Annenberg Institute
BROWN UNIVERSITY

EdWorkingPaper No. 26-1485

Meeting People Where They Are: Experimental Evidence on Embedded Supports, Service Use, and Educational Outcomes

Kelli A. Bird
University of Virginia

Benjamin L. Castleman
University of Virginia

Many public services suffer from persistently low take-up despite high potential returns. A growing body of evidence suggests that information alone does little to close this gap; instead, hassle costs and default access points may be binding constraints on utilization. We test whether reallocating existing services to where people already are – rather than requiring them to seek out centralized offices – can improve utilization and outcomes.

We study this question in community colleges, where even modest frictions can impede students from engaging with available support services. We conducted a multi-semester randomized controlled trial in which treatment sections had tutoring and advising staff embedded directly in the classroom. Faculty and staff targeted these supports toward students at risk of academically struggling. Control sections followed standard instructional practices. Critically, both treatment and control students had access to the same on-demand tutoring and advising through centralized campus offices.

Embedding targeted supports into the classroom generated large increases in utilization: within the pilot course, treated students were 31 percentage points more likely to meet with a tutor, 29 percentage points more likely to meet with an advisor, and 12 percentage points more likely to meet with their instructor. Treatment students were 5 percentage points more likely to pass their courses and 4.6 percentage points less likely to withdraw. Impacts spilled over beyond the focal course: treated students were more likely to meet with faculty and staff outside the pilot course, and showed gains in overall semester GPA and persistence into subsequent terms. Our findings demonstrate that resource reallocation can meaningfully improve outcomes for underserved students.

VERSION: May 2026

Suggested citation: Bird, Kelli A., and Benjamin L. Castleman. (2026). Meeting People Where They Are: Experimental Evidence on Embedded Supports, Service Use, and Educational Outcomes. (EdWorkingPaper: 26-1485). Retrieved from Annenberg Institute at Brown University: https://doi.org/10.26300/pbk1-zf09

Meeting People Where They Are:
Experimental Evidence on Embedded Supports, Service Use, and Educational Outcomes

Kelli Bird and Ben Castleman, University of Virginia

Abstract

Many public services suffer from persistently low take-up despite high potential returns. A growing body of evidence suggests that information alone does little to close this gap; instead, hassle costs and default access points may be binding constraints on utilization. We test whether reallocating existing services to where people already are – rather than requiring them to seek out centralized offices – can improve utilization and outcomes.

We study this question in community colleges, where even modest frictions can impede students from engaging with available support services. We conducted a multi-semester randomized controlled trial in which treatment sections had tutoring and advising staff embedded directly in the classroom. Faculty and staff targeted these supports toward students at risk of academically struggling. Control sections followed standard instructional practices. Critically, both treatment and control students had access to the same on-demand tutoring and advising through centralized campus offices.

Embedding targeted supports into the classroom generated large increases in utilization: within the pilot course, treated students were 31 percentage points more likely to meet with a tutor, 29 percentage points more likely to meet with an advisor, and 12 percentage points more likely to meet with their instructor. Treatment students were 5 percentage points more likely to pass their courses and 4.6 percentage points less likely to withdraw. Impacts spilled over beyond the focal course: treated students were more likely to meet with faculty and staff outside the pilot course, and showed gains in overall semester GPA and persistence into subsequent terms. Our findings demonstrate that resource reallocation can meaningfully improve outcomes for underserved students.

Acknowledgements: We are very grateful for the collaboration of VCCS and PVCC administrators, instructors, advisors, and tutors, particularly Cat Finnegan, David Tod, and Marina Bagrev at the VCCS; and John Donnelly, Andrew Renshaw, David Lerman, and Brittany Resmann at PVCC. We are grateful to the Ascendium Education Group for financial support of the project, under grant number G-202209-17857. We are grateful to Yifeng Song and Renzhe Yu for their contributions to the project, and to Elizabeth von der Heydt for research assistance. Any errors or omissions are our own.

I. Introduction

Many public services suffer from persistently low take-up despite high potential returns. Across domains — from social safety net programs to preventive healthcare to job training — eligible populations systematically underutilize available services (Currie 2006). For instance, approximately 20–25 percent of eligible taxpayers fail to claim the Earned Income Tax Credit despite the credit being worth hundreds or even thousands of dollars (Bhargava and Manoli 2015). Under 70 percent of eligible adults receive recommended colorectal cancer screening, even when fully covered by insurance (Davidson et al. 2021). These patterns persist despite the substantial documented benefits of participation, suggesting that standard models of rational behavior do not fully capture the barriers individuals face in accessing available services.

A natural policy response has been to provide better information about existing programs. Yet experimental evidence increasingly suggests that information alone is rarely sufficient to close large utilization gaps. Informing unemployed workers about available job search assistance did not significantly improve employment outcomes (Altmann et al. 2018). Explicitly informing employees they could increase their retirement savings by leveraging employer contribution matches had no effect on contribution rates (Choi, Laibson, and Madrian 2011). Where information does increase take-up, the effects are typically modest relative to the additional gains from reducing the hassle of acting on it. Finkelstein and Notowidigdo (2019) found that informing individuals of their likely SNAP eligibility increased enrollment by 5 percentage points, but adding application assistance tripled this effect. Similarly, Bettinger et al. (2012) found that information about financial aid eligibility had no effect on college enrollment, but pairing it with help completing the FAFSA substantially increased enrollment.

These results point to a consistent pattern: the binding constraint on utilization is often not a lack of knowledge but a lack of action. A substantial literature in behavioral economics formalizes this insight, demonstrating that seemingly minor frictions — hassle costs, time constraints, the cognitive burden of navigating unfamiliar systems — can suppress participation even when benefits far exceed costs (Bertrand, Mullainathan, and Shafir 2006). The Creating Moves to Opportunity experiment offers a clear illustration: when offered housing vouchers to move to higher-opportunity neighborhoods, only 15 percent of eligible families successfully relocated, but adding search assistance increased this to 53 percent (Bergman et al. 2024). For many populations, particularly those facing severe time and resource constraints, the transaction costs of navigating to and accessing services may weigh more heavily than the magnitude of the direct financial benefits.

This insight raises a potentially important but understudied policy question: can we improve the relative efficiency of public service expenditures by reallocating existing resources to where people already are – thereby changing the default access point – rather than requiring them to seek out centralized services? Across multiple policy domains, embedded service models exist and have expanded substantially, though their effectiveness relative to centralized delivery has not been rigorously evaluated. School-based health centers, which place primary care and mental health services directly in schools rather than requiring families to travel to community clinics, now serve over 6 million students nationally and have expanded rapidly over the past two decades (Love et al. 2019). These centers are associated with increased preventive care utilization and reduced emergency department visits, though existing evidence comes primarily from observational studies (Knopf et al. 2016; Ran, Chattopadhyay, and Hahn 2016). In K-12 education, high-impact tutoring programs that embed individualized support directly in the school day -- rather than requiring students to attend after-school sessions -- have produced large gains in academic performance. Guryan et al. (2023) found that Saga Education's model of in-school, during-the-day tutoring in small consistent groups increased math achievement by 0.37 standard deviations in a randomized trial with over 2,700 Chicago high school students, with effects persisting into future years. However, the existing experimental literature on tutoring compares embedded tutoring to business as usual -- not to the same services offered through centralized channels (Nickow, Oreopoulos, and Quan, 2024).

The absence of experimental evidence on whether embedding services causally improves outcomes represents an important gap in our understanding of effective service delivery design. Without randomized variation in service location, it is difficult to distinguish whether improved outcomes reflect the accessibility benefits of embedding services or simply selection into locations where embedded services are placed. Moreover, existing studies of embedded services typically measure utilization of the embedded service itself but not utilization of alternative centralized services, leaving ambiguous whether embedding increases total service use or simply shifts the location where services are consumed.

We provide an experimental test of whether embedding existing support services improves service use and in turn outcomes by reducing barriers to access. We conducted a multi-semester randomized controlled trial at a community college, where students in courses with high historic failure rates were randomly assigned to sections with or without embedded tutoring and advising support. In both treatment and control sections, students had identical access to centralized tutoring and advising services through the college's learning center and student success office -- services that required students to travel to a central campus location during operating hours. Each treated section was assigned an embedded tutor who attended class regularly throughout the semester and a team of advisors who similarly maintained regular in-class presence. Working in coordination with section instructors, embedded tutors and advisors proactively identified and reached out to students showing early signs of academic struggle to provide both academic support and assistance addressing non-academic barriers to course success.

This design allows us to isolate the causal effect of embedding and targeting services on both resource utilization and academic outcomes, while holding constant the availability of centralized support services. Critically, we collected parallel data on resource use in both treatment and control sections, allowing us to observe not only whether students in treated sections used more embedded support, but also whether this came at the expense of reduced utilization of centralized services. This measurement of utilization patterns across both service locations is uncommon in existing studies of service delivery, where researchers typically observe use of only the focal intervention.

We conducted the experiment at Piedmont Virginia Community College (PVCC) across three semesters (Fall 2023, Spring 2024, and Fall 2024), randomizing a total of 52 course sections to treatment or control. Participating sections spanned multiple entry-level courses in mathematics, English, Spanish, and biology -- high-enrollment courses with historically high failure and withdrawal rates.

