Playing with Matches: Adopting Gale–Shapley for Managing Student Enrollments Beyond CS2

Anna N. Rafferty
arafferty@carleton.edu
Carleton College
Northfield, MN, USA

David Liben-Nowell
dln@carleton.edu
Carleton College
Northfield, MN, USA

David R. Musicant
dmusicant@carleton.edu
Carleton College
Northfield, MN, USA

Emy Farley
e.farley@bowdoin.edu
Bowdoin College
Brunswick, ME, USA

Allie Lyman
alyman@carleton.edu
Carleton College
Northfield, MN, USA

Ann May
amay@carleton.edu
Carleton College
Northfield, MN, USA

ABSTRACT

Enrollment in computer science has increased dramatically in recent years, straining capacities and leading to various strategies for managing enrollment. But some strategies increase student competition and may have disproportionate negative impacts on students from underrepresented groups. We believe success in computing education necessitates a more equitable approach to course enrollment. In this experience report, we describe our new enrollment mechanism, "the Match." Building on the Gale–Shapley stable matching algorithm, the Match was designed to encourage a liberal arts approach to course selection and attempt to broaden participation in computing. Drawing on data from three years of use, we find high student participation, with the vast majority of students having their enrollment preferences met. With Match registration, our courses have tended to be a bit more inclusive of younger students. The Match appears not to have disparate negative impacts like those of competitive enrollment, but has increased workload in the Registrar's Office. Overall, we believe the Match has decreased student and faculty angst around registration, and we argue that systems like the Match can help manage enrollment pressures in ways that are consistent with educational values.

CCS CONCEPTS

• **Social and professional topics → Computing education**.

KEYWORDS

course registration, enrollment management, stable matching

ACM Reference Format:
Anna N. Rafferty, David Liben-Nowell, David R. Musicant, Emy Farley, Allie Lyman, and Ann May. 2024. Playing with Matches: Adopting Gale–Shapley for Managing Student Enrollments Beyond CS2. In *Proceedings of the 55th ACM Technical Symposium on Computer Science Education V. 1 (SIGCSE 2024), March 20–23, 2024, Portland, OR, USA.* ACM, New York, NY, USA, 7 pages. https://doi.org/10.1145/3626252.3630906

SIGCSE 2024, March 20–23, 2024, Portland, OR, USA

ACM ISBN 979-8-4007-0423-9/24/03
https://doi.org/10.1145/3626252.3630906

1 INTRODUCTION

Undergraduate computer science courses and majors have experienced huge growth in enrollment in the past fifteen years, but growth in computer science faculty has not kept pace [9]. Handling enrollment growth has been challenging: strategies have included limiting the number of majors (via lottery, competitive enrollment, or first-come-first-served systems); redirecting student interest; increasing course sizes; and trying to leverage external teaching resources [17, 22]. But these approaches can have unintended consequences and can differentially impact students. Competitive enrollment decreases feelings of belonging and self-efficacy [18]. Larger classes tend to have higher dropout rates [5]. Major-specific admissions policies or grade-threshold restrictions are associated with decreased participation in more lucrative majors by students in underrepresented minority groups [6, 8]. Women and students from underrepresented racial minority groups have often had decreased representation after past enrollment booms [16, 25], and policies enacted to cope with these booms can negatively impact departmental cultures in ways that disproportionately harm students from underrepresented groups [21, 22]. We believe that effective computing education demands more equitable approaches. We thus developed a new registration initiative that was intended to better meet our institutional goals as a liberal arts college, and decrease the effort and anxiety from the allocation process for the scarce seats in CS classes, while avoiding the inequities of some other mechanisms.

At Carleton College, mirroring many institutions, a student can register for courses only after all more-senior students select all of their courses. Our post-CS2 courses often filled to capacity with older students. Younger students were left with no CS options, and many added themselves to a multitude of CS waitlists.

Forcing students to try to enroll solely via waitlists for many terms seemed unwelcoming and potentially particularly alienating to groups of students already underrepresented in computing. Students spoke to us about their registration anxiety and difficulties. Individual faculty attempted to address these issues informally, including by increasing course sizes and some backchannel centralization of waitlist management (to avoid duplication in accommodating shut-out students). These mechanisms were time consuming and failed to address student registration angst or fully mitigate the negative effects of waitlists, leaving students unable to predict which course they might get into (if any) and unable to plan their course selection, inside and outside of CS.

