should report their true preferences without fear.*

Regarding (2) in Design Insight 6, we note that there are over 500 available programs. Typically, candidates have not been constrained on the number of preferences they can enter. Indeed, candidates submit extremely long preference lists—of average length about 80 in practice. Constructing such detailed preferences is a Herculean task; therefore, published cutoff ranks are very helpful in this regard.


As a result of (3) in Design Insight 6 (and because of explanatory material provided to candidates), we are not particularly concerned that candidates are misreporting their preferences, even though this has been observed in other centralized matching systems (Hassidim et al. 2017). We have heard very few anecdotes about candidates being unsure whether they should report their true preferences.

4.1 Software Development and Testing

Figure 3 summarizes the timeline of the development of the new process. Two versions of software implementations of our algorithm were developed. One was a database-based version developed by NIC, and the other was a main-memory version we developed at IIT Kanpur. No data set of integrated candidate preferences (over both the IIT and non-IIT programs) was available for 2014 or earlier; therefore, testing teams from three IITs prepared synthetic data sets for software validation. The teams also prepared a set of validation modules to verify that the allocations

generated satisfied all business rules. They tested and validated 50 test cases of varying sizes over two months. The IIT Kanpur software took 30-40 minutes to run on a data set with one million candidates. It took 10 minutes per run and typically three to four runs were needed (recall that multiple runs were used to implement dereservation of unfilled seats; see Section 3.3).

Figure 3. The timeline for developing the new process extended from March 2015 to July 2015




The preference filling started on July 1. Thereafter, every day until the conclusion of the process, two or three snapshots of the candidate preferences were taken and both versions of the software were run on that data and matched with each other. The matched allocations were then passed through validation modules prepared by the testing teams. This ensured there would be no surprises on allocation days (one for each round) when we had to publish the results.

In 2016, the software was modified to accommodate the modified business rules, including the introduction of the Withdraw option. The running time of the software also improved, decreasing from 40 minutes to between 3 and 4 minutes. One of the major problems in 2016 was the late addition of candidates and changes in their marks because of reevaluations. New candidates were added even after the third round of allocation.

4.2 Allocation Process: Timeline and Details

Figure 4. The timeline of the 2017 multiround seat allocation process extended from early April through the end of July



Note: The timeline was similar in 2015 and 2016, except the Special Round occurred much later in 2015, and not at all in 2016

Figure 4 shows the timeline of the process in 2017. Candidates filled in their preferences before the first round. A few days after the preference-filling portal opened, two successive “mock allocation” rounds were conducted online to ensure that candidates filled their preferences correctly and thus obtained their most desired program for which they cleared the cutoff. Analysis of the data reveals the benefits of the mock rounds. In 2017, 2,063 candidates who did not get a seat in the first mock allocation properly updated their preferences and were allotted a seat in a program

in the actual first round. Furthermore, the mock rounds (especially the second mock round) were found to provide fairly accurate guidance in the sense that the closing ranks of the programs were close to those in the actual first round. The median errors in the estimated closing ranks (for the open category) was 13.8% for the IITs (18.4% for the non-IITs) in the first mock round, and only 2.5% for the IITs (6.1% for the non-IITs) in the second mock round.

Design Insight 7 *Conducting a mock allocation based on tentative preferences and revealing the results to candidates helps them understand the purpose of the preference list they are submitting and the allotments they may receive as a result, and hence mitigates misreporting.*

Shortly after the mock allocation rounds, candidates could “lock” their preferences or they could allow the system to automatically lock them. Candidates were not allowed to change their preferences after the deadline. (Allowing a change of preferences in subsequent rounds would conflict with the requirement to have a single closing rank for each program in each category. It may also have produced other kinds of confusion in an environment with high stakes.) After being allotted to a program, the candidates were asked to pay the fee and physically report to a reporting center for document verification and their acceptance of the allotted seat. At the reporting center, each candidate had to choose an option—Freeze, Slide, or Float (see Section 2). If a candidate did not report, it was classified as a reject. In 2016, an additional option, Withdraw, was added to the list. The Withdraw option was added to address candidates who accepted a seat in a past round but did not want it anymore. Such candidates were given a refund of their seat acceptance fee and asked to fill out a short survey regarding their reasons for withdrawing, including the university at which they were planning to study and whether they intended to write JEE again.