Our results demonstrate that embedding and targeting services generates large increases in utilization. Students in treated sections were 30.8 percentage points more likely to meet with a tutor, 29.4 percentage points more likely to meet with an advisor, and 12.2 percentage points more likely to meet with an instructor within the pilot course. These increases in staff resource use are substantially larger than typical effects from information-based interventions, suggesting that accessibility -- not just awareness -- is the binding constraint on utilization for time-constrained populations. These effects on service utilization extended well beyond the pilot course sections. Treated students were also significantly more likely to meet with tutors and advisors in their other courses -- a 12.1 percentage point increase in tutoring engagement and a 19.6 percentage point increase in advisor engagement outside the pilot course. Treated students were also 15.5 percentage points more likely to use other campus resources. These spillover effects suggest that embedding and targeting supports does not simply shift where services are consumed but may change students' broader orientation toward help-seeking, perhaps by familiarizing them with support staff and lowering the perceived costs of initiating contact in other settings.

Embedding targeted supports also improved academic outcomes. Treatment students were 5 percentage points more likely to pass their pilot courses and 4.6 percentage points less likely to withdraw. These impacts spilled over beyond the focal course: treated students were more likely to pass courses across their full course load, less likely to withdraw from other courses, and showed gains in overall semester GPA. Treated students were 10.1 percentage points more likely to persist into the subsequent semester, and 10.8 percentage points more likely to persist into the following year. Impacts on pilot course passing were especially concentrated among new students -- population least familiar with campus resources -- who were 31.5 percentage points more likely to pass in treated sections.

These findings make several contributions to existing literature. First, we provide novel experimental evidence on a fundamental but understudied dimension of program design: where services are physically located and how this shapes default patterns of access. While prior research has examined how defaults affect whether people participate in programs (Johnson and Goldstein 2003; Madrian and Shea 2001), we show that the spatial placement of services represents an equally important default. When the default is "services available if sought," utilization remains low even when services are free and benefits are high. Embedding services and targeting them to people likely to benefit changes the default to "services proactively offered and easy to use," dramatically increasing take-up.

Second, our results speak to a different mechanism than prior experimental work on hassle costs and administrative burden. Existing studies demonstrate that simplifying application processes, reducing paperwork requirements, and providing assistance with complex forms can increase program participation (Bhargava and Manoli 2015; Bettinger et al. 2012). Our intervention changed *where* services were offered, eliminating the time cost of traveling to centralized locations and the psychological cost of initiating contact with unfamiliar staff. This distinction matters because resource reallocation may be feasible even in contexts where administrative simplification is difficult, and because spatial accessibility may be particularly important for populations facing severe time poverty.

Finally, our findings have important implications for cost-efficiency. Even if embedding services increases total demand for staff time, the large utilization gains we document imply that embedded delivery is likely more efficient per dollar spent than centralized services that few students use. In our setting, embedding nearly quadruples tutoring engagement in the pilot course and increases overall engagement in student success advising by nearly sixfold – implying that embedded delivery remains more cost-efficient unless per-section costs increase by factors of four and six, respectively, a threshold that likely exceeds any plausible cost premium associated with classroom-based delivery.[1] While implementation quality may be more difficult to maintain at scale, the magnitude of both the utilization and academic impacts we observe suggests that reallocating resources to meet people where they are may be substantially more cost-efficient than expanding centralized service capacity.

The remainder of this paper proceeds as follows. Section II provides more detail on the intervention design. Section III describes the data, randomization, and empirical strategy. Section IV presents our main results. Section V discusses implications for service delivery design and concludes.

[1] Overall engagement in student success advising includes sessions based on the pilot course, non-pilot courses, and non-course specific sessions. The estimates specific to pilot courses and non-pilot courses are in Table 4, while the overall student success advising estimate is in Table A1.

II. Setting and Intervention Design

A. Piedmont Virginia Community College

Piedmont Virginia Community College (PVCC) is a public two-year institution in the Virginia Community College System (VCCS) located in Charlottesville, Virginia. PVCC enrolls approximately 5,300 students, roughly 80 percent of whom attend part-time (IPEDS, 2023). The student body is approximately 60 percent White, 12 percent Black, 11 percent Hispanic, and 5 percent Asian, with 57 percent female and 43 percent male. Approximately 46 percent of undergraduates receive Pell Grant aid, and 28 percent of students in our experimental sample are first-generation college students. PVCC's demographic profile is broadly consistent with the VCCS system overall, where approximately 82 percent of students attend part-time (JLARC, 2025). PVCC's Pell share is somewhat above the statewide average of 34 percent across all Virginia public institutions (SCHEV, 2024).

Like most community colleges, PVCC provides a range of centralized academic support services to all enrolled students. The college's Tutoring Center and Writing Center offer walk-in and appointment-based tutoring in most subjects, with experienced tutors providing assistance. General academic advising (e.g., course registration, selecting degree pathways) is provided through the Admissions & Advising Center, while the Student Success Office houses success coaches who provide holistic, ongoing academic and personal support. Success coaches help students to: define, clarify, and achieve educational and personal goals; find appropriate resources to address any potential obstacles or barriers; make informed choices regarding academic decisions; and provide encouragement and support throughout their educational pursuits.[2]

Accessing these tutoring and student success services requires students to navigate to specific campus locations during limited operating hours, identify the appropriate office, and initiate contact with staff – steps that, while individually modest, represent meaningful friction for PVCC's larger part-time, commuter student population As we document below, only a small fraction of students in our control group engaged with centralized tutoring or advising services during the study period, despite these services being available and free.

[2] https://www.pvcc.edu/student-success-office

B. Gateway Courses

Our study focuses on gateway courses in mathematics, English, biology, and Spanish – high-enrollment courses that are required for upper-level coursework and degree completion and that historically have high rates of failure and withdrawal. Nationally, 42 percent of students in gateway math courses and 33 percent in gateway English courses earn a D, F, or withdrawal (Koch and Gardner, 2017). At PVCC, 33.7 percent of students earned a D, F, W, or I (incomplete) in the focal courses in the year prior to our pilot.[3]

C. The Embedded Supports Model

The embedded supports intervention placed tutoring and advising staff directly in treated course sections and paired them with course instructors in a coordinated support model. The intervention had three core components: (1) a structured early-term identification process in which instructors worked with the Student Success Office to flag students likely to struggle ("priority" students); (2) embedded tutors assigned to each section who provided proactive, course-specific academic support; and (3) a team of dedicated Student Success advisors who rotated across treated sections and addressed non-academic barriers to student success. The instructor, tutor, and advisor for each section met periodically throughout the semester to coordinate outreach and reassess priority students. We describe each component in turn.

<u>Instructor-led identification of priority students</u>. Prior to the start of each semester, we conducted a training for participating instructors in which we covered the overall intervention design, expectations for tutors, advisors, and staff coordination throughout the semester, reinforcement of PVCC expectations for interaction reporting and Canvas LMS use, and behaviorally-informed strategies for proactive outreach to students.[4] During the first three weeks of the semester, instructors in treated sections met with staff from the Student Success office to identify students at elevated risk of academic difficulty. The Student Success Office prepared a starter list for these meetings by analyzing administrative data on students' prior course performance (for returning students) and early-term Canvas LMS engagement patterns across courses. Instructors refined this list based on their own classroom observations and interactions during the opening weeks,

[3] Authors' calculations using VCCS administrative data described in Section III.A.
[4] We are grateful to Greg Walton for his collaboration and guidance on developing training materials on evidence-based approaches for instructor outreach.

with the goal of identifying approximately half of the section for priority outreach. Beginning in week four, instructors coordinated with the embedded tutors and advisors to contact each priority student. This early identification process ensured that outreach was targeted to students likely to benefit from the additional support.

Embedded tutors. Each treated section was assigned an embedded tutor who attended every class session throughout the semester. Tutors circulated during group and individual work time, worked with students individually when instructors made time available during class, and were accessible before and after class and during outside office hours. Having tutors attend every session served a dual purpose: it made them a visible, familiar presence that students could approach without the psychological costs of initiating contact with an unfamiliar staff member at a separate campus location, and it strengthened tutors' own grasp of the specific instructional content covered in the course, which could in turn improve the quality and relevance of their academic support. Embedded tutors provided proactive, one-on-one academic tutoring to priority students, focusing on course-specific content rather than generic study skills. PVCC conducted new rounds of tutor hiring each semester.

Dedicated advisors. A team of Student Success advisors was dedicated to the treated sections and rotated across them on a regular schedule. Unlike the tutors, whose role focused on course-specific academic content, advisors proactively addressed non-academic barriers to student success – including challenges with balancing work and academic schedules, food insecurity, and transportation difficulties (such as a car breaking down) that could prevent students from attending class. Advisors provided both direct support and referrals to existing campus and community resources. The advising staffing model evolved across semesters: in Fall 2023, one full-time equivalent (FTE) advisor was assigned to the treated sections; in Spring 2024, the model shifted to multiple part-time advisors including the FTE from the first semester; and in Fall 2024, the team comprised several part-time advisors including those from prior semesters, providing some continuity in the advising relationship.

Coordination across the support team. A distinguishing feature of the intervention was the structured coordination among the instructor, tutor, and advisor assigned to each section. The three met periodically to review the priority student list, discuss individual students' progress and emerging challenges, and adjust support strategies as needed. Beginning around week eight, the team reassessed the priority list to add students who had begun struggling mid-semester and to recalibrate support for students whose needs had changed. This coordination model ensured that academic and non-academic support was integrated rather than siloed, and that the three staff members were aware of each other's interactions with students.

D. Control Condition

Students in control sections received standard instructional practices. They had access to the same centralized tutoring and advising services available to all PVCC students through the Tutoring and Writing Centers and the Student Success office. The experiment did not restrict students' access to any campus resources; it only varied whether tutoring and advising were additionally embedded in the classroom. Control instructors were aware that other sections included embedded supports.

