Building Success: The Effect of On-Campus Housing Expansions on Student Outcomes

Ciarán Murphy[a,1]

[a] *University College Dublin, Dublin, Ireland*

Abstract

Housing shortages are increasingly cited as barriers to degree completion. Using a 17-year panel of 1,208 U.S. colleges, I estimate the effect of dormitory expansions on retention and graduation via within-college differencing models, confirmed by stacked difference-in-differences. Large expansions do not affect first-year retention but raise four-year graduation by 1–2.3 percentage points; effects grow with expansion size and cohort exposure and reflect higher total completion rather than re-timing. Mechanism analyses find no rent-market spillovers, pointing to campus housing access as the primary channel. Expanding on-campus capacity is a tractable, institution-controlled policy to raise graduation rates.

Keywords: higher education, student housing, college completion, housing supply, human capital

JEL Classification: I20, I21, I23, R31.

*Corresponding author

[1] ORCID: 0000-0002-1416-9515. Email: ciaran.murphy10@ucdconnect.ie.

1. Introduction

Rising tuition and living costs present one of the most significant barriers to higher education attainment in the United States and globally. High housing costs are widely cited as obstacles to effective learning, due to long commutes, the need to work more hours, and reduced attendance.[2] As rental markets tighten, these financial and logistical pressures on persistence and completion intensify. Colleges have a policy lever at their disposal: the expansion of on-campus housing capacity. When institutions build new residence halls or substantially enlarge existing ones, they not only offer more students the convenience and support of campus living but also relieve pressure on the local rental market. The dual impact of such an expansion may translate into improved retention and graduation outcomes, directly through greater campus residency and indirectly through changes in the local housing market.

To my knowledge, this is the first paper to estimate the causal effect of college housing expansions on student outcomes. I investigate how large expansions in on-campus student housing affect two critical measures of student success: first-to-second-year retention and four-year graduation rates. Completion is a particularly important margin: a large literature shows that obtaining a degree is a primary determinant of long-run earnings, labor market attachment, and economic mobility (Arcidiacono 2004; Arcidiacono et al. 2010; Oreopoulos and Salvanes 2011). The results address a practical policy question: can colleges improve student outcomes by strategically investing in housing infrastructure?

The primary empirical strategy is within-college difference models spanning 17 years and 1,208 colleges. Average effects on first-year retention are indistinguishable from zero, while four-year graduation rises by 1-2.3 percentage points, with effects increasing in the size of the expansion. A stacked difference-in-differences applied to a restricted subset of colleges with single, clearly timed capacity jumps confirms these patterns. These results are consistent with a simple explanation: because most institutions already prioritize first-year housing, additional beds primarily relax constraints for upper-year students, improving completion. Comparing cohorts with different degrees of exposure to the additional capacity shows that effects are larger for students exposed for longer, while similar effects on four-, six-, and eight-year graduation rates suggest that expansions raise total completion rather than merely accelerate it. I also examine heterogeneity across institutional subgroups and test for spillovers to nearby colleges; these results provide little systematic evidence of effects operating through local rent markets, pointing instead to a primary role for the benefits of campus residence.

[2] See, for example, https://apnews.com/article/rent-us-college-students-1cc3bc02f25f8a628f828902029bad9b in the U.S., https://www.theguardian.com/money/article/2024/aug/18/have-laptop-can-travel-the-rise-and-rise-of-the-commuter-student in the UK, and https://www.irishtimes.com/ireland/education/2025/02/11/high-rental-costs-are-pricing-out-thousands-of-students-and-risk-increased-college-dropout-rates/ in Ireland.

2. Conceptual Framework & Institutional Background

College housing expansions can affect student success through two distinct channels, depending on whether a student takes up one of the new beds or remains off campus. First, the "campus residence channel": when a new residence hall opens, more students can live on campus. This is expected to improve those students' academic outcomes in several ways that prior work has shown to matter for student success: peer interactions, greater academic structure and monitoring, access to support services, and increased likelihood of attending class (see, for example, Angrist et al. (2009); Dobkin et al. (2010); Sacerdote (2001); Stinebrickner and Stinebrickner (2006) for related evidence on these mechanisms). On-campus residence could also harm academic performance if greater involvement in campus social life crowds out study time. Because U.S. colleges typically prioritize first-years for dorm beds,[3] one would expect an effect on first-year retention rates through the campus residence channel. However, if prioritization means that most first-years are already housed, expansions may instead benefit older students at the margin: the proportion of first-years accommodated remains constant, while additional spaces become available for upper-year students.

Second, the rent market channel: by increasing on-campus capacity, colleges effectively remove a number of students from the external market, decreasing the demand for private rentals near the campus. All else equal, this would be expected to reduce equilibrium rents in the off-campus market, and improve the availability of housing near the campus (reducing search costs and commute times) for students who remain off campus (disproportionately those in upper years). The reduction in rental costs may also allow these students to substitute study for part-time work,[4] while shorter commutes should make attendance easier and free up time for study. Less competition for nearby housing may also leave students in accommodation better suited to studying.

Allocation policies vary widely: many institutions guarantee (or in some cases even require) on-campus housing for first-year students; continuing-student housing is commonly allocated through lotteries with class-year priorities and limited guarantees. Policies frequently reserve a portion of beds for specific programs or groups (such as international students), with the residual stock assigned by lottery among second-years and above.[5] This institutional variation determines which cohorts receive the marginal beds when capacity expands and therefore which outcomes can move first. For example, expansions that primarily ease first-year constraints are

[3] Most of the colleges whose capacity expansions are used as treatment events in the difference-in-differences analysis (Section 9) give first-year students priority access to housing: 86% of these 63 colleges have a first-year housing policy, defined as a requirement to live on campus, a guarantee of a bed, priority in allocation, or designated first-year residence halls. Further information is in Appendix A.

[4] In the 2012 National Postsecondary Student Aid Study, approximately two-thirds of undergraduates report paying at least some of their own rent or mortgage while enrolled (National Center for Education Statistics 2014).

[5] When students do not receive an institutional bed, U.S. colleges typically do not subsidize off-campus rent nor provide special assistance; the outside option is the private market.

likely to be reflected in first-year retention outcomes via the campus residence channel, whereas expansions that free up places for upper-year students will initially affect upper-year outcomes through a combination of the campus residence and rent market channels.