Our ad hoc process was serving no one well, and we wanted to develop a better way of handling these registration challenges, and to ensure that whatever system we developed met both institutional and departmental values. Three goals were especially important:

Goal 1: Make the Tent Bigger**.** *Enable a greater number of interested students to take a class in the CS department.*

As a small liberal arts college, we value and encourage curricular exploration. We strongly believe that all students who want to take CS should be able to do so. We saw any competitive enrollment approach as both inconsistent with our values and likely to narrow participation in computing rather than broaden it. We thus wanted a system that would maximize the set of interested students who could take a CS class, regardless of their major or intended major.

Goal 2: Reduce Bingeing**.** *Encourage students interested in CS, especially majors, to take a diverse range of courses each term.*

A liberal arts worldview champions its students pursuing a broad course of study, both in their selection of their three courses each term and overall in their college careers. But, in one term with large waitlists, we found some students taking *only* CS courses, while waitlists were stuffed with others enrolled in *no* CS courses. Transcripts showed a small minority of students taking far more elective CS courses than required. This suggested that there were not simply insufficient CS seats available, but rather that some students take a disproportionate number of CS courses.

Goal 3: Meet Needs (and Most Desires)**.** *Ensure that students who need a particular course for graduation can enroll, and, when possible, enable students to take the CS classes they most prefer.*

Faculty may be unaware of students' requirements or preferences when managing waitlists, and we worried that some students (and some groups of students) were more likely to advocate for themselves than others. Our third goal encompassed our desire to more equitably accommodate student needs and desires.

The present work. In this experience report, we describe "the Match," the registration system we developed to address our goals, while seeking to decrease time and angst for students and faculty. Our system algorithmically matches students to a single CS course based on their preferences, and is used alongside the existing system, without college-wide changes. We are unaware of other institutions managing enrollment pressures using such an algorithmic approach to meet student desires. We report on students' participation in the initiative, registration patterns before and after implementation, and effects on faculty and registrar workloads. While the impacts on registration are necessarily correlational and come with significant caveats—especially the confounds of COVID and the fact that our system is only three years old—our results suggest that the system has helped us meet some of these goals. We hope our experience offers inspiration for others grappling with enrollment challenges.

2 THE MATCH

We begin by describing the Match in the context of our institution. The extant registration system has two phases: first, each student must meet with their academic advisor; and, second, the students (in priority order, descending by seniority, with a pseudorandom tiebreaker) register for their full slate of three classes. (See Figure 1.)

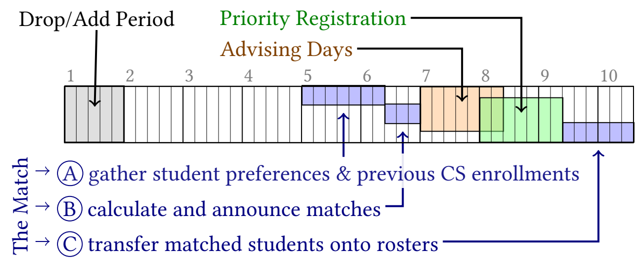

Figure 1: The existing registration system (above) and the Match (below) within the 10-week terms at Carleton.

The Match system allows any individual student to request a reserved seat in a single post-CS2 class. Capacity constraints mean some classes cannot accommodate all interested students; as such, students submit a rank-ordered preference list of CS classes, and a seat is reserved in their most-preferred class that still has space available. (In a minority of cases, every course on a student's preference list is full; in this case, no seat is reserved for them.)

Algorithmically, we allocate seats to students using the "deferred acceptance" mechanism of David Gale and Lloyd Shapley [12], which efficiently computes a *stable* matching. Most famously, the Gale–Shapley algorithm is used, with some modification, for the National Resident Matching Program (NRMP) that matches doctors to hospitals for their post-medical-school residencies [24]; it has also been adapted to other contexts, including public school choice [2, 3, 20], college admissions [23], and military career placement [26]. In public schools, impacts have included improved student welfare and achievement [1]. We believe that the widespread external use of Gale–Shapley was key to gaining institutional approval for the Match: we could merely cite the NRMP to justify the specific mechanism, allowing us to spend our political capital on arguing that there was a problem that needed to be solved.

2.1 Implementation Details

To implement this mechanism, we added three steps to the existing registration process (again, see Figure 1), as follows.

Ⓐ *Gather preferences, exception requests, and past enrollments.* Well before Advising Days, we send an email (and subsequent reminder) to the ≈750–1000 students who have met the prerequisites for any Match-eligible (i.e., post-CS2) course, inviting them to submit their preferences for CS courses via a Google Form. When a submission is made, students choose a subset of courses to rank, and then express their preference order of that subset. We also invite petitions for individual exceptions. (Granted petitions typically enable a student to match to multiple CS courses to stay on pace for graduation.)