5 Impact of Joint Seat Allocation

In this section, we critically analyze the impact of the joint seat allocation during 2015, 2016, and 2017, relative to 2014 when the seat allocation for the IITs and the non-IITs were separate processes. We focus our analysis on seat vacancies.

Table 2 shows the number of vacancies in the IITs in each year during 2014-2017. In 2015, there was a significant reduction in the vacancies in the IITs. This reduction can be attributed to combining the two seat allocation processes. It occurred primarily because in the prior system, the IIT admissions were completed first and candidates who subsequently obtained a preferable allocation at a non-IIT surrendered their IIT seat, which then remained unfilled. We confirmed this benefit to the IITs by running a rigorous counterfactual experiment to estimate the number of vacancies that would have resulted from continuing with two separate allocation processes. We describe this experiment below in Section 5.4. We have omitted IIT Dhanbad from the list in Table 2. It had 7 vacancies in 2014, 33 in 2015, 37 in 2016, and 37 in 2017. The reason for the apparent increase is that this institute filled many vacant seats locally via a “spot” round until 2014 but did not feel the need to fill the few vacant seats from 2015 onward.

In 2016, there was an additional reduction in the vacancies in the IITs as a result of introducing the “Withdraw” option for candidates. The final number of vacancies was reduced by over 70% (relative to 2014) in 2016 and 2017, because most of these vacated seats were reassigned successfully.

Table 3 shows the number of vacancies in the non-IITs during 2014-2017 after the main rounds. A plausible explanation for the reduction in vacancies in 2015 is that until 2014, some candidates would list programs under the non-IIT process they did not truly prefer to the IIT allocation they had already obtained a few days earlier, and later reject such non-IIT programs if allotted. The new process eliminated such waste. In 2016, after the introduction of the Withdraw option, there was an additional reduction in the number of vacancies at the non-IITs (by 22% relative to 2014). However, in 2017, vacancies returned to nearly 2014 levels. We discuss possible reasons later in this section and propose further changes to reduce vacancies in Section 6.

5.1 Joint Seat Allocation: 2015

For the academic programs offered by 87 institutes in the joint seat allocation in 2015, 153,000 candidates listed 85 programs per candidate on average in their preference lists. In addition, more than one million candidates write the JEE Main each year but only about 150,000 of the best-performing candidates qualify to fill in their preferences. The institutes are comprised of 19 IITs (including ISM Dhanbad, which was designated an IIT in 2016), 31 NITs, 18 IIITs, and 18 other-CFTIs. The seat allocation was carried out in four rounds between July 1 and July 21, 2015.

Based on our advice, the non-IITs, after classes began, conducted a systematic centralized special round for the first time; the objective was to fill the nearly 5,697 vacant non-IIT seats. Previously, until 2014, after the final round of admission, the vacant seats in non-IIT institutes were filled locally by each institute by a spot round. In some years, the spot round was centrally organized; however, candidates who had already been given seats could only be given upgrades within the same institute. This led to unfairness and incentive issues in the overall process. Unassigned candidates would physically visit as many institutes as possible (i.e., only a few institutes) as required to participate in their spot rounds. This setup led to misallocation of seats: a seat could be given to a candidate with worse rank than another candidate because only the former could be present physically during that institute’s spot round. In the 2015 special round, unassigned candidates were required to pay fees to participate to reduce the number of frivolous applications. The best possible seat was assigned to each candidate in a fair and efficient manner. Candidates were allowed to submit new preferences for the special round. After the round, 2,148 candidates migrated to a different institute after joining, while 1,492 stayed in the same institute but migrated to a different program. As a result, 5,354 new allotments were made; of these, 2,683 candidates did not report physically to a reporting center, thus leading to a final vacancy count of 2,883 at the non-IITs.