III. Empirical Strategy

A. Data

Our primary data source is administrative data collected and maintained by the Virginia Community College System. VCCS provided us with de-identified student-by-course level data for all credit-bearing courses in the pilot terms, which allow us to observe enrollment in a course section in our experimental sample as well as enrollment in other courses, basic student demographic characteristics, and grades earned. We also incorporate VCCS administrative data from terms prior to the pilot program, which allows us to observe prior student academic history (cumulative GPA and earned credits) and prior instructor experience at VCCS. We also incorporate data from terms after the pilot program; we use these data along with VCCS graduation records and National Student Clearinghouse matches (maintained within the VCCS administrative data system) to construct a measure of student persistence.

One limitation to the VCCS administrative data is that we do not observe students who withdrew or changed course sections within the add/drop period. However, the college's add/drop period ends roughly 2-3 weeks after the start of the term, which is before instructors and advisors identified priority students. So, being identified as a priority student would not cause students to be more likely to drop the course and not appear in our data. Still, there is a concern that students in treated sections would be more likely to stay enrolled once learning about the embedded resources in their course section, while similar students in the control sections dropped the course and do not appear in our data. If we assume that the students induced to stay enrolled are less likely to achieve a positive outcome in the course or less likely to persist, then this selection issue would bias our results toward zero.

We also use data from PVCC's EAB Navigate system on students' use of staff resources. This data includes records of instructor meetings and tutoring and advising sessions with staff embedded in the treated course sections, as well as instructor meetings for students in control sections and tutoring and advising sessions with staff in PVCC's tutoring centers and Student Success office. In many cases, these meetings are tied to a specific course, though students do sometimes seek out general advising or tutoring, for example, "Writing Assistance (any PVCC class)." We also observe meetings with the Academic Advising office (which focuses on course registration and program selection), Career Services, Financial Aid, Student Accessibility and Accommodations, and the Student Resource Center. We aggregate these meetings into an "Other Resources" category. In supplemental analysis, we also use data from PVCC's learning management system's (Canvas) user course access logs, which include the total number of views of each specific subpage within the course section's Canvas site. We construct measures by summing the total number of views within a specific content type.[5]

B. Randomization

We worked with PVCC to implement the pilot program over the course of three terms: Fall 2023, Spring 2024, and Fall 2024. Shortly before the start of each term, PVCC identified a set of gateway courses to focus on and recruited instructors slotted to teach those courses to participate in the pilot program. Once instructors signed on to the study, and there were two or more participating instructors within a particular course (e.g., ENG 111), we then randomly assigned

[5] There are three key limitations to this data: (1) we are unable to observe the total amount of time spent with the Canvas resources or which actions, if any, the students performed on the site (e.g. assignment submission); (2) potential variation across instructors in how they structure their Canvas sites; and (3) we do not have these data for the Fall 2023 cohort.

instructors *within-course* to either the treatment or control condition. If an instructor was teaching multiple sections of the course, we randomly selected one of their course sections to be treated, and removed the other course sections from the experimental sample.[6] Each term, 17-18 instructors participated in the pilot; instructors were allowed to participate in multiple terms, but did not carry their experimental status from one term to the next. Twenty-five unique instructors participated in the pilot across the three terms.[7]

Table 1 shows the number of sections by course and term in the experimental sample. These cells also represent our randomization blocks. Our final experimental sample includes 11 biology sections, 19 English sections, 18 math sections, and 2 Spanish sections, for a total of 52 course sections. Twenty-seven sections were randomly assigned to treatment, and the remaining 25 were assigned to a "business as usual" control condition.

As we describe above, selection into the experimental sample was not random, as instructors were actively recruited to participate in the pilot program prior to randomization. In Table 2, we compare the characteristics of course sections and instructors in the experimental sample versus the non-experimental sample. Panels A and B of Table 2 show that experimental course sections were broadly similar to the non-experimental sample. Experimental sections were slightly more likely to enroll female students, white students, and students with more accumulated college credits and higher GPAs. Panel B of Table 2 shows that compared with non-experimental course sections, experimental sections had slightly higher enrollments and were more likely to be a full

[6] If an instructor taught more than one of the gateway courses, we assigned them to one gateway course prior to the randomization to maximize the within-course variation we could achieve in the sample. For an illustrative example, suppose three instructors were teaching ENG111 sections, one instructor was teaching ENG112 sections, and one instructor was teaching both ENG111 and ENG112 sections. The last instructor was assigned to ENG112 for the purposes of randomization so that ENG112 could be included in the experimental sample.

[7] Due to some fluidity of PVCC course section scheduling up until the first week of courses, we performed the randomization as close to the start of each term as possible while also providing sufficient time to hire tutors for the treated course sections. There were a few instances where the initial random assignments had to be altered. In Spring 2024, one instructor's course section (assigned to treatment) was dropped from the schedule and that instructor was removed from the pilot that term. In Spring 2024 there were three instances where instructors were reassigned to a different section of the same focal course; in these cases, the instructors maintained the experimental status to their new section (one control and two treatment sections). Finally, in Fall 2024 there was one instance where the original participating instructor was replaced with a separate instructor for an experimental course section. The new instructor agreed to participate in the pilot and was kept in the experimental sample with the originally assigned experimental condition (treatment).

session (i.e., 16-17 week course) versus a shorter session.[8] Because only one section per participating instructor was chosen for the experimental sample, several of the non-experimental sample sections were taught by participating instructors.

In Panel C, we compare instructors who ever participated in the experimental sample to those who never did. Half of all instructors teaching at least one section of a focal course participated in the study. Participating instructors had more prior experience teaching one or more of the focal courses (88% versus 64% with any prior experience in the past five years; 8.2 versus 4.8 terms of prior experience in the past five years). However, participating instructors awarded lower grades in prior sections of focal courses compared with non-participating instructors (2.56 versus 2.85 GPA for Fall 2023). Finally, participating instructors were much more likely to be full-time compared with non-participating instructors (76% versus 12%).

C. Estimating Impacts

We estimate the intent-to-treat impact of enrolling in a treated course section using the following OLS regression model:

$$Outcome_{isnc} = \beta_0 + \beta_1 Treated_{snc} + \Gamma StudentChars_i + \Delta SectionChars_s + \\ + \Omega InstructorChars_{nc} + CourseFixedEffects_s + \epsilon_{isnc}$$

where the data is at the level of student i in course-section s with instructor n in course c. To improve precision, we include the following covariates: student characteristics measured at the start of the term (cumulative GPA at PVCC; credits accumulated at PVCC; dual enrollment indicator; first-generation status, in-state status; race; gender; new student status; and whether the student is pursuing a transfer-oriented associate degree); course-section characteristics (course enrollment; average cumulative GPA at PVCC of the other enrolled students); and instructor characteristics (number of prior terms the instructor taught course s; the average GPA of students

[8] This was by design in the first pilot term, where we only considered full session courses to give the instructors, advisors, and tutors sufficient time to identify priority students. We relaxed this requirement for the following terms.

in the instructors' prior sections of course s for the past five years).[9] Finally, as randomization occurred within courses, we include course fixed effects, and we cluster the standard errors at the course-section-level.

We focus on four categories of outcomes: (1) student use of faculty, tutoring, or student success advising resources, separated by resource use for the pilot course versus other courses or general use[10]; (2) student academic outcomes in the pilot course, including course grade, pass rate, and whether students withdrew from the course; and (3) student overall academic outcomes in the pilot term, and (4) persistence to the next term. We pre-registered the experiment at the AEA RCT Registry and pre-specified the following primary outcomes: whether students pass the course, either with a "C" or higher or a "B" or higher, and whether students withdraw from the course prior to the end of the semester.[11] We estimate these outcomes within the pilot course and separately estimate these outcomes averaged across courses in which the student was enrolled in the term. Finally, we define persistence as whether the student either: enrolled at VCCS in one or two terms following the pilot term (excluding summer terms); graduated from VCCS in the pilot term; or enrolled at a non-VCCS college in one or two terms following the pilot term (excluding summers), as observed in the National Student Clearinghouse matches. The vast majority of students who persisted by this definition re-enrolled at VCCS in the first following term (94 percent of persisters) or the second following term (83 percent of persisters). Because the third and fourth categories of outcomes are calculated at the student x term level (instead of student x course section x term level), we drop the 34 students (3.6 percent of unique students) who were enrolled in more than one pilot course section in the same term, and did not have the same experimental status across sections. Our results are very similar if we include these students.[12]

---

[9] We do not include the binary course section and instructor characteristics shown in Table 2 as randomization controls or in baseline equivalence tests because there is insufficient variation of these variables within randomization blocks.

[10] The Navigate records resource use at the student-by-meeting, so we can differentiate tutor, advisor, and instructor meetings. Meetings could be linked to a specific course, so we can differentiate between meetings for the pilot courses versus other courses. We include meetings coded as "general", i.e. that do not specify a course in our outcomes measuring students' use of resources outside of the pilot course.

[11] https://www.socialscienceregistry.org/trials/15886

[12] Of the 933 unique students in the sample, 173 enrolled in multiple course sections over the three pilot terms. Twenty-eight students enrolled in multiple experimental course sections in the same term with the same experimental status; 82 students enrolled in multiple experimental course sections across different terms with the same experimental status; and 52 enrolled in multiple experimental course sections across different terms with different experimental status. Our results are very similar when we exclude the last group of students.

We also test whether there are differential impacts by estimating subgroup-specific impacts based on student enrollment status (new or returning), parental education (first generation college-goer or not), and baseline academic achievement (returning students with a GPA above or below a 2.0).