The CS department and the Registrar's Office compile a data file that lists each student's class year and all CS courses successfully completed. These data are used to generate the "preferences" that classes have for students: classes "prefer" more senior students, with elective classes breaking ties based on a count of fulfilled CS major requirements, and breaking remaining ties (pseudo)randomly.

Ⓑ *Execute Gale–Shapley and announce results.* Using the submitted student preferences for classes, and the calculated class "preferences" for students, we run the Gale–Shapley deferred-acceptance algorithm to compute an assignment of students to classes, respecting enrollment caps.[1] We configure the algorithm to prioritize students' preferences over classes' (i.e., students propose to classes). This means the best strategy for students is to accurately report their preferences. The implementation disallows matches that violate prerequisite requirements. For students with approved exceptions, we "force" a particular course match or allow multiple matches.

Students are informed of the class to which they matched (or that they failed to match), and a seat is reserved for them by the Registrar's Office by granting permission to enroll in a "shadow" section of the class (and reducing capacity of the non-shadow section).

Students then follow the normal registration process: they first meet with their advisor (with knowledge of their matched CS class), and then register for courses. A seat is reserved for each matched student during priority registration, but they may choose to enroll or not. Any unreserved seats are available for anyone to claim (including those who matched to some other CS course). Students may add themselves to the waitlist of any fully enrolled CS course.

Ⓒ *After Priority Registration, transfer students onto course rosters.* Soon after Priority Registration ends, unclaimed Match reservations expire, and students are transferred to non-shadow sections.

2.2 Alternative Mechanisms

We considered several other registration-based mechanisms to achieve our main goals, but rejected them—often for reasons of external campus considerations or constraints. Approval for the Match hinged on compatibility with existing registration mechanisms, eliminating many possible solutions. (For example, a round-robin system, in which *all* students choose a first course [in seniority order] before *any* student chooses a second course, accomplishes our main goals—but at such a cost in duration and complexity of the registration system that it was a nonstarter to consider.)

(1) We would have preferred for Step Ⓑ to *directly* enroll students in their matched courses, decreasing complexity. But this was infeasible due to student advising policies.

(2) We considered seeking a *ceiling* rather than a (near-universal) *floor* on enrollment in CS courses. But per-term or career-long enrollment limits would have been a more radical departure from existing campus policies, and we wanted students to (accurately) see our goal as *expanding students' access* to courses, rather than view the system as *restricting (different) students' access.*

(3) A flow-based maximum-matching algorithm, in lieu of Gale–Shapley, would provably maximize the number of matched students. But this would decrease attention to student preferences and eliminate the institutional principle of prioritizing more senior students. It would also mean students could game the system [3].

3 SOURCES OF DATA FOR ANALYSIS

The Match was approved by Carleton in early 2020, beginning with a "roll out" term: during registration for Spring 2020, a limited number of seats (128 in total) were held out for the Match. All seats in post-CS2 computer science courses were available in the Match for Fall 2020 registration and beyond.[2] A section is a single offering of a course in a particular term and timeslot.

Throughout, *the pre-Match period* denotes Fall 2015–Winter 2020 (a pre-intervention comparison period); *the Match period* denotes Fall 2020–Spring 2023, when the Match was used fully. Spring 2020, the capacity-limited roll-out term, is omitted from both. When averaging across all terms of an academic year to smooth out predictable seasonal variations, we omit the 2019–2020 year. Almost all sections had 34 seats, and there were slightly fewer Match-eligible sections per term in the pre-Match period (mean: 7.1 vs. 8.2).

In the rest of this paper, we report on students' engagement with the Match, registration outcomes, and waitlist and workload impacts of the Match system, drawing on the following data sources.

Data collected through the Match. We consider student-submitted preferences and their matched courses (or lack of match) in registration for each term of the Match period.

Registrar data. We record for each course its title, the term, and the list of students who were enrolled in it at the end of the one-week Drop/Add Period (see Figure 1). We classified each CS course as *Match-eligible,* or not, based on whether it would have been included in the Match had the Match been in use at the time.

Demographic data from our U.S.-based institution include a student's entering cohort (e.g., Fall 2020), actual or expected graduation year (e.g., 2024), gender, and race/ethnicity.[3] Our data do not indicate a student's class year when they enroll in a course, nor do they note leaves of absence. Carleton typically permits only 12 terms in residence, so analyses involving class year are limited to the ≈90% of students whose graduation year is four years after their entering cohort, meaning their class year for each course is known. All results follow the same trends if all students are included.