Table 2: The table shows vacancies when classes began at the IITs in 2014 (separate seat allocation), 2015 (joint seat allocation introduced), 2016 (Withdraw option introduced), and 2017

IIT	2014	2015	2016	2017
BHU	197	72	55	51
Bombay	4	2	0	4
Bhubaneshwar	21	18	12	12
Delhi	12	7	2	4
Gandhinagar	10	7	6	8
Guwahati	32	17	6	12
Hyderabad	15	13	0	2
Indore	2	6	4	13
Jodhpur	28	8	2	2
Kanpur	18	10	5	0
Kharagpur	97	34	22	22
Mandi	6	16	1	3
Madras	37	29	19	4
Roorkee	84	27	16	16
Patna	17	12	3	7
Ropar	7	8	6	3
Pallakad		11	7	3
Tirupati		11	4	7
Jammu			7	8
Dharwad			7	7
Goa			4	5
Bhilai			2	5
Total vacancies	587	308	190	198
Vacancies prevented by DA (see Section 5.4)	-	373	379	629
Vacancies in pre-2014 IITs	587	286	159	163
Reduction vs 2014		51%	73%	72%

Notes. There are blanks for years in which that IIT did not exist. At the bottom, we make a before-and-after comparison based on only the IITs that existed in 2014. The IITs had a total of 9,784 seats in 2014 across 17 IITs. This increased to 10,988 seats in 2017 across 23 IITs.

Table 3: The table shows vacancies across the non-IITs when classes began in 2014, 2015, 2016, and 2017

	2014	2015	2016	2017
NITs	3,208	3,209	2,613	3,244
IITs	578	709	666	1,292
Other GFTIs	1,710	1,779	1,622	1,974
Total vacancies	5,596	5,697	4,901	6,510
pre-2014 NITs	3,208	3,111	2,530	3,112
pre-2014 IITs	578	444	347	528
pre-2014 Other GFTIs	1,710	1,632	1,502	1,740
Total vacancies	5,596	5,141	4,379	5,380
Reduction vs 2014		8%	22%	4%

Notes. To ensure a fairer comparison, the bottom set of numbers excludes institutes that were not a part of the system (or did not exist) in 2014. There were 21,285 seats across the non-IITs in 2014; these were across 30 NITs, 12 IITs, and 16 other CFTIs. By 2017, there were 25,220 seats across non-IITs. A large part of the increase in seats was because of adding one NIT, 11 IITs, and four other CFTIs during this period.

5.2 Joint Seat Allocation: 2016

In 2016, 92 institutes participated in the joint seat allocation. This includes 23 IITs, 31 NITs, 20 IITs, and 18 other-CFTIs. The seat allocation was carried out in six rounds. The major change introduced in 2016 was to allow the candidates a Withdraw option: a candidate who had previously accepted a seat allotted by JoSAA could withdraw by appearing at a reporting center before the last round of seat allocations. This option provided candidates with the ability to choose freely between their allotted CFTI seats and outside (non-CFTI) options before the last round: a candidate who withdrew was refunded the seat acceptance fee and was allowed to sit for the JEE in the following year.

Design Insight 8 *A majority of candidates may willingly share they have decided not to accept their allotted seats if they are provided some monetary or other incentive.*

Withdrawal by 3,762 candidates allowed us to reallocate the surrendered seats in the later rounds of allocation, further reducing vacancies; however, significant room for further improvement remained because of the limited efficiency of filling vacant seats in the last rounds, as we describe in Section 6.

In 2016, the non-IIT umbrella did not conduct a centralized special round. We later found that some of these institutes conducted “spot” rounds on their own to fill their vacant seats. We believe a special round should have been conducted as in 2015.