The two channels imply different timing patterns for observable effects. First-year retention effects should arise as soon as new capacity comes online, assuming dorms are fully operational by the start of the academic year. Graduation effects will emerge more gradually, as successive cohorts benefit from expanded capacity across multiple years of study. The impact on four-year graduation rates will therefore depend on both the timing of expansions and the institution's housing allocation policies, which determine how capacity is distributed across student years.

The U.S. undergraduate sector is not uniformly residential: among full-time undergraduates in 2015-16, roughly 32% lived on campus, 44% lived off-campus away from parents, and 24% lived with parents.[6] Patterns differ somewhat outside the U.S. In Eurostudent countries, for example, an average 34% of students live with their parents and only about 15% reside in designated student housing (Hauschildt et al. 2024). Where a large share of students commute from the parental home or rely on the local private rental market, the rent market channel is likely to be relatively more salient than the campus residence channel. The framework is therefore applicable outside the U.S., although the relative magnitudes of effects will depend on local housing institutions.

3. Literature Review

Research on students' residential arrangements dates back more than five decades (see Blimling (1989) for an early review), yet robust causal evidence is limited. Broadly, the empirical literature divides into three strands: (i) single-institution causal studies, (ii) multi-institution evidence using national datasets, and (iii) work on institutional spending that implicitly touches on housing issues. None of these strands, however, evaluates the consequences of expanding residence-hall capacity, the policy lever which is the focus of this study.

Several papers exploit within-college variation to estimate the effect of living on campus. Reynolds (2020) uses a natural experiment at Kent State whereby some students are exempt from a residency rule, and shows that on-campus residence raises grade point average (GPA). Similarly, Araujo and Murray (2010) instrument for students' residence choices at a large public university and find GPA gains of 0.2–1.0 points. In a randomized trial, Perez-Felkner et al. (2022) offer rent-free apartments to community-college students. Participants work fewer paid hours and are more likely to persist to second year. These studies deliver credible estimates of the campus residence channel (the benefits that arise from proximity to peers, classes and campus support) but say nothing about the effects of how many beds a college provides or about spillovers to the local rental market.

[6] See National Center for Education Statistics Profile of Undergraduate Students: https://nces.ed.gov/pubs2019/2019467.pdf.

Moving beyond single campuses, Schudde (2011) combines the Educational Longitudinal Study with propensity-score matching and finds that students who live on campus are three percentage points more likely to enroll in their second year. Descriptive national analyses yield mixed results: positive associations are found by Bozick (2007) and Thompson et al. (1993), but no significant effects remain once covariates are added in Lopez Turley and Wodtke (2010). Although these studies cover hundreds of institutions, their treatment remains residence status, not housing supply.

Two papers that use IPEDS (Integrated Postsecondary Education Data System) data on U.S. colleges address housing indirectly. Webber and Ehrenberg (2010) relate functional expenditures to graduation and persistence, finding that spending on student services matters, but auxiliary spending (including housing) is omitted from their models, likely because auxiliary units (such as residence halls and dining) are operated as self-supporting enterprises whose expenditures are offset by their own revenues. Zong and Davis (2024) include dormitory variables in an attempt to model retention and graduation, but conclude that the effect of capacity is controversial and requires further study.

Guzman (2021) is to the best of my knowledge the only study to exploit national IPEDS panel data to address a question of this nature with causal intent. Guzman analyzes mandatory residency policies for 100 top universities and, after adding institution fixed effects, finds no impact on retention or six-year completion. The bed-to-enrollment ratio appears as a control, but dormitory construction or expansion does not serve as the treatment.

Across all of this work, treatment variation comes from whether students reside on campus, not from increases in the number of beds. The literature therefore provides evidence on the campus residence channel but is silent on the rent market channel (the idea that new dorms ease pressure on local housing markets and benefit off-campus students). By estimating the effect of residence-hall openings, this study identifies the combined effects of both channels and relates directly to a policy lever (expanding capacity) that colleges can control.

4. Data

4.1. Overview

I construct a panel of 1,208 U.S. four-year colleges[7] from 2003 to 2019 (20,536 observations), drawing on three data sources. The core dataset is IPEDS (National Center for Education Statistics 2026), which covers institutions participating in federal student aid programs accounting for almost all public and private colleges. I supplement this with American Community Survey (ACS) rent data (U.S. Census Bureau 2020) to measure local housing market conditions and Census Bureau Building Permits Survey (BPS) data (U.S. Census Bureau 2021) to proxy

[7]Specifically, the sample is restricted to colleges which are classified as four-year in every year of the sample period. When applied to two-year colleges, the analysis is severely underpowered because relatively few two-year institutions offer on-campus housing.

private rental supply. The IPEDS panel is constructed by merging the following IPEDS survey components:

- IPEDS Institutional Characteristics data files: Provide information on each institution's housing capacity, annual room charges, and mandatory residence rules.
- IPEDS Institutional Directory data files: Provide institution-level characteristics including public/private control, urban/rural locale classification, and ZIP code.
- IPEDS Admissions and Test Scores data files: Provide median and quartile ranges for SAT scores and ACT scores for incoming students.
- IPEDS Graduation Rate data files: Provide the percentage of students in each cohort in each institution to graduate within 100% of normal time (4 years), 150% of normal time (6 years) and 200% of normal time (8 years). Graduation rates for each cohort are reported 8 years after enrollment, meaning that this dataset captures cohort graduation outcomes through 2015 entrants.
- IPEDS fall enrollment data files: Provide first-to-second-year retention rates, total number of incoming first-years, and total headcount for each institution.
- IPEDS finance data files: Provide total spending on instruction, academic support, student services, and scholarships in each year.
- IPEDS Student Financial Aid files: Provide information on the percentage of students in each institution in receipt of Pell grants.

I restrict the sample to relatively large institutions (average enrollment $\geq 1,000$, and further dropping another 50 colleges which report an incoming cohort of less than 50 students in some year over the sample period). Institutions that never offer housing during the sample period are dropped. I exclude large institutions which mainly cater to students online (Arizona State Digital Immersion, Southern New Hampshire University, and Liberty University). I drop 2020 onward to avoid COVID-19 confounding.