Waitlist data. During and after Priority Registration, a student who wishes to enroll in a full course may join its waitlist. After Priority Registration, instructors have discretion regarding which and how many waitlisted students to admit.[4] Waitlist data was available for registration for Fall 2017–Spring 2023.

4 STUDENT PARTICIPATION IN THE MATCH

With data in hand, we now turn to analyzing the use of (Section 4), and result of using (Section 5), the Match. In regard to student participation, many students—including non-CS majors—participate, and we accommodate their preferences fairly well. Most students enroll in their matched course, but a sizable minority do not.

4.1 Who participates in the Match?

Between 225 and 326 students, about 11–16% of Carleton's student body, have participated in the Match every term since it was introduced: . Student participation has modestly increased year over year, but, more saliently, the number of participating students

[1]Python implementation: https://github.com/annarafferty/carleton-cs-match

[2]One section in Spring 2021 was omitted due to being added after Priority Registration.

[3]This dataset has notable limitations: it allows only for binary gender identification and stores identities as reported in students' admission applications. Race/ethnicity information follows conventions from the Common Dataset Initiative [13].

[4]The process for admitting students from the waitlist is multistage and largely manual.

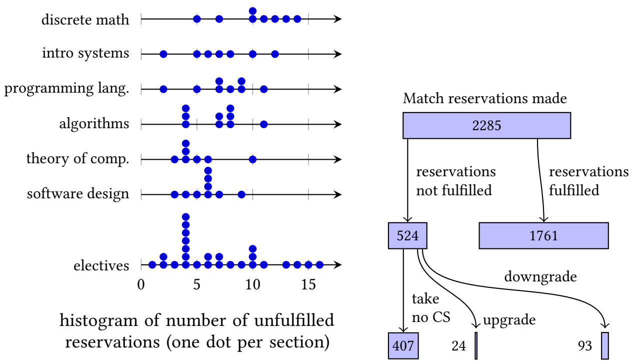

Figure 2: Students who match to classes, but opt not to enroll.

trends upwards *within* the three terms in each academic year (2020–2021; 2021–2022; 2022–2023). This pattern is consistent with the fact that, as the year goes on, more and more students complete the prerequisites for Match-eligible courses beyond CS2. (First-year students also can't participate until Winter.)

Of Match submissions from students with a known class year (see Section 3), just under half (995 of 2186, 45.5%) were made by CS majors.[5] Many fewer submissions came from students who have declared majors but are not CS majors (286, 13.1%); of these, students from STEM majors are predominant, with social science disciplines and studio art also frequent. A larger fraction of submissions came from students who have not yet declared a major (905, 41.4%).

4.2 Do we meet submitted student preferences?

Consistent with our goal Meet Needs (and Most Desires), we find student preferences are largely being met: 74% of students matched to their top-ranked course, and 88% matched to either their first- or second-ranked course. Nearly all students (94%) who submit preferences successfully matched to one of the courses in their submitted lists. (The lengths of students' preference lists varies, with a mode of 3 courses and wide spread.) The 6% of submissions that did not result in a match included 18 instances in 2020–2021, 47 in 2021–2022, and 98 in 2022–2023. It is clear that the number of Match failures is increasing: the high of 26 in Winter 2023 was quickly eclipsed by the 66 failures in Spring 2023. This was due in part to a particular high-demand course both terms, but also reflects increasing interest in CS overall.

4.3 Do students enroll in their matched course?

The Match is not a *registration* system but rather a *reservation* process (see Section 2): students may (or may not) choose to fulfill their reservation by registering for their matched course. As measured by registration at the end of the start-of-term Drop/Add Period, 1761 of 2285 (77.1%) of reservations were fulfilled (i.e., the students followed through and enrolled). The choice not to enroll in a matched course is spread across nearly every course (see Figure 2), ranging from ≈10–30% of the typically 34 seats in each class. In 22% of the cases of unfulfilled reservations, the student instead enrolled in one (or more) other Match-eligible courses. (Most of these students "downgrade" to a course they originally ranked lower than their matched course.)

Why do so many students not enroll in their matched course? One reason is that students lose nothing by submitting a Match form, even if they are uncertain about their plans. Their only cost is completing a Google Form; their potential benefit is a free-to-discard guarantee of a seat in a broadly sought-after course. Uncertainty about plans likely causes our discrete math course to be a significant outlier in non-enrollment, as it is typically a student's 2nd or 3rd CS course and has an alternative in the math department.

[5]Students declare majors at the end of their second year. We count submissions as from CS majors if our data indicate the student majored in CS (including if they double-majored) and the preference form would have been filled out after they had declared a major. This may miscount slightly due to students changing majors.