5.3 Joint Seat Allocation: 2017

In 2017, 97 institutes participated in the joint seat allocation process. This includes 23 IITs, 31 NITs, 23 IIITs, and 20 other-CFTIs. The seat allocation was carried out in seven rounds. For the first time, the non-IITs did not give weightage to the board exam (i.e., high school graduation exam) marks in constructing their merit lists. This possibly resulted in more candidates wanting to sit for the JEE again the following year, thus leading to a large increase in the number of withdrawals (i.e., over 6,000), and hence an increase in the final number of vacancies in the non-IITs, almost to the same level as in 2014. Additionally, the increase in withdrawals could have been caused partially by more candidates listing and accepting programs they did not really want, because: (1) candidates who did not receive any allocation during the mock seat allocation were encouraged to list more programs, and (2) awareness regarding the option to withdraw later may have increased (in 2016 the option was introduced for the first time), leading candidates to list and accept programs more liberally. We propose some solutions to this issue below in Section 6.

Happily, a special round for non-IIT institutes was conducted centrally again in 2017 as in 2015. Of the 6,510 vacant seats, 5,352 were allotted in the special round, and 3,830 of these candidates actually joined classes, resulting in a final vacancy count of 2,680 at the non-IITs.

5.4 Counterfactual Experiment to Assess Impact of Joint Seat Allocation

To conclusively establish the reduction in vacancies was due to the superiority of our DA-based joint seat allocation process over the legacy process, we simulated the legacy allotment process used until 2014 using data for each year since 2015. We then compared the resulting allocation with the allocation for that year under our DA-based process.

Until 2014, the IITs conducted their allocation before the non-IITs did. Because IIT seat allocations were frozen by the time non-IIT allotments took place, many candidates abandoned their IIT seats in favor of non-IIT seats. The abandoned IIT seats then remained vacant. To simulate the allocation that would have been produced under the legacy process, we proceeded as follows. We first allotted only IIT seats, using the preferences submitted by the candidates over the IIT programs (ignoring the non-IIT programs they had listed) and the IIT merit lists. Once the IIT seats were allotted, we then considered the submitted preferences of the candidates over non-IIT programs only, after removing their preference entries below their allotted IIT program (if any) and allotted the non-IITs based on the non-IIT merit lists. (This simulation setup was based on the optimistic assumption candidates would list only those programs in the old non-IIT process they preferred to their IIT allocation. If candidates listed additional programs, this would only serve to further increase the number of seats that would have been wasted under the legacy process.) We then compared the final allocation with the first-round allocation obtained under DA.

One can show, mathematically, that *every single candidate obtains a weakly better allocation under the new process than under the legacy process*. The seat allocation produced by the legacy process is identical to the one the new joint process would produce if it acted on a reduced seat matrix, with the reduction in seat counts defined as follows: There is no change for the non-IITs;

however, for the IITs, we reduce the seat counts by the number of candidates who would have abandoned their IIT seats for non-IIT seats under the legacy process. Thus, the new process produced a candidate-optimal allocation with respect to a larger seat matrix relative to the legacy process. It follows from Theorem 2.25 in Roth and Sotomayor (1990) that all candidates are weakly better off under the new process.

We also see this in the simulation results. Our simulation further proves that our DA-based joint process caused a *large reduction in the number of vacancies at the IITs*; our process prevented 373 vacancies in IITs in 2015, 379 in 2016, and as many as 629 in 2017 (see Table 2 and Figure 1). Last but not least, we also found as a result of the new process, 1,890 candidates received a more preferred program in 2015, 1,767 in 2016, and 3,672 in 2017.

6 Vacancies at the Non-IITs: Causes and Proposed Solutions

As we have shown, our joint process has been very successful in reducing vacancies in the IITs, as per our mandate. Meanwhile, vacancies in the non-IITs have decreased only slightly; although this was outside our mandate, it concerns us. The Withdraw option introduced in 2016 produced a modest further reduction in vacancies in non-IITs (see Tables 2 and 3) by reducing the number of vacancies discovered only after classes began. The data reveal the reason the additional reduction was modest: most of the seats from which candidates withdraw are difficult to fill in late rounds. Over 70% of first-time allocations in late rounds were rejected by candidates, and most of the withdrawals occurred in the penultimate round. Almost all of the withdrawals were from the non-IITs; we estimate that, in 2016 and 2017, about 1,816 and 2,918 vacancies, respectively, in the non-IITs were caused by late withdrawals. The rest of the vacancies occurred primarily because an estimated 2,233 and 2,800 candidates in 2016 and 2017, respectively, accepted a seat and did not withdraw, but then did not show up when classes began.