I conduct basic integrity checks across IPEDS files (e.g. internal consistency of year-to-year capacity, retention and graduation). On this basis, one anomalous capacity report was corrected after manual inspection.[8] I drop 51 colleges that never report a retention or graduation rate in the sample period; these are almost all graduate and professional institutions (for example law, medical and theological schools). I further drop 50 colleges where any entering cohort in the sample period has a reported retention or graduation rate of 0%, for two reasons: (i) concerns over misreporting, and (ii) a lack of within-institution variation in the outcomes of interest over time. I also drop 10 colleges whose campus location changed or whose reporting unit was merged

[8]New Jersey Institute of Technology reports capacity of 1,678 in 2012, 11,806 in 2013, and 1,657 in 2014. 2013 capacity is manually imputed as the mean of 2012 and 2014 capacities. Twenty-nine other colleges report zero capacity for one or two years, with positive capacity before and after. I keep these observations because some may be genuine closures. Replacing each gap by a linear interpolation has no material effect on results.

with another institution during the sample.[9] Finally, I retain only colleges that appear in all 17 years of the sample (a balanced panel), which drops a further 30 colleges. Results are robust to all of these sampling restrictions. Classical mismeasurement in the treatment (capacity) would attenuate estimates toward zero, so any positive effects despite potential noise likely understate true impacts. In addition, reported decreases may sometimes reflect ownership or accounting changes (e.g., transfers to affiliated entities) rather than physical closures, which would also bias estimates toward zero. IPEDS non-response arises primarily from administrative reporting gaps (such as institutions missing a filing deadline) rather than underlying student outcomes, so any missingness is plausibly missing at random. Following these sampling restrictions, the sample captures the vast majority of U.S. four-year institutions. For instance, in 2019, it represented 69% of the total undergraduate enrollment at four-year institutions as reported by IPEDS.

4.2. Descriptive Statistics

Table 1 shows summary statistics for the variables used throughout the analysis.

Table 1: Descriptive Statistics for Key Institutional Variables

Variable	Mean	SD	Min	Max	N
Room Capacity	2104.37	2215.22	0.00	19446.00	20531
Total Enrollment	7827.78	9293.61	359.00	96211.00	20533
Four-year Graduation Rate (%)	39.59	21.93	0.00	96.00	14399
Six-year Graduation Rate (%)	55.97	18.01	0.00	100.00	14399
Eight-year Graduation Rate (%)	57.99	17.23	0.00	100.00	14399
First-Year Retention Rate (%)	76.51	11.07	19.00	100.00	20076
Beds per Student	0.40	0.25	0.00	3.54	20137
Beds per First-Year	2.34	1.01	0.00	10.59	20123
First-Year Cohort Size	1104.95	1254.07	50.00	10099.00	20476
Selectivity Index	1071.10	132.96	510.00	1555.00	17928
Annual In-State Tuition ($)	18466.12	13369.75	0.00	58615.00	20472
Instruction Spending ($mn)	86.24	182.31	0.02	3126.10	20073
Academic Support Spending ($mn)	23.62	60.20	0.00	1619.43	20073
Student Services Spending ($mn)	17.06	23.29	0.00	387.40	20073
Scholarships Spending ($mn)	7.01	16.49	0.00	208.84	20073
Mandatory On-Campus Housing	0.21	0.41	0.00	1.00	19326

Notes:
Sample: IPEDS, 2003-2019. Unit of observation is institution-year. Annual in-state tuition is the published in-state charge, in nominal dollars. Mandatory On-Campus Housing is a binary indicator equal to 1 if the institution requires all first-year students to live on campus, and 0 otherwise.

[9]Specifically, the following colleges were dropped on this basis: Albany State University, Atlanta Christian College (now Point University), Baker College of Flint, California College of the Arts, Davenport University, Indiana Wesleyan University-Marion, Nyack College, The Sage Colleges, Urbana University, and the University of Colorado Denver.

The number of beds offered per college ranges from 0 (as although colleges are dropped from the sample if they never offer housing, some do not offer housing in every sample year) to 19,446 (the capacity in Michigan State University in 2004). Housing capacity refers to institutionally-controlled residential facilities; institutions report whether they provide institutionally controlled housing and the associated capacity. Throughout, I use "on-campus housing" as shorthand for institutionally-owned or -controlled beds, which may be located on or near the main campus. Total enrollment ranges from just 359 (Truett McConnell College in 2004) to 96,211 (Grand Canyon University in 2019), with a mean of 7,828. On average, 77% of first-year students persist to second year, while just 40% of students in four-year courses graduate within normal time.[10] Unfortunately, colleges do not report retention rates on a yearly basis after the first year. 76.51% of students proceed to the second year, 39.59% graduate within 4 years, and a further 18.40% graduate in the following four years. 18.52% of students are therefore unaccounted for, and drop out some time after the first-to-second-year transition (or potentially finish their four-year degree after more than eight years).

IPEDS reports capacity (number of beds), but not occupancy or the share of students living in institution-controlled housing. Changes in capacity are therefore shocks to the potential number of on-campus residents. As descriptive context, the average college has enough beds to cater to all first-year students, but significantly fewer beds than the total number of enrolled students.[11] The absolute number of beds in the college (the number of institution-owned beds reported in IPEDS) is used as the measure of capacity in the main analysis, as it represents the cleanest, most exogenous treatment definition: it captures the physical opening (or closure) of dormitory space rather than endogenous shifts in utilization. The beds per student and beds per first-year ratios are not used as treatment definitions, as the estimated effect of an expansion may be confounded by the effect of an increase in total enrollment or first-year cohort size. Instead, total enrollment and incoming first-year cohort size are used as covariates.

The selectivity index is constructed based on test scores of incoming first-year students. As some colleges report SAT scores and others report ACT scores, ACT scores are converted to SAT scores using ACT concordance tables (ACT Education Corp. 2018). The SAT totals used correspond to the combined verbal and math sections, on a 1600-point scale. A single measure of selectivity is then constructed as the mean of the 25th percentile and 75th percentile of SAT scores if available, and converted ACT scores if not. The index is therefore a proxy for

[10] Graduation rates are cohort measures; denominators are entering cohorts, not conditional on second-year progression. The graduation rates for this sample following the imposition of the previously outlined sampling restrictions remain in line with national averages reported by the National Center for Education Statistics: https://nces.ed.gov/FastFacts/display.asp?id=40.