D. Baseline Equivalence Tests

We test for baseline equivalence by regressing each of the Student, Course, and Instructor characteristics listed above in Section III.C. on an indicator for treatment, and include randomization block fixed effects. We present the results in Table 3. In Panel A, we find two statistically significant differences in the student-level characteristics: students in treated sections have lower accumulated college credits (column 12) and are more likely to be a new student (column 13); these two characteristics are necessarily related as all new students have accumulated zero college credits. In Panel B, we also observe that treated course sections have slightly smaller enrollments (1.5 fewer students, compared to a mean of 24.3 in the control sections) and are taught by instructors with fewer prior terms teaching the pilot course (2.2 fewer terms, compared to a mean of 9.2 in control sections). Given both the small number of randomization blocks within each pilot term and the small number of course sections within each randomization block, the modest imbalance at the section/course level is predictable. We control for all these as covariates in our impact models.

IV. Results

A. Impacts on Resource Use

We first examine whether the embedded supports models in treated course sections increased students' use of staff resources. Table 4 presents the intent-to-treat impacts on whether students had any meeting with a tutor, student success advisor, or instructor, separately for meetings related to the pilot course (Panel A) and for meetings related to other courses or general support (Panel B).[13]

[13] We repeat our main impact analyses removing the regression covariates (student, instructor, and course section characteristics), and display the results in Table A1. Our main impact estimates in Tables 4, 5, and 6 are quite similar to those shown in Panels A, B, and C of Table A1, respectively, though those estimated without the regression covariates are less precise.

Panel A of Table 4 shows that embedding supports generated large increases in students' engagement with all three types of staff for the pilot course. Students in treated sections were 30.8 percentage points (403 percent) more likely to meet with a tutor for the pilot course, compared to 7.7 percent of students in control sections (column 1). Virtually no students in control sections met with a student success advisor specifically for the pilot course (0.4 percent), while embedding advising supports led to a 29.4 percentage point increase among students in treated sections (column 2). Students in treated sections were 12.2 percentage points (343 percent) more likely to meet with the instructor of the pilot course, compared to 3.6 percent of students in control sections (column 3).

Embedding tutors and advisors in pilot courses led to positive spillovers in treated students' resource use outside the pilot courses, either for other courses or general tutoring. Students in treated sections were 12.1 percentage points (79 percent) more likely to meet with a tutor outside the pilot course, compared to 11.7 percent of students in control sections (Table 4, Panel B, column 1). Students in treated sections were also 19.6 percentage points (315 percent) more likely to meet with a student success advisor outside the pilot course, compared to only 6.2 percent of students in control sections (column 2). We find no spillover effects on faculty meetings outside the pilot course (column 3).[14] Treated students were also substantially more likely to engage with other campus resources, including academic advising, career services, financial aid, and the student resource center. Students in treated sections were 15.5 percentage points (35 percent) more likely to use any of these other resources, compared to 44.3 percent of control students (Appendix Table A2, column 7).

These spillover effects suggest that embedding supports does not simply shift where services are consumed but may change students' broader orientation toward help-seeking. Familiarizing students with support staff in one setting appears to lower the perceived costs of initiating contact in other settings, consistent with the hypothesis that psychological barriers – not just logistical ones – may suppress service utilization.

---

[14] We find similar impacts when we consider the number of meetings within the pilot course (Appendix Table A2, columns 1-3) and outside the pilot course (columns 4-6).

Within treated sections, the large increases in resource use were concentrated among students flagged as priority for proactive outreach. Appendix Table A3 presents descriptive evidence on the association between priority status and resource use among students in treated sections only (N = 571). Priority students were 31.1 percentage points (407 percent) more likely to meet with a tutor for the pilot course than non-priority students in the same treated sections, 7.7 percent of whom met with a tutor (column 1). Virtually no non-priority students in treated sections met with a student success advisor (0.4 percent), compared to over 50 percent of priority students (column 3). These results indicate that the coordinated outreach process — through which instructors, tutors, and advisors jointly identified and contacted students at elevated risk of academic difficulty — was appropriately targeted, directing the most intensive support toward students identified early in the term as most likely to struggle in the course.

As we report in Bird and Castleman (forthcoming), treated students appear to have substituted away from asynchronous online course resources in favor of direct staff supports. Appendix Table A4 presents impacts on students' page views of their course-specific Canvas sites (available for one of three cohorts; N = 331).

B. Impacts on Pilot Course Performance

We next examine whether the large increases in resource utilization documented in Section IV.A translated to improved academic performance. Table 5 presents the intent-to-treat impacts on academic outcomes in the pilot course.

Students in treated sections were 6.6 percentage points (13 percent) more likely to pass the pilot course with at least a B, compared to 49.3 percent of control students (column 1). Treated students were 5.0 percentage points (8 percent) more likely to pass with an A, B, or C, compared to 64.6 percent of control students (column 2). The larger impact on passing with at least a B than on passing with at least a C suggests that the treatment not only moved some students from failing to passing but also shifted the grade distribution upward among students who would have passed regardless. Treatment reduced the probability of withdrawing from the pilot course by 4.6 percentage points (34 percent reduction), from the control section mean of 13.7 percent withdrawal (column 3). Treatment increased average course grades by 0.13 grade points on a 4-point scale, though this estimate is imprecisely estimated (column 4).[15]

C. Impacts on Term-Level Performance and Persistence

Table 6 presents impacts on students' academic performance across their full course load in the pilot term and on their persistence into subsequent terms. As noted in footnote 12, because these outcomes are measured at the student-by-term level, we exclude the 34 students (3.6 percent of unique students) who were enrolled in multiple pilot course sections with different treatment statuses in the same term.

The academic gains from embedded supports extended beyond the pilot course. The A/B pass rate was 4.2 percentage points (9.2 percent) higher among students enrolled in treated sections, compared to 45.6 percent A/B pass rate among students in control sections (column 1). Similarly, the A/B/C pass rate was 4.0 percentage points (7.1 percent) higher among students in treated sections, compared to 56.7 percent A/B/C pass rate for students in control sections. The share of courses from which students in treated sections withdrew decreased by 2.2 percentage points (21.2 percent reduction), compared to 10.4 percent among students in control sections (column 3). Treatment increased semester GPA by 0.16 points, from a control group average of 2.58 (column 4).

The most striking results emerge for persistence. Students in treated sections were 10.1 percentage points (13 percent) more likely to persist to the next term, compared to 76.9 percent of students in control sections (column 5).[16] This effect continued over time: students in treated sections were 10.8 percentage points (18 percent) more likely to persist for at least two subsequent terms, compared to 60.6 percent of students in control sections (column 6). The persistence effects are larger in magnitude than the impacts on academic performance in the pilot course, consistent with the spillover patterns documented in Section IV.A — embedding supports

[15] The number of observations in column 4 (N=996) excludes students who withdrew and did not receive a course grade

[16] We define persistence as whether the student either: enrolled at VCCS in the following term (excluding summer terms); graduated from VCCS in the pilot term; or enrolled at a non-VCCS college in the following term (excluding summers), as observed in the National Student Clearinghouse matches. The vast majority of students who persisted into their first a second terms by this definition re-enrolled at VCCS (94% and 83%, respectively)

in a single course appears to have changed students' broader engagement with the institution, increasing both their use of campus resources and their likelihood of continued enrollment.

D. Heterogeneous Impacts

We next examine whether treatment effects differed across student subgroups. Table 7 presents impacts separately by first-generation status, enrollment status (new versus returning), and prior academic performance (cumulative GPA above or below 2.0). For each subgroup, we focus on four key outcomes: any tutoring, any student success advising, passing the pilot course with at least a B, and one-term persistence.

1. First-generation students

Treatment increased resource use among both first-generation and non-first-generation students, with broadly similar effects on the extensive margin of tutoring (35.7 versus 33.2 percentage points) and student success advising (43.1 versus 34.1 percentage points). However, the effects on academic outcomes were substantially larger for first-generation students. First-generation students in treated sections were 10.9 percentage points (31 percent) more likely to pass the pilot course with at least a B, compared to 35.0 percent of first-generation control students (column 3). The effect for non-first-generation students was a statistically insignificant 4.2 percentage points. The contrast is especially pronounced for persistence: treated first-generation students were 20.5 percentage points (31.5 percent) more likely to persist into the next term, compared to 65.1 percent of first-generation control students (column 4). The persistence effect for non-first-generation students was considerably smaller at 5.6 percentage points (7 percent increase; 80.2 percent of non-first-generation control students persisted). The similar resource use effects but divergent academic impacts are consistent with first-generation students having more to gain from each additional interaction with support staff – these students may lack the informal networks and institutional knowledge that continuing-generation students acquire through family experience with higher education, making the guidance provided by embedded tutors and advisors particularly valuable for translating effort into academic success.

2. New versus returning students

Treatment increased resource use among both new and returning students, with similar effects on tutoring (31.3 versus 26.3 percentage points) and somewhat smaller effects on student success advising for new students (18.7 versus 27.3 percentage points). Effects on academic outcomes were concentrated, however, among new students – those least familiar with campus resources and institutional norms. New students in treated sections were 31.5 percentage points (73 percent) more likely to pass the pilot course with at least a B, compared to 43.1 percent of new students in control sections. The effect for returning students was a statistically insignificant 4.3 percentage points. Persistence effects were also substantially larger for new students (27 percentage points; 34.9 percent increase, compared to 77.4 percent of new control students) than for returning students (10.7 percentage points; 13.9 percent increase, compared to 76.9 percent of returning control students). These results suggest that the intervention's greatest academic impacts are concentrated among students for whom campus services are most unfamiliar, reinforcing the interpretation that the status quo of centralized, seek-it-out service delivery is particularly costly for students who have not yet developed the institutional knowledge and relationships needed to navigate campus support systems.