In 2017, there was an unexpected increase in the number of withdrawals. The number of withdrawals was 3,762 in 2016 (of which 1,114 said they planned to write JEE again) but rose to 5,525 in 2017 (of these, 1,684 said they wanted to write JEE again). Over 4,000 of these withdrawals occurred in the penultimate round. Consistent with the pattern in 2016, most of these seats remained vacant at the end of Round 7 (the last round) due to rejection of the majority of first-time allocations. This leads us to the question: How can seats be effectively filled in the main rounds despite the significant number of rejects and withdrawals, which are concentrated in a subset of the programs? Data indicate most of the vacancies at the end of the main rounds are *avoidable* in the sense that there are eligible candidates who want those seats. It is clear that these vacancies can be avoided if we somehow reduce unnecessary withdrawals (caused by allotting candidates to programs they do not really want) and increase the seat acceptance rate in the last round. It would also help to add more rounds to fill up the seats vacated by withdrawals

We now argue for modification of candidate *incentives* to improve the efficiency of the seat allocation process. Currently, there is no bar on writing JEE Main again even if a candidate accepts a seat, does not report for classes, and does not withdraw. In our view, this leniency

may be causing an unnecessarily large number of candidates to accept seats in programs they have no intention of joining. The problem is exacerbated because the non-IITs no longer use high school graduation exam marks to compute their merit lists, causing more candidates to consider the option of writing the JEE again. The policy of refunding the seat acceptance fee to a candidate who withdraws may be amplifying the problem further (now that the awareness of the Withdraw option is increasing). Note that the IITs do not allow candidates who accept a seat in an IIT program (and then do not withdraw) to apply for an IIT seat in the future, whereas there is no analogous rule for the non-IITs. We believe that this asymmetry between IIT and non-IIT programs should be removed. We advocate for the following approach:

1. **Any candidate who is allotted a seat (whether in an IIT or a non-IIT) and then accepts that seat should *not* be permitted to apply for admission to any CFTI in the future.** In 2017, 5,525 candidates withdrew after accepting a seat (almost all from non-IIT seats). In a survey these candidates filled at the time of withdrawal, 1,684 (30.5%) of them indicated that they intended to apply again the following year. With the proposed rule change, these candidates would not have accepted the seat in the first place, enabling us to fill most of these 1,684 seats by the end of the seat allocation process. In contrast, because of the low success rate of reallocating seats freed by last-minute withdrawals, as we discuss in Baswana et al. (2019, Appendix 10.9), the current process filled only about 30% (i.e., 500) of those seats. Thus, this rule change can be expected to save close to 1,100 vacancies.
2. **The Withdraw option can be retained (primarily for candidates who just received better offers from outside the system and are no longer interested in the seat they were allotted), and incentivized (e.g., with a quick return of fees). However, candidates who withdraw should not be allowed to apply for admission to any CFTI in the future, for the reason we state above. (Currently, a candidate who withdraws from an IIT seat is allowed to reapply in future).** Of the candidates who withdrew, 3,197 told us they had received an outside offer of admission and did not plan to apply again the following year. Because most outside admissions offers occur after the start of JoSAA admissions, allowing these candidates to withdraw is essential to filling up these seats. Because their early withdrawal would give us more time (i.e., rounds) to fill these seats, it can be incentivized (e.g., via monetary incentives, a technique that was successful in the special round).
3. **Messaging should *not* encourage candidates to blindly list *more* programs. Instead, candidates should be encouraged to list programs in which they are truly interested and clearly made aware of the consequences of accepting a program they do not intend to join.** This was the only one of our suggestions implemented in 2018 and it resulted in desirable changes in key metrics. (i) The average length of candidate preference lists in 2018 was 17% smaller than in 2016 and 2017; candidates apparently avoided listing some programs in which they were not interested, thus reducing the number

of unnecessary withdrawals. (ii) The downstream consequence in the non-IITs was vacancy reduction of about 6% (before the special round)—from 6,510 in 2017 to 6,133 in 2018.