[11] There are 311 college-year observations in 63 colleges where beds outnumber total enrollment. 57 of these 63 colleges, at some stage in the sample period, have a requirement for all first-year students to live on campus. The phenomenon of beds outnumbering students typically arises after the removal of such a requirement, which is generally followed by a decrease in the number of beds available on campus. Removing these institutions from the sample has no material effect on results.

cohort academic performance that varies over time within colleges; this variable captures shifts in entrant quality even when GPA distributions and application essays are unobserved.

Colleges report annual spending in various areas, including instruction, academic support, student services and scholarships; instruction spending is typically the highest. On-campus housing is mandatory for some undergraduate students in 21% of college-year observations. Such policies have become less common over the sample period; 38.2% of institutions reported having such a policy in 2004, compared to just 3.4% in 2019 (although this partly reflects a change in IPEDS reporting instructions in 2012, which tightened the definition of a "mandatory" policy).

To contextualize on-campus room charges in the local housing market, I compare IPEDS annual room charges to local rents measured from the American Community Survey (see Appendix B for full details). The median on-campus room charge over 2015-2019 is approximately $674 per month. This is below the median one-bedroom rent in the census tract containing the campus ($709 per month), suggesting on-campus housing provides a saving relative to sole-occupancy rental. However, it is substantially above the per-person cost of sharing a two-bedroom unit ($447 per person per month) - a more realistic off-campus alternative for most students - and at roughly 85% of institutions in the sample, on-campus charges exceed this per-person benchmark.

4.2.1. Descriptive Evidence on Housing Supply Expansions

Figure 1 shows the housing capacity over time for 12 randomly selected institutions in the sample. In a number of cases, the room capacity is steady with some sharp increases. Notably, however, some of the colleges exhibit sharp decreases as well as increases. It is also notable that in many cases, sharp increases in room capacity occur over a period of more than 1 year.

Table 2 presents descriptive statistics regarding which colleges tend to experience changes in capacity, where a change in capacity is defined as an expansion or reduction in total capacity by at least 250 beds in one year. 567 institutions in the sample increased their capacity by 250 beds in one year, while 354 reduced by 250 beds in one year (of these 354, 318 had also had an increase of 250 beds within the sample period, indicating that many reductions in capacity are temporary and followed by corresponding expansions). Of the 567 colleges that expanded, 335 had two or more expansions during the sample period. The average enrollment of the institutions which increased/decreased their capacity over the sample period is higher than those which did not (this is to be expected due to the definition of treatment as an absolute increase/decrease). Colleges which experienced expansions tended to be public (64% of expanding colleges, compared to 20% of non-expanding colleges) and urban (78% of expanding colleges, compared to 68% of non-expanding colleges). There were no differences in the mandatory housing policies for first-years in expanding and non-expanding colleges, nor do these policies tend to change at the time of an expansion or reduction itself.

Figure 1 and Table 2 make two facts clear. First, capacity paths are, in many cases, erratic. Many colleges experience multiple large increases and decreases. Many expansions are preceded or followed by a contraction of a similar magnitude. These temporary closures may be due

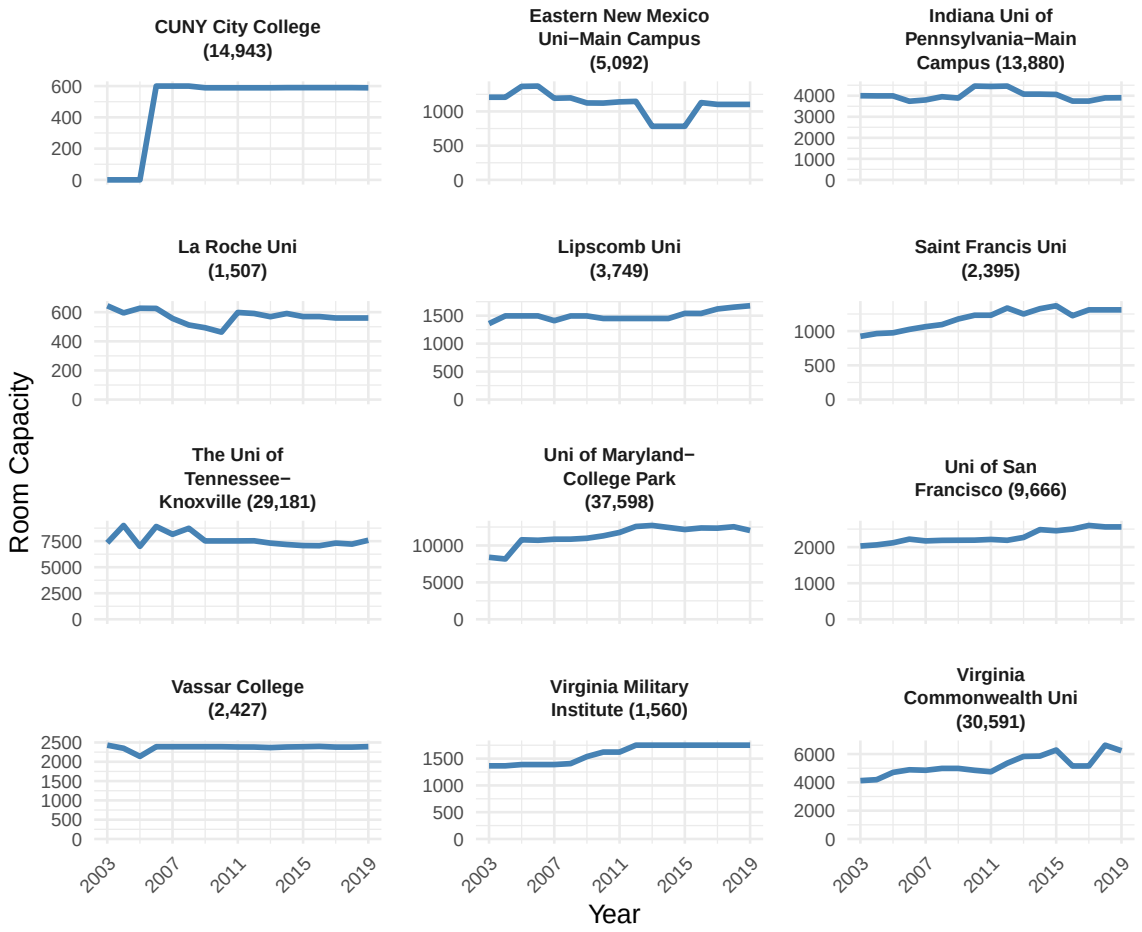


Figure 1: Housing capacity over time for 12 randomly selected institutions. Average enrollment over the sample period shown in parentheses.