3. Prior academic performance

Among returning students, we further examine whether effects varied by baseline cumulative GPA. Students entering with a GPA below 2.0 showed especially large treatment effects. The student success advising effect is particularly striking: treatment increased the share of students meeting with a student success advisor from 11 percent to over 77 percent (column 2), suggesting that embedded advising was especially effective at reaching students who were otherwise very unlikely to seek out support on their own. These academically vulnerable students were 14.6 percentage points (53 percent) more likely to pass the pilot course with at least a B, compared to only 27.4 percent of below-2.0 control students (column 3), and 21.2 percentage points (39 percent) more likely to persist, compared to 54.3 percent of below-2.0 control students (column 4). By contrast, among students with a GPA of 2.0 or above, treatment effects on passing are small and insignificant (3.1 percentage points), though persistence effects are economically meaningful at 10.1 percentage points (12.5 percent increase, compared to 81.1 percent of above-2.0 control students). This pattern suggests that embedded supports generate the largest academic gains for students with the most room for improvement and the greatest unmet need for intervention.

Across all three dimensions of heterogeneity, a consistent pattern emerges: while treatment generates broadly similar increases in resource use across subgroups, the translation of that increased engagement into improved academic outcomes is concentrated among the most vulnerable students — those who are first-generation, new to the institution, or entering with low prior academic performance. One explanation for this pattern is that the marginal returns to support are highest for students starting from a lower baseline of both academic performance and institutional integration. Students who already have familiarity with campus resources, established study habits, and networks of peers and mentors may benefit less from additional tutoring or advising interactions than students who lack these foundations. Embedded supports may thus function as a partial substitute for the institutional knowledge and social capital that more experienced or more advantaged students have already accumulated, explaining why similar increases in service use produce substantially larger academic gains for the most vulnerable populations.

E. Cost effectiveness comparison

We next explore the cost-effectiveness of the PVCC embedded supports program in comparison to other evidence-based efforts to improve persistence among community college students. The total cost of the embedded supports was $203,797: $108,632 to fund the embedded student success advisors; $41,165 to fund the embedded tutors; and $54,000 in instructor stipends.[17] This total cost translates to a per-treated student cost of $357 (based on 571 unique students enrolled in treated sections). Given the 10.8 percentage point increase in persistence into the next academic year among treated students, this translates to a cost per additional persister of $3,306.

[17] It is common practice at PVCC to provide instructors with stipends for participating in programs outside standard practices. In this case, instructors teaching control sections received a stipend of $1,000 each, and treated sections received a stipend of $3,000. We count the difference of these amounts toward the total program cost. If we instead include all stipends as program cost, then the total program cost would be $255,797 ($4,148 per each induced persistence). If we exclude stipends as a program cost (the more realistic scale-up cost), then the total program cost would be $149,797 ($2,429 per each induced persistence).

We compare in Table 8 the cost-effectiveness of the embedded supports program to other recent interventions that have similarly focused on improving academic outcomes, persistence, and degree attainment among current college students. We limit our comparison to studies that were also evaluated through RCTs, which provide sufficient detail on program costs, and which report impact estimates on persistence one year after the intervention (the most distal outcome we observe in our experiment).[18] Some interventions, like ASAP, have generated large impacts on persistence (and subsequent degree attainment), but they are also quite expensive. For instance, the ASAP replication in Ohio generated a similar impact on persistence as the PVCC embedded supports program (10.4 vs. 10.8 percentage point increase in persistence), yet the direct costs per student were an order of magnitude higher ($6,478 vs. $357). Across the programs we study, the implied cost per induced persister (column 7) ranges from $15,230 for Opening Doors Louisiana to $241,277 for CUNY ASAP. This comparison implies that the embedded supports are substantially more cost effective at improving persistence than the other programs considered in Table 8. An important note about this comparison is that we cannot speak to potential degree impacts of the embedded supports model, whereas other models like ASAP have demonstrated (and large) impacts on degree attainment.

V. Discussion

Our paper provides experimental evidence that reallocating existing support services to where people already are can generate significant increases in resource use that in turn leads to improved outcomes. We find that embedding tutors and advisors directly in classrooms rather than housing them in centralized offices, combined with targeted outreach from faculty and staff, generated large increases in service utilization and meaningful improvements in academic outcomes. These findings have implications for how institutions design service delivery, for our understanding of the behavioral barriers that suppress take-up of beneficial programs, and for the broader policy question of how to improve the efficiency of public service expenditures.

---

[18] We excluded three interventions due to no statistically significant overall impacts on persistence: Stay the Course (Evans et al, 2020); Opening Doors Learning Communities (Weiss et al 2014); and Learning Communities Demonstration (Visher et al, 2012). We excluded STAR due to lack of comparable persistence outcome measure (Angrist, Lang, & Oreopoulos, 2009). We excluded CUNY Start due to lack of sufficient information on program cost (Weiss et al, 2021). Finally, we excluded the Performance-Based Scholarship Demonstration due to lack of sufficient information on program cost and treatment sample size (Mayer et al, 2015).

The magnitude of the utilization effects we document is striking: within the pilot course specifically, students in treated sections were 30.8 percentage points more likely to meet with a tutor, 29.4 percentage points more likely to meet with an advisor, and 12.2 percentage points more likely to meet with an instructor over the course of the semester. We also find large and significant increases in use of tutoring and student success advising outside the pilot courses. These effects occur in a context where the same services were freely available to control students through centralized campus offices. The binding constraint on utilization in our setting was not awareness or cost, but accessibility – specifically, the logistical and psychological frictions involved in seeking out services at a separate campus location, during limited hours, from unfamiliar staff.

Our results extend the behavioral economics literature on defaults and hassle costs to a new dimension of program design: the physical location of services. We show that the spatial default – whether services are located where people already are or require them to seek out a separate location – represents an important and understudied margin. When the default is "services available if sought," utilization remains low even when services are free and benefits are potentially high. Embedding changes the default to "services delivered in place," dramatically increasing take-up.

A natural question is whether targeted, proactive outreach alone could generate similar effects. While we cannot empirically isolate the causal contribution of the outreach, we believe that the physical embedding of the staff was necessary to achieve the impacts we estimate. A robust literature demonstrates that informational outreach and nudge interventions in postsecondary education have limited effects on students' engagement with campus resources and on academic outcomes, even when carefully designed and personalized (Oreopoulos and Petronijevic 2019; Bird et al. 2021; Avery et al. 2021; Page et al. 2023). Most directly relevant is Bettinger et al. (2022), which reports null impacts from a large-scale, multi-site text-based outreach campaign ("Nudges to the Finish Line") that explicitly encouraged students at broad-access institutions to connect with campus tutoring, advising, and financial aid resources – an intervention designed and implemented at the same type of institutions we study, including several Virginia Community College System campuses. Despite high student response rates and personalized messages from campus advisors, the campaign produced no detectable effects on academic performance, credit accumulation, or degree attainment across over 21,000 students. These findings suggest that informing students about available resources and encouraging them to seek out support is insufficient when the services remain housed in centralized offices that students must navigate to on their own.

On the other hand, would embedded supports alone generate the same impacts? Our results suggest that embedded supports are a necessary but potentially not sufficient component of intervention design: the concentration of resource use effects among priority students (Appendix Table A2) indicates that both components – physical accessibility and targeted outreach – contributed to the large increases in engagement we observe. Proactive identification and outreach amplified the effects of accessibility by directing the most intensive support to students identified as most likely to struggle.

The absence of spillover effects on faculty meetings outside the pilot course (Table 4, column 3, bottom panel) further underscores the importance of proactive faculty outreach. While students in treated sections were more than four times as likely to meet with the instructor in their pilot section, treated students were no more likely to meet with faculty in their other courses, despite large increases in other resource use outside the pilot course. This pattern suggests that barriers to initiating contact with faculty are qualitatively different from–and likely higher than–barriers to engaging with tutors or advisors. Students may perceive faculty as less approachable due to differences in status, age, and perceived role (Cox 2009; Cotten and Wilson 2006), and may be particularly reluctant to reveal academic difficulty to the person evaluating their performance. These barriers may be especially acute at community colleges, where faculty-student interactions outside of class are limited by the commuting patterns of both students and instructors. Experimental evidence from Carrell and Kurlaender (2023) demonstrates that when faculty proactively reach out to students, students rate the instructor as being more approachable, available, and caring about their students, though faculty outreach on its own did not lead to significant improvements in academic performance.[19]

A related mechanism that may have contributed to the utilization effects within treated sections, though one we cannot empirically substantiate with our data, is positive social norming. When

[19] In small pilot studies (N=53 and N=138) the authors do find positive effects on academic outcomes, but these do not persist in the full-scale study (N=3,729).

embedded tutors and advisors interacted with students before, during, and after class – and when priority students visibly engaged with these staff members in the classroom setting – other students in the section may have updated their beliefs about the prevalence and social acceptability of help-seeking. A substantial literature in behavioral economics demonstrates that observing peers engage in a behavior can shift individuals' own willingness to participate, both through social learning (individuals infer the behavior's value from observing others) and through descriptive norms (Banerjee 1992; Cialdini and Goldstein 2004; Duflo and Saez 2003). In educational settings, peer behaviors have been shown to influence students' own academic and campus engagement (Sacerdote 2001; Carrell, Fullerton, and West 2009). Making help-seeking visible and normal within the classroom – rather than a private, stigmatized act conducted at a centralized office – may have lowered psychological barriers for students who had not been proactively contacted, potentially contributing to the elevated resource use rates we observe even among non-priority students in treated sections.

The timing of outreach may have also been an important factor in improving student performance. The coordinated outreach occurred during the fourth week of a 15-week semester, when students may not have realized they were struggling with the material and still had sufficient time to adjust their approach to class and improve academic performance.