4. **Stop withdrawals two rounds before the last round instead of one round before the last round.** This change will significantly increase the fraction of seats vacated by the withdrawing candidates, which we will be able to fill from the current approximately 30% by giving us two opportunities to fill them; see our analysis in Baswana et al. (2019, Appendix 10.9). Also, the loss in flexibility to candidates will be minimal since the interval between later rounds is only two to three days anyway
5. **Try to ensure only serious candidates participate in the last two rounds, for example, by using monetary incentives or asking candidates to indicate if they want to participate.** This proposal is inspired by the special rounds, which successfully use monetary incentives (e.g., fee deposit and penalties for rejecting a seat) and require candidates to sign up afresh to ensure their seriousness, and benefit from a higher percentage of seat acceptance. In 2017, approximately 59% of the seats were accepted in the special round compared to about 30% in the last main round.

In Baswana et al. (2019, Appendix 10.9), we make a further suggestion to conservatively admit some additional candidates in excess of program capacity based on somewhat predictable rejections of fresh allocations (the challenge is that many virtual programs are very small, and overage in specific categories can have political ramifications in addition to causing logistical issues). This suggestion can be coupled with the ones listed above to maximize efficiency gains.

Design Insight 9 *The incentive properties (and related messaging) of the overall dynamic process may play a crucial role in determining how participants interact with it, and hence greatly impact its allocative efficiency. Monetary incentives can have an impact. Incentives perceived to affect a candidate's career options may be yet more powerful.*

We advocated to include these changes (with the exception of overbooking) into the business rules for 2018; however, only Suggestion 3 was partly accepted. Our other suggestions were rejected citing, for example, that the proposed penalty for backing out was “too harsh” on candidates. Consequently, we predicted several thousand vacancies when classes began in 2018 as a result of the persistent improper incentive structure, and indeed, 6,133 seats remained vacant in the non-IITs at the end of the main rounds of admission, only slightly less than in 2017. We continue to advocate that our suggested changes be made in future years. Some popular articles covering our recommendations have been published recently (Chhaphia 2018, IFORS News 2019).

7 Discussion

The theory and practice of one-shot seat allocation using the DA algorithm is well developed. Nevertheless, we faced many challenges in bringing it to the high-stakes setting of seat allocation for

the most prestigious engineering colleges in India. Challenges included the complexity of business rules and other requirements such as a dynamic multiround process to fill rejected seats. Since its implementation in 2015, our new joint process has provably reduced the vacancies at the IITs, which previously conducted their admissions independently of, and prior to, those of the non-IIT CFTIs. The reduction in vacancies at the IITs was more than 50% in 2015 and has been in excess of 70% since 2016, relative to 2014. Although vacancies at the non-IITs have been reduced compared to 2014, a significant fraction of seats remain vacant at the end of the multiround admissions process. We are advocating for multiple changes to address this issue; a major goal being to change the incentives of candidates so that fewer seats remain unfilled because of being vacated at a late stage in the admissions process.

Our algorithmic innovations in this project include (1) a practical heuristic for incorporating a nonnested common quota, (2) a method to dereserve seats with no modifications to the core software, and (3) a robust approach to reduce variability in the number of reserved-category candidates admitted, while retaining fairness. Theoretically formalizing our insights regarding the first and the third innovations may be interesting for future work.

Overall, our experience in developing, executing, adapting, and improving this centralized seat allocation process has taught us many practical lessons, including those we highlighted as “Design Insights” and which we support by analysis and statistics throughout the paper. We are optimistic that many practitioners faced with similar problems can benefit from our learnings and that more countries will be inspired to collect the efficiency gains that come from centralizing admissions.

Acknowledgement

We are thankful to all those administrators who provided us their support in the execution of the Joint Seat Allocation process, and permitted us to conduct research on the admissions data. Kanoria was supported by the National Science Foundation grant CMMI-1653477.