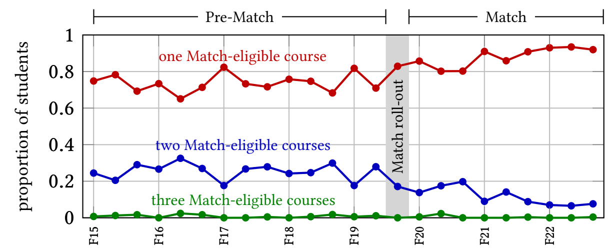

Figure 3: The term-by-term distribution of the number of Match-eligible courses taken by each student.

5 CHANGES IN ENROLLMENT PATTERNS

We hoped the Match would encourage enrollment patterns more consistent with liberal-arts-college values: increase how many students can take a CS course and decrease the number taking many simultaneously. We thus compare enrollment patterns in the pre-Match period relative to those during the Match, where changes may be caused directly by Match reservations and potentially by Match-related cultural shifts. As we expand on below, our classes are more diverse in terms of students' class years post-Match, and students tend to take fewer CS courses at once. We do not have any evidence that the Match is driving away underrepresented groups.

In keeping with success in our goal to Reduce Bingeing, *fewer students now take multiple simultaneous CS classes.* In the pre-Match period, 73.8% of students enrolled in Match-eligible courses in a term were taking one such course, rising to 88.1% of students during the Match period (Figure 3). Students averaged significantly fewer Match-eligible courses per term during the Match period (GLM with fixed factors for Match status and term within the year, random factor for student; $\beta(\text{Match}) = -.14$, $t = 10.6$, $p < .001$).

In accordance with our goal to Make the Tent Bigger*, more (and younger) students now take CS classes each term.* Because more students are taking only a single Match-eligible course (Figure 3), our courses have more *distinct* students. Moreover, Match-eligible courses now enroll students from a greater variety of class years (Figure 4), with the proportion of enrolled students in each class year having changed with the implementation of the Match ($\chi^2(3) = 69.2$, $p < .001$). The proportion of first-year students in Match-eligible courses more than doubled to 7.3% (cf. 3.3% pre-Match); the proportion of second-year students rose from 24.2% to 30.5%.

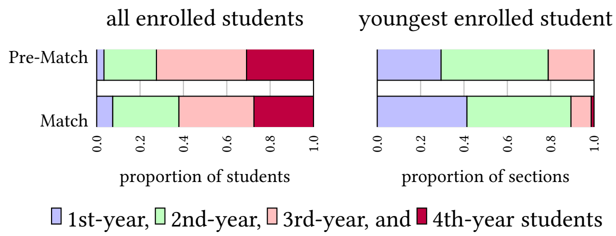

Figure 4: Student enrollment in Match-eligible courses, by class year (left); and the class-year composition of sections, by the seniority of the youngest enrolled student (right).

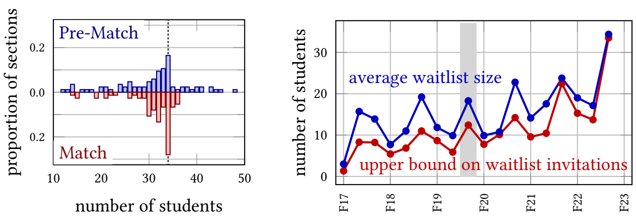

(a) Course enrollments. **(b) Waitlist sizes over time.**

Figure 5: The size of (a) courses and (b) waitlists before and after the Match. The dotted line in (a) is the standard nominal course capacity; the gray bar in (b) is the roll out of the Match.

The increased share of younger students impacted many courses: the number of course sections with at least one first-year student grew from 29% pre-Match to 41% in the Match period (again, see Figure 4). Overall, 89% of Match sections had at least one first- or second-year student (cf. 79% pre-Match). We see these changes in enrollment as beneficial for all students: a diversity of class years brings variety in perspectives to discussions and collaborative work, and we hope that increasing the accessibility of courses to less-senior students leads to CS feeling more welcoming and inclusive.

The increased number of distinct, and younger, students has not come at the cost of higher class sizes: average roster size following the first-week Drop/Add Period has remained between 29 and 32 in all years from 2015–2016 through 2022–2023. Almost all sections have capacity 34, but faculty may individually increase enrollment beyond this limit. Increasing enrollment in individual courses has historically been a fraught topic, as faculty disagree about its costs and benefits. Perhaps because the Match provides a different way to accommodate students—goal Meet Needs (and Most Desires)—faculty now more rarely choose to exceed nominal capacities, and those rarer instances have been by a smaller margin (Figure 5a).