Table 2: Summary of Institutional Characteristics by Expansion/Reduction Status

	Institutional Characteristics				
Treatment Status	Number	Avg Enroll	Public	Urban	Mandatory Residence
Did Not Expand 250+	641	3344	20%	68%	39%
Expanded 250+	567	12896	64%	78%	37%
Did Not Reduce 250+	854	5279	31%	71%	36%
Reduced 250+	354	13975	64%	76%	43%
Expanded & Reduced	318	14933	65%	76%	43%
Neither Expanded Nor Reduced	605	3216	18%	67%	39%
Expanded 250+ (More than once)	335	16450	70%	83%	38%

Notes:
Table shows descriptive statistics for colleges that expanded (adding 250 beds in a year) or reduced capacity (decreasing by 250 beds). Mandatory residence is defined as requiring first-year students to live on campus in 2004 (the first year in which this information is available).

to maintenance, or scenarios whereby new dorms are built to allow for the closure of an old facility. Second, treatment propensities are heterogeneous. Colleges that add 250 beds in a single year are, on average, larger, more likely to be public, and more likely to be located in urban

areas. These features complicate a conventional staggered difference-in-differences design: units receive treatment at different times (and sometimes repeatedly), and "never-treated" colleges are structurally different. It is therefore difficult to clearly label colleges as treated, with a distinct pre-treatment and post-treatment period, and equally challenging to define appropriate control units when their housing capacities also fluctuate significantly. For example, defining treatment as any increase above a given bed-threshold would classify many colleges as treated multiple times, while others just below the cutoff might still undergo substantial, and sometimes sharp, changes in capacity. A simple binary treatment definition fails to capture this complexity and risks misrepresenting the true variation in housing trajectories.

5. Methodology

5.1. Identification Strategy

Because standard two-way-fixed-effects staggered difference-in-differences assumes a binary treatment, and that each unit is treated once and remains treated, the IPEDS housing data lend themselves poorly to that framework. The primary empirical strategy therefore employs institution-level differencing. Differencing eliminates all time-invariant college characteristics (the chief source of bias given the heterogeneity documented above) and is particularly suited to data of this nature: it exploits the rich within-institution variation (increases, decreases, and periods of stability) without requiring an explicit control group. As a robustness check, a stacked difference-in-differences applied to a restricted subset of colleges with single, discrete capacity jumps is described and estimated in Section 9; this design validates the main results in a cleaner potential-outcomes setting, but without the large sample size and external validity of analysis based on the entire sample.

There are two endogeneity issues: colleges sorting into treatment, and students sorting into treated colleges. Regarding colleges sorting into treatment, both estimators eliminate all time-invariant differences across colleges by construction. The relevant question for identification is therefore not whether colleges that choose to expand are different from colleges that do not - Table 2 shows that they are - but whether the periods in which housing comes on stream are systematically different from other periods at the same institution in ways that independently affect student outcomes.

The central identification argument rests on the long and largely exogenous production lag between the decision to expand and the delivery of new beds. Building a new residence hall is not a short-run decision: the process typically involves capacity planning, capital budgeting and financing, architectural design and regulatory approval, and a construction phase subject to its own delays. From initial approval to completion, timelines commonly span several years. As a result, conditional on the decision to expand, the exact timing of delivery is unlikely to be systematically related to short-run fluctuations in institutional quality or student composition which could independently affect student outcomes. Identification therefore comes from comparing

adjacent cohorts within the same institution, where the expansion decision is typically already fixed but cohorts differ in whether the additional housing has actually come on stream. This isolates variation in realized capacity that is plausibly orthogonal to other short-run shocks.[12]

Nevertheless, there may be residual concerns about concurrent institutional changes. Housing expansions may coincide with broader initiatives aimed at improving student outcomes, or may crowd out other forms of expenditure. I address these concerns by controlling for changes in spending on instruction, academic support, student services, and scholarships. These controls capture contemporaneous institutional investments.

The second concern is that changes in housing capacity coincide with changes in the scale or composition of enrolled students. Colleges may expand capacity alongside increases in enrollment, or improvements in housing availability may make the institution more attractive to prospective students. I address changes in scale by controlling for cohort size and total enrollment, ensuring that estimates are not driven by mechanical changes in student numbers. Changes in composition are addressed using a selectivity index based on SAT and ACT distributions of entering cohorts. As an additional check, Appendix C reports event studies for a wider set of cohort-composition variables around expansion events, including Pell receipt, gender, race/ethnicity, recent high-school graduation, and in-state status. Moreover, given the unpredictable multi-year lag between planning and delivery, it is unlikely that entering cohorts can respond to the precise timing of housing completion. Appendix D shows that, at least with the data available, the year in which housing comes on stream cannot be predicted from an institution's own recent trends in enrollment, spending, selectivity, tuition, or retention and graduation rates. The covariates referenced above are described in detail in Section 5.2, and additional robustness checks are carried out in Section 8.

5.2. Difference Models

I exploit within-college variation by estimating difference models. Let Δ_1 and Δ_4 denote one-year and four-year temporal differences, respectively. Outcomes are measured in percentage points. I begin with a linear benchmark and then allow for non-linearities using bins. For

[12]For example, the residence hall at the College of Staten Island (CUNY), the college's first on-campus housing, took six years and four months from the first board action in April 2007 to opening in August 2013. In June 2008 CUNY's trustees approved the development agreements for 607 beds, with financing planned for Fall 2008 and opening in Fall 2010 (see https://policy.cuny.edu/policy/wp-content/uploads/sites/6/page-assets/minutes/Minutes06232008.pdf). Financing was delayed "due to poor market conditions", and in September 2011 the trustees re-approved a smaller project of 458 beds (see https://policy.cuny.edu/wp-content/uploads/sites/6/page-assets/minutes/2011Minutes.pdf). Bonds were issued in March 2012 (see https://comptroller.nyc.gov/wp-content/uploads/documents/Housing_Development_Financial_Statements_FY_2013__Year_Ended_Oct._31_2012_.pdf) and the hall opened with 454 beds, three years after the date announced in 2008 (see https://csitoday.com/2013/10/new-dorms-at-the-college-of-staten-island-create-a-new-world-for-student/).