The spillover effects we document suggest that embedding supports in a single course may change students' broader orientation toward help-seeking, perhaps by familiarizing them with support staff and lowering the perceived costs of initiating contact in other settings. If so, embedded supports could generate returns well beyond the focal course, a possibility consistent with the large persistence effects we observe. This interpretation aligns with research on the role of social integration and sense of belonging in student retention (Tinto 1993; Walton and Cohen 2011): for students who lack prior experience navigating institutional support systems, a single positive interaction with a tutor or advisor may serve as an entry point to a broader network of campus resources.

A natural concern about the embedded supports model is that dramatically increasing service utilization will require proportional increases in staffing and expenditure. But this concern overlooks an inefficiency in the status quo: under centralized delivery, institutions already invest substantially in tutoring and advising infrastructure that reaches only a small fraction of the students it is intended to serve. The relevant question is not whether embedded delivery costs more in total, but whether it generates more impact per dollar spent. The large utilization gains we document imply that embedded delivery is likely more efficient per dollar spent than centralized services that few students use. The magnitude of both the utilization and academic impacts we observe suggests that reallocating resources to meet people where they are may be substantially more cost-efficient than expanding centralized service capacity. Our cost effectiveness analysis moreover indicates the intervention is significantly more cost-effective than numerous other experimentally-evaluated strategies to improve college persistence.

Finally, our findings have broad implications for the design of public service delivery. The core insight – that reallocating existing resources to where people already are can dramatically increase utilization and improve outcomes – is not specific to higher education. Across domains where eligible populations underutilize available services, from preventive healthcare to social safety net programs to workforce development, embedding services in settings that people already frequent may be substantially more effective than expanding centralized capacity or providing additional information about existing services. The magnitude of the effects we observe suggests that meeting people where they are may represent one of the more cost-effective strategies for improving the reach and impact of public services.

VI. References

Altmann, Steffen, Armin Falk, Simon Jäger, and Florian Zimmermann. 2018. "Learning about Job Search: A Field Experiment with Job Seekers in Germany." *Journal of Public Economics* 164: 33–49.

Angrist, Joshua, Daniel Lang, and Philip Oreopoulos. 2009. "Incentives and Services for College Achievement: Evidence from a Randomized Trial." *American Economic Journal: Applied Economics* 1 (1): 136–163.

Avery, Christopher, Benjamin L. Castleman, Michael Hurwitz, Bridget Terry Long, and Lindsay C. Page. 2021. "Digital Messaging to Improve College Enrollment and Success." *Economics of Education Review* 84: 102170.

Banerjee, Abhijit V. 1992. "A Simple Model of Herd Behavior." *Quarterly Journal of Economics* 107 (3): 797–817.

Bergman, Peter, Raj Chetty, Stefanie DeLuca, Nathaniel Hendren, Lawrence F. Katz, and Christopher Palmer. 2024. "Creating Moves to Opportunity: Experimental Evidence on Barriers to Neighborhood Choice." *American Economic Review* 114 (5): 1281–1337.

Bertrand, Marianne, Sendhil Mullainathan, and Eldar Shafir. 2006. "Behavioral Economics and Marketing in Aid of Decision Making Among the Poor." *Journal of Public Policy and Marketing* 25 (1): 8–23.

Bettinger, Eric P., Benjamin L. Castleman, Alice Choe, and Zachary Mabel. 2022. "Finishing the Last Lap: Experimental Evidence on Strategies to Increase Attainment for Students Near College Completion." *Journal of Policy Analysis and Management* 41 (4): 1040–1059.

Bettinger, Eric P., Bridget Terry Long, Philip Oreopoulos, and Lisa Sanbonmatsu. 2012. "The Role of Application Assistance and Information in College Decisions: Results from the H&R Block FAFSA Experiment." *Quarterly Journal of Economics* 127 (3): 1205–1242.

Bhargava, Saurabh, and Dayanand Manoli. 2015. "Psychological Frictions and the Incomplete Take-Up of Social Benefits: Evidence from an IRS Field Experiment." *American Economic Review* 105 (11): 3489–3529.

Bird, Kelli A. and Benjamin L. Castleman (forthcoming). Do Embedded Supports Promote Engaged Learning? Experimental Evidence on Resource Use Among Community College Students. *American Economics Association: Papers and Proceedings.*

Bird, Kelli A., Benjamin L. Castleman, Jeffrey T. Denning, Joshua Goodman, Cait Lamberton, and Kelly Ochs Rosinger. 2021. "Nudging at Scale: Experimental Evidence from FAFSA

Completion Campaigns." *Journal of Economic Behavior and Organization* 183: 105–128.

Carrell, Scott E., Richard L. Fullerton, and James E. West. 2009. "Does Your Cohort Matter? Measuring Peer Effects in College Achievement." *Journal of Labor Economics* 27 (3): 439–464.

Carrell, Scott E., and Michal Kurlaender. 2023. "My Professor Cares: Experimental Evidence on the Role of Faculty Engagement." *American Economic Journal: Economic Policy* 15 (4): 113–141.

Choi, James J., David Laibson, and Brigitte C. Madrian. 2011. "$100 Bills on the Sidewalk: Suboptimal Investment in 401(k) Plans." *Review of Economics and Statistics* 93 (3): 748–763.

Cialdini, Robert B., and Noah J. Goldstein. 2004. "Social Influence: Compliance and Conformity." *Annual Review of Psychology* 55: 591–621.

Cotten, Shelia R., and Bonnie Wilson. 2006. "Student–Faculty Interactions: Dynamics and Determinants." *Higher Education* 51 (4): 487–519.

Cox, Rebecca D. 2009. *The College Fear Factor: How Students and Professors Misunderstand One Another.* Cambridge, MA: Harvard University Press.

Currie, Janet. 2006. "The Take-Up of Social Benefits." In *Poverty, the Distribution of Income, and Public Policy,* edited by Alan Auerbach, David Card, and John Quigley, 80–148. New York: Russell Sage Foundation.

Davidson, Karina W., Michael J. Barry, Carol M. Mangione, et al. (US Preventive Services Task Force). 2021. "Screening for Colorectal Cancer: US Preventive Services Task Force Recommendation Statement." *JAMA* 325 (19): 1965–1977.

Duflo, Esther, and Emmanuel Saez. 2003. "The Role of Information and Social Interactions in Retirement Plan Decisions: Evidence from a Randomized Experiment." *Quarterly Journal of Economics* 118 (3): 815–842.

Evans, William N., Melissa S. Kearney, Brendan Perry, and James X. Sullivan. 2020. "Increasing Community College Completion Rates Among Low-Income Students: Evidence from a Randomized Controlled Trial Evaluation of a Case-Management Intervention." *Journal of Policy Analysis and Management* 39 (4): 930–965.

Finkelstein, Amy, and Matthew J. Notowidigdo. 2019. "Take-Up and Targeting: Experimental Evidence from SNAP." *Quarterly Journal of Economics* 134 (3): 1505–1556.

Guryan, Jonathan, Jens Ludwig, Monica P. Bhatt, Philip J. Cook, Jonathan M. V. Davis, Kenneth Dodge, George Farkas, Roland G. Fryer Jr., Susan Mayer, Harold Pollack, Laurence Steinberg, and Greg Stoddard. 2023. "Not Too Late: Improving Academic Outcomes among Adolescents." *American Economic Review* 113 (3): 738–765.

Integrated Postsecondary Education Data System (IPEDS). 2023. National Center for Education Statistics. Data for Piedmont Virginia Community College (Unit ID: 233116). Washington, DC: U.S. Department of Education. https://nces.ed.gov/ipeds/.

Johnson, Eric J., and Daniel Goldstein. 2003. "Do Defaults Save Lives?" *Science* 302 (5649): 1338–1339.

Joint Legislative Audit and Review Commission (JLARC). 2025. *Virginia's Community Colleges and the Changing Higher Education Landscape.* Report 608. Richmond, VA: JLARC.

Knopf, John A., Ramona K. C. Finnie, Yinan Peng, Robert A. Hahn, Benedict I. Truman, Mary Vernon-Smiley, Veda C. Johnson, Robert L. Johnson, Jonathan E. Fielding, Carles Muntaner, Pete C. Hunt, Camara Phyllis Jones, and Mindy T. Fullilove; Community Preventive Services Task Force. 2016. "School-Based Health Centers to Advance Health Equity: A Community Guide Systematic Review." *American Journal of Preventive Medicine* 51 (1): 114–126.

Koch Andrew K. and John N. Gardner. 2017. Transforming the "Real" First-Year Experience: The Case for and Approaches to Improving Gateway Courses. In: Feldman RS, ed. *The First Year of College: Research, Theory, and Practice on Improving the Student Experience and Increasing Retention*. Cambridge University Press:126-154.

Love, Hayley E., John Schlitt, Samira Soleimanpour, Nirmita Panchal, and Caroline Behr. 2019. "Twenty Years of School-Based Health Care Growth and Expansion." *Health Affairs* 38 (5): 755–764.

Madrian, Brigitte C., and Dennis F. Shea. 2001. "The Power of Suggestion: Inertia in 401(k) Participation and Savings Behavior." *Quarterly Journal of Economics* 116 (4): 1149–1187.

Mayer, Alexander K., Reshma Patel, Timothy Rudd, and Alyssa Ratledge. 2015. *Designing Scholarships to Improve College Success: Final Report on the Performance-Based Scholarship Demonstration.* New York: MDRC.

Nickow, Andre, Philip Oreopoulos, and Vincent Quan. 2024. "The Promise of Tutoring for PreK–12 Learning: A Systematic Review and Meta-Analysis of the Experimental Evidence." *American Educational Research Journal* 61 (1): 74–107.