It is less clear if the Match helps to Make the Tent Bigger *with respect to demographics other than age, but we do not have evidence that it hurts.* Among enrollments by students who had declared a major, there was little change in the proportion from non-CS majors (17.4% pre-Match vs. 15.0% in the Match period). The Match period has also seen substantial increases in the number of CS majors, meaning there are fewer interested students who are non-majors.

We examined associations between enrollment/Match participation and available gender and race/ethnicity information to determine if there was evidence that the Match was alienating underrepresented students. We first investigated if demographics were related to whether students choose to fulfill their Match reservations. Using registrar data of student gender (which is restricted to be binary), we found that 74.6% of matches for female students were fulfilled, compared to 78.4% of matches for male students. While statistically significant (GLM with fixed effect for gender and random effect for student; $\beta(\text{Male}) = .30$, $Z = 2.1$, $p = .038$), the effect is small and there are multiple interpretations: this difference could indicate that the extra step after the Match is a barrier to female students, or that female students are more likely to participate in the Match even if they are not committed to registering.

Because registrar data about race/ethnicity is coarse, lumping together all international students, our analyses related to race/ethnicity consider three groups of students: U.S. students of color, U.S. [non-Hispanic] white students, and international students. We found no difference across groups in reservation fulfillment.

We next examined enrollment proportions based on demographic categories. Match-eligible courses enrolled a similar proportion of female students before (29.9%) and after (30.8%) the Match. There were changes in the composition of our courses by race/ethnic group: the proportion of seats filled by U.S. white students decreased markedly from 62.7% before the Match to 46.9% after the Match. This large shift is likely driven by a similar change in the student body at Carleton: from Fall 2015 to Fall 2022, the proportion of U.S. white students went from 63.4% of 52.4% [10].

Finally, we looked for demographic associations with whether students enrolled by the Match or by standard registration/waitlists. No significant differences by demographic category were found. Overall, these results suggest that, even as enrollment pressures increase, the Match avoids the negative impacts on underrepresented groups that can occur with competitive enrollment schemes.

6 IMPACT ON WAITLISTS

The Match was developed in part because of increasing faculty time and student anxiety devoted to managing waitlists—the one pre-Match place in which the extant registration system allowed for values-driven decision-making by faculty, seeking to achieve goals like those detailed here but with limited information and tools. Inviting a student from the waitlist requires manual processing by the Registrar's Office, so large numbers of invitations are costly.

Historical waitlist data are coarse, allowing us only to partition students on each waitlist into three categories: *droppers* (who drop themselves from the waitlist without enrolling); *invitees* (those offered a seat in the class [they may accept or not]); and *nonactors* (who neither receive an invitation nor leave the waitlist).

While these three categories do not directly translate to the number of waitlist invitations per term, we do know that all *invitees* and some (but not all) *droppers* were invited, and no *nonactors* were. We upper bound the number of invitations by assuming that all *droppers* are invited.[6] There is significant seasonal variability, but overall both average waitlist length and the number of invitations were increasing prior to the Match and have further increased

[6]The same trends in results hold if we instead assume no *droppers* are invited.

during the Match period (Figure 5b). These increases are likely due to both increased enrollment and the Match.

Why should the Match increase waitlist size and *movement*? The Match offers *reservations*, and some students do not enroll in their matched course (see Section 4.3). Because all or most seats are assigned in the Match, waitlists form earlier than before the Match. Scarcity of seats and a lack of cost encourages Match participation, further increasing waitlists. One solution would be to "overbook" classes, matching 1.1–1.3× as many students as the class's capacity based on historical non-enrollment. If institutionally viable, moving away from a reservation process might also reduce waitlist movement. Decreasing such movement would decrease workload and be beneficial for maintaining the stability properties of Gale–Shapley.

7 ADDITIONAL IMPACTS AND DISCUSSION

Some consequences of the Match are hard to assess but are still important for understanding how the department and institution have been impacted. Here are a few such areas.

Culture and experiences within the department regarding waitlists. Before the Match, there was increasing informal faculty coordination around waitlist management, seeking (in an ad hoc way) to admit more unique students to CS classes without enrolling other students in multiple CS classes. Such efforts occurred with minimal information about student preferences. Now information from Match submissions can inform any coordination, which is primarily focused around students who participated in the Match but failed to match. Some faculty further prioritize waitlisted students who are in fewer (or no) CS classes over those in more. Faculty are more likely to avoid overenrollment (see Section 5). Some faculty report less angst around waitlists, as there is less need for individual solutions to systemic challenges.