first-to-second-year retention, which is observed after one year, I estimate:[13]

$$\Delta_1 \text{Retention}_{ic} = \alpha + \beta^{(1)} \frac{\Delta_1 B_{i,c}}{100} + \Delta_1 \mathbf{X'}_{ic}\,\gamma + \Delta_1 \mathbf{Z'}_{ic}\,\theta + \Delta_1 \varepsilon_{ic} \tag{1}$$

Here $\Delta_1 \text{Retention}_{i,c} = \text{Retention}_{i,c} - \text{Retention}_{i,c-1}$ is the change in the first-year retention rate between cohort c and cohort $c-1$, and $\Delta_1 B_{ic} = B_{ic} - B_{i,c-1}$ is the change between the number of beds available when cohort c was in its first year and when cohort $c-1$ was in its first year. The coefficient $\beta^{(1)}$ is interpreted as the percentage-point change in retention per +100 beds. Because very small adjustments (and contractions) may have different effects from large, sustained expansions, I also estimate a flexible binned specification:

$$\Delta_1 \text{Retention}_{ic} = \alpha + \sum_j \beta_j^{(1)}\, \mathbf{1}\{\Delta_1 B_{i,c} \in \mathcal{B}_j\} + \Delta_1 \mathbf{X'}_{ic}\,\gamma + \Delta_1 \mathbf{Z'}_{ic}\,\theta + \Delta_1 \varepsilon_{ic} \tag{2}$$

$\mathcal{B} = \{B_1, ...B_6\}$ denote six mutually exclusive bins for capacity changes: a large increase ($\Delta_1 B > 750$), a moderate increase ($500 < \Delta_1 B \leq 750$), a small increase ($250 < \Delta_1 B \leq 500$), a small decrease ($-500 \leq \Delta_1 B < -250$), a moderate decrease ($-750 \leq \Delta_1 B < -500$), and a large decrease ($\Delta_1 B < -750$). For each bin j, the model includes the indicator $\mathbf{1}[\Delta_1 B_{i,c} \in \mathcal{B}_j]$; the omitted category is $|\Delta_1 B_{ic}| \leq 250$. The coefficients $\beta_j^{(1)}$ therefore capture the effect of being in bin j relative to small/no change. The non-linear binning of capacity allows the effect of large expansions or contractions to deviate from that of minor adjustments, a flexibility that would be lost in a purely linear differencing specification.

$\Delta_1 \mathbf{X}'_{i,c}$ contains covariates which control for the possibility that the college may make other changes which affect student outcomes. Specifically, I include the one-year differences in total spending on instruction, academic support, student services and scholarships. $\Delta_1 \mathbf{Z}'_{i,c}$ contains one-year differences in covariates which control for cohort composition effects: the selectivity index constructed based on SAT and ACT scores and the size of the enrolling cohort c.

Standard errors are clustered by college.[14] In the main text, I present two fully specified models - one including the size of the entering first-year cohort and one in which cohort size is replaced by total institutional enrollment (which includes students of all levels) for robustness. Including these covariates in separate specifications avoids potential multicollinearity concerns; results are similar when both are included jointly. Appendices E and F report additional specifications that incrementally add controls.

Graduation is analyzed with four-year differences that match the normal time to degree.

[13] The intercept α captures a common linear trend in levels; in practice, I also estimate a variant that replaces α with a full set of cohort-year dummies (equivalent to first-differenced year fixed effects), yielding nearly identical results.

[14] I also rerun the main results, clustering by core-based statistical area (CBSA, the Census definition of a local urban labour and housing market), which allows for correlation among colleges in the same local market. There are no material changes to results.

Because capacity may influence graduation through multi-year exposure, I define treatment at the cohort level using the average number of beds available over the four academic years a cohort is enrolled. For the cohort entering in year c, exposure is $\overline{B}_{ic} = \frac{1}{4}\sum_{k=0}^{3} B_{i,c+k}$. Treatment is the change in this four-year average between cohort c and cohort $c-4$ in the same institution ($\Delta_4\overline{B}_{ic} = \overline{B}_{ic} - \overline{B}_{i,c-4}$). Spending variables enter as four-year differences in cohort-average exposure ($\Delta_4\overline{X}_{ic}$) to mirror treatment timing, whereas selectivity and cohort size (or total enrollment) are defined at cohort entry and therefore enter as four-year start-to-start differences ($\Delta_4 Z_{ic}$).[15] The linear benchmark model is:

$$\Delta_4 \,\text{Graduation}_{ic} = \alpha + \beta^{(4)} \frac{\Delta_4\overline{B}_{ic}}{100} + \Delta_4\overline{\mathbf{X}}'_{ic}\,\gamma + \Delta_4\mathbf{Z}'_{ic}\theta + \Delta_4\varepsilon_{ic}. \tag{3}$$

The coefficient $\beta^{(4)}$ is interpreted as the percentage-point change in the four-year graduation rate per +100 average beds over the cohort's four years. For flexibility, I also estimate the binned analogue, using the same bins as in the retention models:

$$\Delta_4 \,\text{Graduation}_{ic} = \alpha + \sum_j \beta_j^{(4)}\,\mathbf{1}\Big\{\Delta_4\overline{B}_{ic} \in \mathcal{B}_j\Big\} + \Delta_4\overline{\mathbf{X}}'_{ic}\,\gamma + \Delta_4\mathbf{Z}'_{ic}\theta + \Delta_4\varepsilon_{ic}. \tag{4}$$

Table 3 reports the distribution of $\Delta_1 B_{ic}$ and $\Delta_4\overline{B}_{ic}$ and the share of observations in each bin.

Table 3: Observation counts and percentage shares by capacity-change bin (Graduation and Retention)

Bin	N (Grad)	Share, % (Grad)	N (Ret)	Share, % (Ret)
Increase >750 beds	463	5.1	261	1.7
Increase 500–750 beds	396	4.4	218	1.4
Increase 250–500 beds	1068	11.7	717	4.7
Change within ±250 beds (reference)	6839	75.2	13669	88.9
Decrease 250–500 beds	212	2.3	298	1.9
Decrease 500–750 beds	69	0.8	82	0.5
Decrease >750 beds	52	0.6	124	0.8

Note:
This table reports the number of institution-year observations falling into each capacity-change bin for both samples. Shares are percentages of total observations within each sample. Positive values indicate increases in housing capacity; negative values indicate decreases.