References

- Abdulkadiroglu A, Pathak PA, Roth AE (2005) The New York City high school match. *American Economic Review* 95(2):364–367.
- Ashlagi I, Braverman M, Hassidim A (2014) Stability in large matching markets with complementarities. *Operations Research* 62(4):713–732.
- Ashlagi I, Kanoria Y, Leshno JD (2015) Unbalanced random matching markets: The stark effect of competition. *Journal of Political Economy* 25(1):65–98.
- Baswana S, Chakrabarti PP, Kanoria Y, Patange U, Chandran S (2019) Joint seat allocation 2018: An algorithmic perspective. *CoRR* abs/1904.06698, URL <http://arxiv.org/abs/1904.06698>.
- Biró P (2008) Student admissions in Hungary as Gale and Shapley envisaged. Technical report, University of Glasgow, Glasgow, UK.
- Biró P, Fleiner T, Irving RW, Manlove DF (2010) The college admissions problem with lower and common quotas. *Theoretical Computer Science* 411(34-36):3136–3153.

- Biró P, Kiselgof S (2015) College admissions with stable score-limits. *Central European Journal of Operations Research* 23(4):727–741.
- Biró P, McBride I (2014) Integer programming methods for special college admissions problems. Zhang Z, Wu L, Xu W, Du DZ, eds., *Combinatorial Optimization and Applications*, 429–443 (Cham: Springer International Publishing), ISBN 978-3-319-12691-3.
- Chhappia H (2018) Senior IIT faculty propose curbs on withdrawal to prevent seat vacancies. *The Times of India* URL <https://www.pressreader.com/india/the-times-of-india-mumbai-edition/20180713/282071982669973>.
- Dubins LE, Freedman DA (1981) Machiavelli and the Gale-Shapley algorithm. *The American Mathematical Monthly* 88(7):485–494.
- Dur U, Kominers SD, Pathak PA, Sönmez T (2018) Reserve design: Unintended consequences and the demise of Boston’s walk zones. *Journal of Political Economy* 126(6):2457–2479, URL <http://dx.doi.org/10.1086/699974>.
- Gale D, Shapley LS (1962) College admissions and the stability of marriage. *American Mathematical Monthly* 69(1):9–15.
- Hassidim A, Marciano D, Romm A, Shorrer RI (2017) The mechanism is truthful, why aren’t you? *American Economic Review* 107(5):220–24.
- IFORS News (2019) Centralized admissions for engineering colleges in India. *IFORS News, March 2019* 6–8, URL https://issuu.com/ifors/docs/ifors-news-march2019_ad0d0d56b6c74d/6, accessed June 2019.
- JoSAA (2018) Business rules for joint seat allocation for academic programs offered by the IITs, NITs, IEST, IIITs and other-GFTIs. URL <https://josaa.nic.in/webinfocms/Handler/FileHandler.ashx?i=File&i=109&i=109&i=Y>, Accessed March 1, 2019.
- Kagel JH, Roth AE (2000) The dynamics of reorganization in matching markets: A laboratory experiment motivated by a natural experiment. *The Quarterly Journal of Economics* 115(1):201–235.
- Kojima F, Pathak PA, Roth AE (2013) Matching with couples: Stability and incentives in large markets. *The Quarterly Journal of Economics* 128(4):1585–1632.
- Manjunath V, Turhan B (2016) Two school systems, one district: What to do when a unified admissions process is impossible. *Games and Economic Behavior* 95:25–40.
- Rios I, Larroucau T, Parra G, Cominetti R (2019) Improving the Chilean college admissions system. *Operations Research* Forthcoming.
- Roth AE (2002) The economist as engineer: Game theory, experimentation, and computation as tools for design economics. *Econometrica* 70(4):1341–1378.
- Roth AE, Peranson E (1999) The redesign of the matching market for american physicians: Some engineering aspects of economic design. *American Economic Review* 89(4):748–780.
- Roth AE, Sotomayor MAO (1990) *Two-sided Matching: A Study in Game-Theoretic Modeling and Analysis*. Number 18 (Cambridge University Press, Cambridge, UK).