Match failures, especially in repeated terms, are felt negatively by students. Rather than only being waitlisted, as occurred prior to the Match, students first are told they have failed to match and then later must (try to) enroll via waitlist.

Workload. The Match imposes significant workload. In the CS department, one faculty member manages the Match each year, coordinating with students and the Registrar's Office. In addition to collecting preferences and reporting Match results to the Registrar's Office, this coordinator manages an email list for student petitions and questions. The number of petitions varies (≈3–20), and each is considered individually. Pre-written responses handle many FAQs (e.g., *I missed the deadline*) but not all (e.g., *I'm not sure whether to take X or Y*). An extra source of advising benefits students—especially those less connected to CS—but requires faculty time.

The Match creates substantial extra work for the Registrar's Office, as alluded to in Section 2, with software limitations forcing manual entry of reservations. Staff in the Registrar's Office report that the Match is a significant burden to execute, but that they appreciate how the CS department has handled the process, with an established point of contact, clear deadlines, and congeniality. Informal requests by other departments to use a similar system have been denied in part due to workload constraints. Other schools may face smaller hurdles: an upcoming institutional software switch is anticipated to notably reduce workload by simplifying reservations.

Rising numbers of CS majors. At Carleton, the number of computer science majors has repeatedly reached record highs, increasing 30% from the pre-Match comparison period to the Match period (average of 58 to 75 majors per year). As of 2023, over 15% of junior and senior Carleton students have declared a CS major.

We are concerned about this trend not just in terms of staffing but also because we see diversity in areas of study as essential for a vibrant liberal arts culture. Against our will, we have become a hegemon, and the increasing number of CS majors is viewed negatively by some colleagues in other departments. In an unfortunate coincidence, use of the Match coincides precisely with the onset and continuing impact of COVID-19. The pandemic, as well as fears of recession, may have led students towards majors that are perceived as more career-oriented and have higher median pay [4, 7]. And yet the Match may have contributed to this increase, too. One of our goals was to reduce barriers to post-CS2 courses. Increasing the number of majors was neither a goal nor anticipated. But, in hindsight, we should have anticipated it: making individual CS courses more accessible is likely to make majoring in CS more accessible.

To the extent that the number of majors was previously depressed by enrollment barriers, we view the increase in majors positively. However, it is possible that the Match's structure leads students to disproportionately prioritize CS. Match participation precedes advising meetings and registration. Matching to a course may result in students being more emotionally attuned to losing a guaranteed seat than in choosing to enroll in an alternative course [14, 15]. Further, Match participation could nudge students towards CS courses by forming an implementation intention [19]. More generally, while declaring one's preferences requires some reflection, matching could circumvent later or more prolonged reflection; students may implicitly view a match as a recommendation to be followed without critically considering other options [11].

8 CONCLUSION

The Match is our response to increasingly untenable enrollment pressures: student angst was high, faculty attempts to address the problem were limited by lack of information and control, and the system was serving no one well. We believe the Match has mitigated some of these challenges, although we acknowledge a number of confounding factors. Due to COVID-19, all CS courses from Spring 2020–Spring 2021 were held online. Prior to the Match, faculty members used many ad hoc solutions to address waitlists: our comparison is not to doing nothing but to these (unsustainable) efforts. Further, our comparison implicitly assumes similar enrollment pressures in the pre-Match and Match periods, but, especially in the last year, these pressures have grown. We have not systematically examined student or faculty feelings about the system.

The context of our institution impacted the details of our implementation of the Match, but we are excited about the potential for it to be adapted to other institutions. The exact mechanism for and ease of integrating the Match process into an existing registration system will depend on the underlying registration software. The petition system could be made more scalable, such as by shifting load to advisors. We hope this experience report provides a case study in how to instantiate a department's values when managing enrollment as an alternative to competitive enrollment approaches.