Oreopoulos, Philip, and Uros Petronijevic. 2019. "The Remarkable Unresponsiveness of College Students to Nudging and What We Can Learn from It." NBER Working Paper No. 26059. Cambridge, MA: National Bureau of Economic Research.

Page, Lindsay C., Bruce I. Sacerdote, Sara Goldrick-Rab, and Benjamin L. Castleman. 2023. "Financial Aid Nudges: A National Experiment with Informational Interventions." *Educational Evaluation and Policy Analysis* 45 (2): 195–219.

Ran, Tao, Sajal K. Chattopadhyay, and Robert A. Hahn; Community Preventive Services Task Force. 2016. "Economic Evaluation of School-Based Health Centers: A Community Guide Systematic Review." *American Journal of Preventive Medicine* 51 (1): 129–138.

Sacerdote, Bruce. 2001. "Peer Effects with Random Assignment: Results for Dartmouth Roommates." *Quarterly Journal of Economics* 116 (2): 681–704.

State Council of Higher Education for Virginia (SCHEV). 2024. *Key Facts about Higher Education in Virginia.* Richmond, VA: SCHEV.

Tinto, Vincent. 1993. *Leaving College: Rethinking the Causes and Cures of Student Attrition.* 2nd ed. Chicago: University of Chicago Press.

Visher, Mary G., Michael J. Weiss, Evan Weissman, Timothy Rudd, and Heather Wathington. 2012. *The Effects of Learning Communities for Students in Developmental Education: A Synthesis of Findings from Six Community Colleges.* New York: MDRC.

Walton, Gregory M., and Geoffrey L. Cohen. 2011. "A Brief Social-Belonging Intervention Improves Academic and Health Outcomes of Minority Students." *Science* 331 (6023): 1447–1451.

Weiss, Michael J., Alexander Mayer, Dan Cullinan, Alyssa Ratledge, Colleen Sommo, and Jonathan Diamond. 2014. *A Random Assignment Evaluation of Learning Communities at Kingsborough Community College: Seven Years Later.* New York: MDRC.

Weiss, Michael J., Susan Scrivener, Alyssa Slaughter, and Beth Cohen. 2021. *An On-Ramp to Student Success: A Randomized Controlled Trial Evaluation of a Developmental Education Reform at the City University of New York.* New York: MDRC.

Table 1: Number of Course Sections in Experimental Sample

	Total	Fall 2023	Spring 2024	Fall 2024
	(1)	(2)	(3)	(4)
BIO 141	9	3	2	4
BIO 142	2	0	2	0
ENG 111	10	6	2	2
ENG 112	9	2	5	2
MTH 154	9	4	3	2
MTH 155	4	0	2	2
MTH 161	7	2	2	3
SPA 101	2	0	0	2

Table 2: Comparing the experimental sample versus other PVCC course sections

Panel A: Average Student Characteristics in Course Sections

	Experimental (1)	Non-Experimental (2)
Dually enrolled	4.1%	3.2%
First generation college goer	28.1%	25.9%
Missing parental education	9.9%	8.2%
In-state	91.9%	92.4%
Female	63.0%	55.9%
White	56.2%	49.4%
Black	13.1%	19.3%
Hispanic	15.5%	14.8%
Asian	6.1%	5.7%
Other Race	9.1%	10.8%
Pursuing Transfer oriented AS degree	87.3%	87.1%
New Student	16.3%	23.1%
Accumulated college credits	15.23	11.63
Cumulative GPA	2.75	2.57
Student N	1133	2966

Panel B: Course section characteristics

	Experimental (1)	Non-Experimental (2)
Enrollment Count	21.79	18.20
Online	28.8%	33.1%
In-person	67.3%	65.0%
Hybrid	3.8%	1.8%
Full Session	92.3%	71.2%
Three-quarter session	1.9%	20.9%
Half session	5.8%	8.0%
Course Section N	52	163

Panel C: Instructor-level characteristics

	Experimental (1)	Non-Experimental (2)
Any prior experience teaching pilot course	88.0%	64.0%
Number of prior terms teaching pilot course	8.2	4.8
GPA of prior pilot course sections	2.56	2.85
Full-time	76.0%	12.0%
Instructor N	25	25

Notes: based on a sample of the focal courses at PVCC within each of the pilot terms. E.g. only Spanish 101 sections are in the sample for Fall 2024. The top panel compares characteristics of students enrolled in experimental course sections versus course sections outside the experimental sample. The second panel compares the characteristics of experimental versus non-experimental course sections. Note that a particular student can be enrolled in multiple experimental course sections and multiple non-experimental course sections. The bottom panel compares the characteristics of instructors participating in the experiment versus non-participating instructors. The first three instructor characteristics are based on the five years prior to the pilot term. The experimental sample of instructors includes any instructor who was in the experimental sample during at least one of the pilot terms; the non-experimental sample of instructors includes instructors who never participated in the pilot. The instructor characteristics are measured as of the earliest of the pilot terms during which the instructor was teaching at PVCC.

Table 3: Testing Baseline Equivalence

Panel A: Student-level baseline characteristics

	Dual enrolled	First Gen	Missing parental education	In-state	Female	White	Black	Hispanic	Asian	Other Race	Transfer oriented AS	New Student	Accum. credits	Cum. GPA
	(1)	(2)	(3)	(4)	(5)	(6)	(7)	(8)	(9)	(10)	(11)	(12)	(13)	(14)
Treated	-0.003	0.023	-0.014	-0.005	-0.044	-0.019	-0.025	0.035	-0.001	0.009	0.027	0.043**	-2.46***	0.042
	(0.012)	(0.029)	(0.018)	(0.016)	(0.028)	(0.030)	(0.020)	(0.022)	(0.015)	(0.017)	(0.020)	(0.018)	(0.937)	(0.067)
N	1,133	1,021	1,133	1,133	1,133	1,133	1,133	1,133	1,133	1,133	1,133	1,133	1,133	898
R-squared	0.055	0.05	0.057	0.015	0.087	0.025	0.023	0.014	0.014	0.023	0.042	0.27	0.212	0.16
Control mean	0.0463	0.271	0.0996	0.929	0.665	0.582	0.133	0.135	0.0605	0.089	0.852	0.128	16.94	2.859

Panel B: Course section and Instructor-level baseline characteristics

	Enrollment count	Cum GPA of enrolled students	GPA of prior pilot course sections	Prior terms teaching pilot course
	(1)	(2)	(3)	(4)
Treated	-1.491***	0.02	0.016	-2.23***
	(0.230)	(0.017)	(0.012)	(0.221)
N	1,133	1,133	1,061	1,133
R-squared	0.692	0.759	0.773	0.375
Control mean	24.33	2.779	2.606	9.21

Notes: within each panel, each column displays results from a separate regression of the baseline characteristic listed on an indicator for treatment and randomization block fixed effects (course x term). We exclude observations for which the relevant baseline characteristic is missing. Instructor characteristics are based on the five years prior to the pilot term. Robust standard errors in parentheses, *** p<0.01; ** p<0.05; * p<0.1

Table 4: Impacts on student use of staff resources

Panel A: Impacts on service use within the pilot course

	Any tutoring session	Any student success advisor meeting	Any faculty meeting
	(1)	(2)	(3)
Treatment	0.308***	0.294***	0.122***
	(0.029)	(0.027)	(0.031)
N	1,133	1,133	1,133
R-squared	0.264	0.235	0.225
Control mean	0.0765	0.00356	0.0356

Panel B: Impacts on service use outside the pilot course

	Any tutoring session	Any student success advisor meeting	Any faculty meeting
	(1)	(2)	(3)
Treatment	0.121***	0.196***	-0.004
	(0.025)	(0.026)	(0.005)
N	1,133	1,133	1,133
R-squared	0.155	0.143	0.037
Control mean	0.153	0.0623	0.0125

Notes: within each panel, each column displays results from a separate regression of the outcome on an indicator for treatment, the baseline characteristics listed in Table 3, and randomization block fixed effects (course x term). Robust standard errors in parentheses, *** p<0.01; ** p<0.05; * p<0.1

Table 5: Impacts on academic outcomes in the pilot course

	Passed course with A/B	Passed course with A/B/C	Withdrew	Grade
	(1)	(2)	(3)	(4)
Treated	0.066**	0.050*	-0.046***	0.134
	(0.030)	(0.029)	(0.016)	(0.108)
N	1,133	1,133	1,133	996
R-squared	0.225	0.143	0.058	0.23
Control mean	0.493	0.646	0.137	2.51

Notes: each column displays results from a separate regression of the outcome on an indicator for treatment, the baseline characteristics listed in Table 3, and randomization block fixed effects (course x term). Robust standard errors in parentheses, *** p<0.01; ** p<0.05; * p<0.1

Table 6: Impacts on pilot term academic and persistence outcomes

	% Courses passed with A/B	% Courses passed with A/B/C	% Courses withdrawn	Term GPA	Persisted to next term	Persisted one year
	(1)	(2)	(3)	(4)	(5)	(6)
Treated	0.042*	0.040**	-0.022**	0.157**	0.101***	0.108***
	(0.022)	(0.020)	(0.010)	(0.078)	(0.020)	(0.024)
N	1,055	1,055	1,055	1,004	1,055	1,055
R-squared	0.253	0.223	0.092	0.202	0.111	0.089
Control mean	0.456	0.567	0.104	2.577	0.769	0.606

Notes: each column displays results from a separate regression of the outcome on an indicator for treatment, the baseline characteristics listed in Table 3, and randomization block fixed effects (course x term). Robust standard errors in parentheses, *** p<0.01; ** p<0.05; * p<0.1