Appendix to Section 3.2: Example of a Rare Failure

In this appendix, we provide an example of a potential pathological situation in our heuristic algorithm to incorporate a nonnested common quota, forcing the creation of a supernumerary seat.

Three IITs are involved in the example—Kanpur (IITK), Delhi (IITD), and Bombay (IITB).

Let Amar, Akbar, Chetan, and Dhanush be four DS candidates. At the end of the DA algorithm, Amar, Akbar, and Chetan are allocated a DS seat in IITK-DS (had requested Mechanical), IITD-DS (had requested Metallurgy), and IITD-DS (had requested Electrical), respectively; but Dhanush is given a seat in IITB-Electrical-Open. Moreover, let Dhanush be the last-ranked candidate who is given a seat in IITB-Electrical-Open; in addition, suppose there are already two (unnamed) DS candidates who are ranked above Dhanush and are occupying the two IITB-DS seats. Let Bharat, Krish, and Ekansh be the last-ranked open-category candidates in IITK-Mechanical, IITD-Metallurgy, and IITD-Electrical, respectively. Figure 5 shows the details of these seven candidates and the programs the DA algorithm has allocated.

Figure 5. The figure shows the interim program allocation in the DA algorithm to four DS and three GE candidates.

Candidate	Open Rank	Program	Seat allocated	DS status
Amar	1200	IITK-Mechanical	DS	Yes
Akbar	400	IITD-Metallurgy	DS	Yes
Chetan	500	IITD-Electrical	DS	Yes
Dhanush	200	IITB-Electrical	GE	Yes

Candidate	Open Rank	Program	Seat allocated	DS status
Bharat	600	IITK-Mechanical	GE	No
Krish	350	IITD-Metallurgy	GE	No
Ekansh	150	IITD-Electrical	GE	No

Preference list of Dhanush
IITB-Computer Sc
IITB-Electrical
IITD-Computer Sc
...

Preference list of Ekansh
...
IITD-Electrical
IITB-Electrical
...

We now describe the processing of Amar, Akbar, and Chetan who were allotted DS seats by our heuristic algorithm. Now, these candidates need to be reallocated to open seats. To accommodate Amar, we need to remove Bharat; and this leads to Bharat getting some other less preferred program, and possibly pushing another candidate out, thus starting a rejection chain. In a similar manner, Akbar is allotted a seat in IITD-Metallurgy after the removal of Krish, and Krish is allotted some other less preferred program, again in a rejection chain.

Next, we process Chetan. Because Chetan was allotted a seat in IITD-Electrical through a DS quota, we need to remove Ekansh from IITD-Electrical. The next most preferred program for Ekansh is IITB-Electrical. Recall that Dhanush is the last-ranked candidate being allotted IITB-Electrical-Open. Notice that although Dhanush is a DS candidate who was allocated an open seat as an open-category candidate in IITB-Electrical. Therefore, because Ekansh received a seat in IITB-Electrical-Open, Dhanush will be removed from IITB-Electrical-Open. Dhanush will also be rejected from IITB-DS (because both spots are already occupied), and will then apply for his next preferred program, which is IITD-Computer Science, a program for which he does not clear the open-category cutoff; therefore, he will be rejected by the IITD-Computer Science-Open virtual

program and will then apply to the IITD-DS virtual program. Two DS candidates, Akbar and Chetan, are already occupying the two spots in IITD-DS. Because Dhanush has a better rank than Akbar, and Akbar has a better rank than Chetan, Dhanush will cause Chetan to be removed from the DS virtual program of IITD, and we see that we have created a loop condition.

At this point, we also realize that Ekansh should never have been rejected from IITD-Electrical because no DS candidate is taking a seat there, and Dhanush should not have been rejected from IITB-Electrical; then we ask why Chetan can't keep his spot in IITD-DS, and so on. We encountered a problem, and in fact, such examples may not have a stable matching at all. In this example, our algorithm gives Chetan a supernumerary seat in the IITK-Electrical program.

We expect failures to be rare when the nonnested quota is small, because multiple improbable events, as captured in this example, must occur to produce a failure; see the discussion in Section 3.2. Indeed, we have not observed any failures in practice to date.