REFERENCES

[1] Atila Abdulkadiroğlu, Nikhil Agarwal, and Parag A Pathak. The welfare effects of coordinated assignment: Evidence from the New York City high school match. *American Economic Review*, 107(12):3635–3689, 2017.
[2] Atila Abdulkadiroğlu, Parag A Pathak, and Alvin E Roth. The New York City high school match. *American Economic Review*, 95(2):364–367, 2005.
[3] Atila Abdulkadiroğlu, Parag A Pathak, Alvin E Roth, and Tayfun Sönmez. Changing the Boston school choice mechanism. *National Bureau of Economic Research*, 2006.
[4] Esteban M Aucejo, Jacob French, Maria Paola Ugalde Araya, and Basit Zafar. The impact of COVID-19 on student experiences and expectations: Evidence from a survey. *Journal of Public Economics*, 191:104271, 2020.
[5] Eric P Bettinger and Bridget Terry Long. Mass instruction or higher learning? The impact of college class size on student retention and graduation. *Education Finance and Policy*, 13(1):97–118, 2018.
[6] Zachary Bleemer and Aashish Mehta. College major restrictions and student stratification. Working paper, available at SSRN 3981921, 2022.
[7] Erica Blom, Brian C Cadena, and Benjamin J Keys. Investment over the business cycle: Insights from college major choice. *Journal of Labor Economics*, 39(4):1043–1082, 2021.
[8] Carla E Brodley. Why universities must resist GPA-based enrollment caps in the face of surging enrollments. *Communications of the ACM*, 65(8):20–22, July 2022.
[9] Tracy Camp, W Richards Adrion, Betsy Bizot, Susan Davidson, Mary Hall, Susanne Hambrusch, Ellen Walker, and Stuart Zweben. Generation CS: The growth of computer science. *ACM Inroads*, 8(2):44–50, 2017.
[10] Carleton College Office of Institutional Research and Assessment. Carleton's Common Data Set, 2023.
[11] Mary T Dzindolet, Scott A Peterson, Regina A Pomranky, Linda G Pierce, and Hall P Beck. The role of trust in automation reliance. *International Journal of Human-Computer Studies*, 58(6):697–718, 2003.
[12] David Gale and Lloyd S Shapley. College admissions and the stability of marriage. *The American Mathematical Monthly*, 69(1):9–15, 1962.
[13] Common Data Set Initiative. https://commondataset.org.
[14] Daniel Kahneman, Jack L Knetsch, and Richard H Thaler. Anomalies: The endowment effect, loss aversion, and status quo bias. *Journal of Economic Perspectives*, 5(1):193–206, 1991.
[15] Daniel Kahneman and Amos Tversky. Prospect theory: An analysis of decision under risk. *Econometrica*, 47(2):263–292, 1979.
[16] Kathleen J Lehman, Julia Rose Karpicz, Veronika Rozhenkova, Jamelia Harris, and Tomoko M Nakajima. Growing enrollments require us to do more: Perspectives on broadening participation during an undergraduate computing enrollment boom. In *Proceedings of the 52nd ACM Technical Symposium on Computer Science Education*, pages 809–815, 2021.
[17] National Academies of Sciences, Engineering, and Medicine and others. *Assessing and responding to the growth of computer science undergraduate enrollments*. National Academies Press, 2018.
[18] An Nguyen and Colleen M Lewis. Competitive enrollment policies in computing departments negatively predict first-year students' sense of belonging, self-efficacy, and perception of department. In *Proceedings of the 51st ACM Technical Symposium on Computer Science Education*, pages 685–691, 2020.
[19] Sheina Orbell, Sarah Hodgkins, and Paschal Sheeran. Implementation intentions and the theory of planned behavior. *Personality and Social Psychology Bulletin*, 23(9):945–954, 1997.
[20] Parag A Pathak and Tayfun Sönmez. School admissions reform in Chicago and England: Comparing mechanisms by their vulnerability to manipulation. *American Economic Review*, 103(1):80–106, February 2013.
[21] Elizabeth Patitsas, Michelle Craig, and Steve Easterbrook. A historical examination of the social factors affecting female participation in computing. In *Proceedings of the 2014 Conference on Innovation & Technology in Computer Science Education*, pages 111–116, 2014.
[22] Elizabeth Patitsas, Michelle Craig, and Steve Easterbrook. How CS departments are managing the enrolment boom: Troubling implications for diversity. In *2016 Research on Equity and Sustained Participation in Engineering, Computing, and Technology (RESPECT)*, pages 1–2. IEEE, 2016.
[23] QuestBridge. questbridge.org.
[24] Alvin E Roth. The origins, history, and design of the resident match. *Journal of the American Medical Association*, 289(7):909–912, 2003.
[25] Linda J Sax, Kathleen J Lehman, Jerry A Jacobs, M Allison Kanny, Gloria Lim, Laura Monje-Paulson, and Hilary B Zimmerman. Anatomy of an enduring gender gap: The evolution of women's participation in computer science. *The Journal of Higher Education*, 88(2):258–293, 2017.
[26] Tayfun Sönmez. Bidding for Army career specialties: Improving the ROTC branching mechanism. *Journal of Political Economy*, 121(1):186–219, 2013.