Table 7: Impacts by student subgroups

Panel A: First generation college going students

	Any tutoring session	Any student success advisor meeting	Passed pilot course with A/B	Persisted to next term
	(1)	(2)	(3)	(4)
Treated	0.357***	0.431***	0.109**	0.205***
	(0.039)	(0.036)	(0.045)	(0.047)
N	295	295	295	278
R-squared	0.311	0.358	0.281	0.29
Control mean	0.182	0.0657	0.35	0.651

Panel B: Non-first generation students

	Any tutoring session	Any student success advisor meeting	Passed pilot course with A/B	Persisted to next term
	(1)	(2)	(3)	(4)
Treated	0.332***	0.341***	0.042	0.056**
	(0.037)	(0.034)	(0.032)	(0.022)
N	726	726	726	671
R-squared	0.208	0.231	0.23	0.109
Control mean	0.195	0.065	0.553	0.802

Panel C: New Students

	Any tutoring session	Any student success advisor meeting	Passed pilot course with A/B	Persisted to next term
	(1)	(2)	(3)	(4)
Treated	0.277*	0.174**	0.315**	0.270**
	(0.150)	(0.071)	(0.123)	(0.100)
N	164	164	164	149
R-squared	0.235	0.277	0.422	0.261
Control mean	0.306	0.0972	0.431	0.774

Panel D: Returning students

	Any tutoring session	Any student success advisor meeting	Passed pilot course with A/B	Persisted to next term
	(1)	(2)	(3)	(4)
Treated	0.340***	0.381***	0.043	0.107***
	(0.028)	(0.027)	(0.027)	(0.020)
N	969	969	969	906
R-squared	0.23	0.279	0.237	0.129
Control mean	0.169	0.0612	0.502	0.769

Panel E: Students with GPA below 2.0

	Any tutoring session	Any student success advisor meeting	Passed pilot course with A/B	Persisted to next term
	(1)	(2)	(3)	(4)
Treated	0.365***	0.665***	0.146**	0.212**
	-0.089	-0.061	(0.072)	(0.098)
N	127	127	127	121
R-squared	0.435	0.584	0.4	0.408
Control mean	0.247	0.11	0.274	0.543

Panel F: Students with GPA above 2.0

	Any tutoring session	Any student success advisor meeting	Passed pilot course with A/B	Persisted to next term
	(1)	(2)	(3)	(4)
Treated	0.357***	0.355***	0.031	0.101***
	-0.033	-0.03	(0.031)	(0.021)
N	771	771	771	720
R-squared	0.279	0.328	0.318	0.103
Control mean	0.154	0.0559	0.532	0.811

Notes: within each panel, each column displays results from a separate regression of the outcome on an indicator for treatment, the baseline characteristics listed in Table 3, and randomization block fixed effects (course x term). Robust standard errors in parentheses, *** p<0.01; ** p<0.05; * p<0.1

Table 8: Cost-effectiveness of other interventions

Intervention (1)	Citation (2)	Type (3)	Sample (4)	Impact on 1 year persistence (5)	Total cost per student (6)	Implied cost per induced persister (7)
CUNY ASAP	Scrivener et al. (2015)	Comprehensive Support Program	451 Treated, 445 Control	7.7 pp**	$18,578	$241,277
ASAP Ohio Replication	Miller et al. (2020)	Comprehensive Support Program	806 Treated, 695 Control	10.4 pp***	$6,478	$62,290
InsideTrack Coaching	Bettinger & Baker (2014)	Coaching & Advising	8049 Treated, 5506 Control	5.3 pp***	$1,538	$29,021
Opening Doors Louisiana	Richburg-Hayes et al (2009)	Performance based incentives	505 Treated, 514 Control	11.8 pp***	$1,797	$15,230
Opening Doors Learning Communities	Scrivener et al. (2008); Sommo et al. (2012)	Learning Communities	769 Treated, 765 control	5.6 pp**	$2,161	$38,582

Notes: For CUNY ASAP, the impact on persistence is from Table C.1, and the total cost per student (based on direct costs only) is from Table 5.1. For ASAP Ohio Replication: the impact on persistence is from Figure ES.2, and the total cost per students is based on Table 6.1 (direct costs only). For Inside Track Coaching, the impact on persistence if from Table 3, and the total cost per student is from Footnote 14. For Opening Doors Louisiana, the impact on persistence is from Table 4.3, and the total cost per student is based on Table 3.1 (financial aid only). For Opening Doors Learning Communities, the impact on persistence is from Table 4.7 (Scrivener et al, 2008), and the total cost per student is based on Table 5.1 (Sommo et al, 2012). We converted the dollar amounts in columns (6) and (7) to 2023$.

Table A1: Impacts with no regression covariates

Panel A: Impacts on Resources Use

	Any tutoring session: pilot course	Any student success advisor meeting: pilot course	Any faculty meeting: pilot course	Any tutoring session: other	Any student success advisor meeting: other	Any faculty meeting: other
	(1)	(2)	(3)	(4)	(5)	(6)
Treated	0.267***	0.269***	0.073**	0.112***	0.192***	-0.005
	(0.029)	(0.028)	(0.032)	(0.025)	(0.023)	(0.004)
Observations	1,133	1,133	1,133	1,133	1,133	1,133
R-squared	0.219	0.195	0.133	0.113	0.111	0.024
Control mean	0.0765	0.00356	0.0356	0.153	0.0623	0.0125

Panel B: Impacts on Pilot Course Outcomes

	Passed course with A/B	Passed course with A/B/C	Withdrew	Grade
	(1)	(2)	(3)	(4)
Treated	0.045	0.041	-0.050***	0.065
	(0.031)	(0.027)	(0.015)	(0.103)
Observations	1,133	1,133	1,133	996
R-squared	0.114	0.069	0.037	0.11
Control mean	0.493	0.646	0.137	2.51

Panel C: Pilot term and Persistence Outcomes

	% Courses passed with A/B	% Courses passed with A/B/C	% Courses withdrawn	Term GPA	Persisted to next term	Persisted one year
	(1)	(2)	(3)	(4)	(5)	(6)
Treated	0.037	0.037*	-0.025**	0.12	0.093***	0.087***
	(0.023)	(0.019)	(0.011)	(0.077)	(0.020)	(0.025)
Observations	1,055	1,055	1,055	1,004	1,055	1,055
R-squared	0.127	0.134	0.049	0.062	0.062	0.049
Control mean	0.456	0.567	0.104	2.577	0.769	0.606

Notes: within each panel, each column displays results from a separate regression of the outcome on an indicator for treatment and randomization block fixed effects (course x term). Robust standard errors in parentheses, *** p<0.01; ** p<0.05; * p<0.1

Table A2: Impacts on additional resource use outcomes

	# of tutoring sessions: pilot course	# of student success advisor meetings: pilot course	# faculty meetings: pilot course	# of tutoring sessions: other	# of student success advisor meetings: other	# faculty meetings: other	Any student success advisor meeting	Any other resource meetings	# other resource meetings
	(1)	(2)	(3)	(4)	(5)	(6)	(7)	(8)	(9)
Treated	0.829***	0.893***	0.211***	0.162	0.448***	-0.001	0.356***	0.155***	0.340***
	(0.119)	(0.105)	(0.055)	(0.199)	(0.076)	(0.005)	(0.028)	(0.021)	(0.048)
N	1,133	1,133	1,133	1,133	1,133	1,133	1,133	1,133	1,133
R-squared	0.2	0.179	0.203	0.069	0.127	0.037	0.234	0.14	0.179
Control mean	0.354	0.00356	0.0391	0.804	0.126	0.0125	0.0658	0.443	0.806

Notes: within each panel, each column displays results from a separate regression of the outcome on an indicator for treatment, the baseline characteristics listed in Table 3, and randomization block fixed effects (course x term). Robust standard errors in parentheses, *** p<0.01; ** p<0.05; * p<0.1

Table A3: Differences in resource use between priority and non-priority students in treated sections

	Any tutoring session: pilot course	Any student success advisor meeting: pilot course	Any faculty meeting: pilot course	Any tutoring session: other	Any student success advisor meeting: other	Any faculty meeting: other
	(1)	(2)	(3)	(4)	(5)	(6)
Priority	0.311***	1.117***	0.507***	1.565***	0.022	0.089
	(0.051)	(0.230)	(0.056)	(0.226)	(0.055)	(0.093)
N	571	571	571	571	571	571
R-squared	0.317	0.316	0.429	0.324	0.3	0.264
Non-priority mean	0.0765	0.354	0.00356	0.00356	0.0356	0.0391

Notes: within each panel, each column displays results from a separate regression of the outcome on an indicator for priority student, the baseline characteristics listed in Table 3, and randomization block fixed effects (course x term). The sample is restricted to students enrolled in treated sections. Robust standard errors in parentheses, *** p<0.01; ** p<0.05; * p<0.1

Table A4: Impacts on engagement with Canvas

	Total	Assignments	Attachments	Discussion	Modules	Wiki
	(1)	(2)	(3)	(4)	(5)	(6)
Treated	-230.1**	-37.24**	-18.60**	0.95	-20.66**	-37.63***
	(100.28)	(15.67)	(8.46)	(22.84)	(7.69)	(7.10)
Observations	331	331	331	331	331	331
R-squared	0.31	0.286	0.411	0.301	0.241	0.536
Control mean	929.9	119.1	80.13	68.31	56.94	48.32

Notes: Each column presents the results from a separate student x section level regression of the outcome of interest on an indicator of treatment, the baseline characteristics listed in Table 3, and randomization block fixed effects (course x term). The outcomes based on Canvas records of the total number of views of each page type over the course of the term. The sample is limited to the Spring 2024 term. See Bird & Castleman (2026) for more information.