[15] The cohort-average spending measure includes years c through $c+3$, which overlap with the period over which the graduation outcome is determined. If housing expansions cause concurrent changes in spending, or if spending responds to intermediate student outcomes, these controls may absorb part of the causal effect or introduce collider bias. Two robustness checks address this concern: the start-to-start specification (Section 8.3), in which spending controls are measured at entry only and so cannot respond to post-entry outcomes, and the specifications in Appendices E and F that add controls progressively, which show the estimates are not sensitive to their inclusion.

Standard errors in the graduation models are again clustered at the college level. This design aligns the timing of treatment with the outcome's horizon: graduation is a four-year outcome, so each cohort's full four-year exposure enters the comparison. Mid-program expansions or contractions are therefore included in the treated cohort's average (and in the prior cohort's average if they occurred earlier), avoiding the attenuation that arises when only entry-year capacity is used. An alternative "start-to-start" specification is estimated in Section 8.3.

Identification relies on the absence of contemporaneous, unobserved shocks that are correlated with discrete capacity changes and with student outcomes. Three potential threats merit discussion. First, contemporaneous institutional initiatives that coincide with housing expansions could bias results. This concern is mitigated by the inclusion of detailed spending controls (instruction, academic support, student services, and scholarships), which capture most parallel investments. Second, local shocks may simultaneously affect both construction and outcomes. For example, labor-market changes may influence colleges' ability to build while also altering students' opportunity costs of study; zoning or regulatory shifts may affect both the feasibility of new residence halls and the affordability of private rentals. Granular financial and selectivity controls limit the scope for such shocks to drive the results. Third, colleges might build in anticipation of "better" students, or expansions might attract stronger applicants. Both threats seem implausible given the long planning horizons of residence hall projects; residence hall openings typically follow years of design, financing, and construction. Such horizons make it unlikely that expansions respond to the expected quality of a single incoming cohort. Moreover, the selectivity index (based on SAT/ACT scores) flexibly absorbs any shifts in student composition that could bias the estimates. Some further threats remain; for example, a college might increase spending in an area not covered by the controls (for example, by constructing a parallel capital project such as a new student center), or might change its allocation policies around the time of an expansion (for example, prioritizing weak/strong students for beds). The difference-in-differences analysis in Section 9, which operates in a cleaner causal framework, further supports the robustness of these estimates to potential confounding.

6. Results

6.1. Continuous Difference Models

Table 4 reports the results of the fully specified models showing the effects of a 100-bed capacity increase on retention and graduation rates. The unit of observation is the one-year change in retention, or the four-year change (cohort c versus $c-4$) in graduation. Each coefficient can therefore be interpreted as the percentage point increase in the retention or graduation rate associated with a 100-bed increase in capacity.

Given that first-year retention is measured as an integer between 0 and 100, the point estimates in the first two columns are small. The coefficients are essentially zero and precisely estimated, with confidence intervals that rule out meaningful retention effects. This is what one

Table 4: Linear effects per +100 beds

	Retention		Graduation	
Variable	Retention (1)	Retention (2)	Graduation (1)	Graduation (2)
Per +100 beds (pp)	0.000	-0.002	0.119***	0.101***
	(0.006)	(0.006)	(0.030)	(0.027)
Controls				
Spending controls	✓	✓	✓	✓
Selectivity	✓	✓	✓	✓
Cohort size	✓		✓	
Total enrollment		✓		✓
Sample size				
	15,369	15,369	9,099	9,099

Notes:
Coefficients represent the percentage point change associated with a 100-bed capacity increase. Retention uses 1-year differences; graduation uses four-year cohort-average exposure. Robust SEs clustered by institution; *, **, *** indicate significance at 10%, 5%, and 1%, respectively.

would expect given the institutional practice of housing virtually all incoming first-years: adding a few hundred beds rarely alters first-year housing assignments. When expansions occur, any benefits are likely to affect upper years at the margin.

Columns 3 and 4, which present the results of the four-year difference models for graduation outcomes, show a clear positive effect of capacity expansions on degree completion.[16] An increase of 100 beds raises four-year graduation by roughly 0.1 percentage points, significant at the 1% level. The results are robust to replacing the four-year change in cohort size covariate with four-year change in total enrollment. The full set of specifications in Appendix E show that adding spending, selectivity and cohort size controls does not materially affect the estimates, indicating that the completion effects are not driven by changes in enrollment patterns or institutional spending but by genuine capacity relief.

6.2. Binned Difference Models

The results show that, on average, on-campus housing expansions do not affect first-years but do affect upper-year students across the entire sample. However, there are a number of reasons to suspect that these average effects per 100 beds may not be reflective of the true effects of large expansions. First, expansions and contractions may not have symmetric effects. Second, small

[16]The graduation analysis has fewer observations for three reasons. First, graduation data are only available for cohorts entering through 2015, while retention data extend through 2019. Second, the four-year difference requires non-missing graduation rates in both the focal cohort (t) and the comparison cohort ($t-4$), which is binding at the start of the panel. Third, all spending controls must be non-missing over both cohort windows, causing additional attrition where finance data are sparse; this affects 1.0% of treated and 2.4% of control institution-years. Re-estimating with spending imputed from adjacent years rather than dropped yields almost identical results.

incremental changes in capacity, which the descriptive statistics in Section 4.2 show are common, may have minimal effects on overall student outcomes, which would dampen the estimates in the continuous model. Third, such small changes may not represent actual building events, which is the policy lever of interest in the analysis. For this reason, I estimate binned models that capture the effect of discrete capacity expansions and contractions of various sizes, comparing them to periods of little to no capacity changes (absolute change of at most 250 beds). The results for the preferred specification (the full model with the cohort size covariate) are shown in Figure 2.

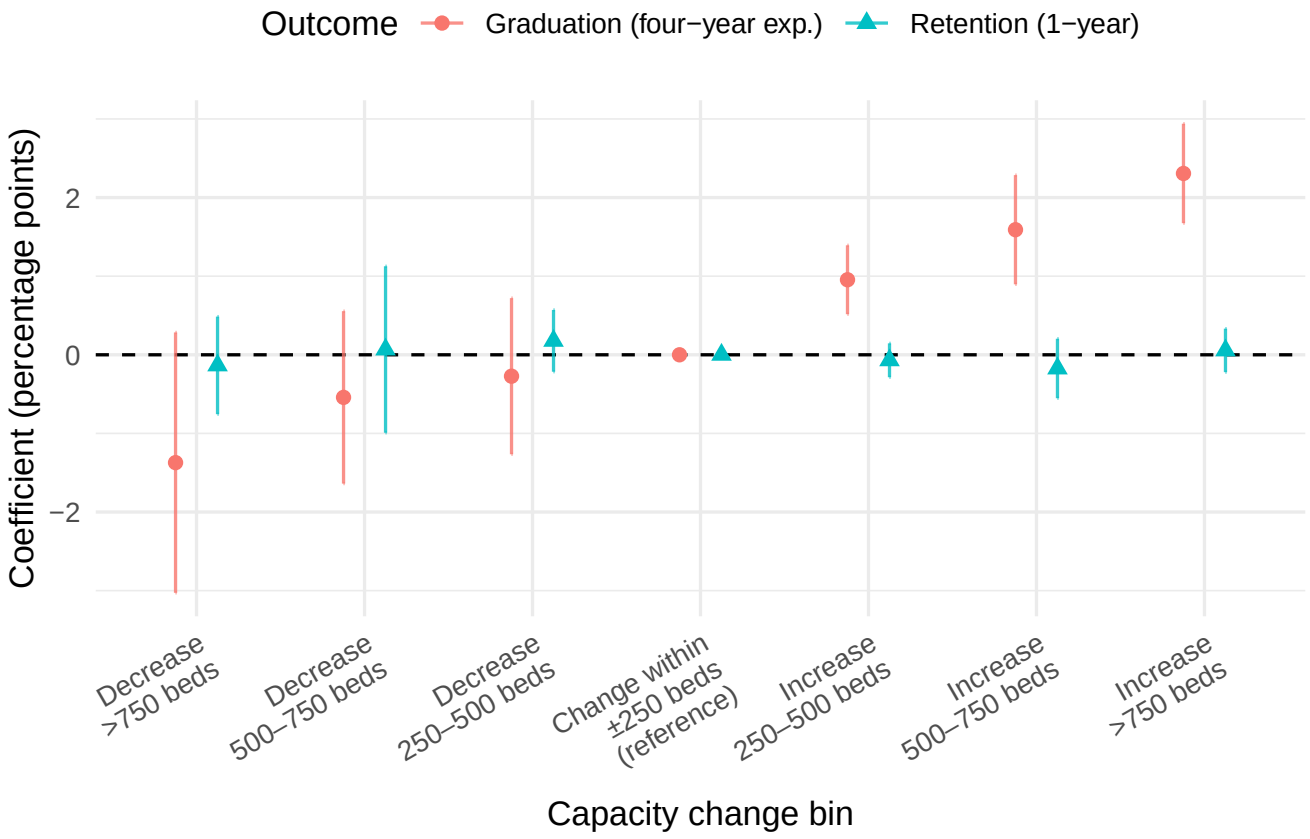


Figure 2: Differencing estimates of the effect of on-campus housing capacity changes on first-year retention and four-year graduation, by size of capacity change. The reference (omitted) category is changes of at most 250 beds. All estimates include controls for institutional spending, selectivity, and cohort size. Points are estimated effects; bars are 95% confidence intervals.

The results of the binned models for retention outcomes are consistent with the continuous model; the estimated coefficients for capacity expansions and contractions of all magnitudes are close to zero. A clearer story emerges regarding graduation effects, however, when bins are used to allow for nonlinearities. First, the effects of contractions in capacity, although directionally negative as expected, are small and statistically indistinguishable from zero. Very large contractions of over 750 beds do show a negative estimate which is large in magnitude, but this coefficient is based on just 52 observations and is not statistically significant. There are a

number of possible explanations for this; it may reflect the fact that closures are often temporary and do not affect students for the duration of their time in college. Alternatively, residences are closed because they are in a condition so poor that students do not in fact benefit from living in them. Whatever the reason, it appears that the effects of dorm closures and openings are not symmetric, and the effect of closures attenuates the results in the continuous model.

The binned model also shows that large expansions have a large and significant effect on graduation outcomes. Increases of 250-500 beds are associated with an increase in graduation rates by roughly 1 percentage point, increases of 500-750 beds are associated with an increase of roughly 1.6 percentage points, and increases of over 750 beds are associated with an increase of over 2 percentage points. These graduation gains are substantially higher than the gains implied by the continuous linear model, and emphasize the importance of allowing for some degree of nonlinearity. The differences can be explained both by the fact that contractions in capacity have practically no effect on outcomes, and by the minor effects of small expansions, which are confined to the reference category in the binned model.

The results show that large discrete expansions in on-campus beds yield large positive effects on four-year graduation rates. This indicates that, even if first-year housing is saturated, boosting supply eases constraints for upper-year students, improving stability and persistence to degree. The retention estimates, meanwhile, imply that incoming cohorts are almost always accommodated, so additional beds rarely alter first-year outcomes. Instead, the graduation gains point to upper-year students as the main beneficiaries: capacity constraints tighten as cohorts progress, and easing them raises completion. The results imply that colleges seeking to improve first-year retention should focus on non-housing factors (such as financial aid, academic support, or orientation), particularly if, as in many colleges, first-years are already sufficiently housed. Colleges aiming to boost graduation rates, on the other hand, may find that targeted investments in upper-year housing capacity pay dividends in timely degree completion.

6.3. Six- and Eight-Year Graduation Rates

An increase in the four-year graduation rate could reflect two distinct effects: students completing on time who would otherwise have finished later (an intensive margin effect on timing), or students completing who would otherwise not have graduated at all (an extensive margin effect, and arguably a stronger welfare claim). As shown in Table 1, 40% of students graduate within four years, a further 16% complete in years five and six, and a further 2% in years seven and eight. Given the low proportion of students graduating in years seven and eight, it is reasonable to assume that the eight-year rate is close to the practical ceiling for total completions. An effect on the four-year rate that does not extend to the eight-year rate would therefore indicate that housing expansions accelerate completion rather than increase it; a significant effect at the eight-year horizon would imply an extensive margin effect: genuine gains in total graduation.

To examine this, I repeat the main binned graduation analysis using the six-year (150% of normal time) and eight-year (200% of normal time) completion rates. The specification is

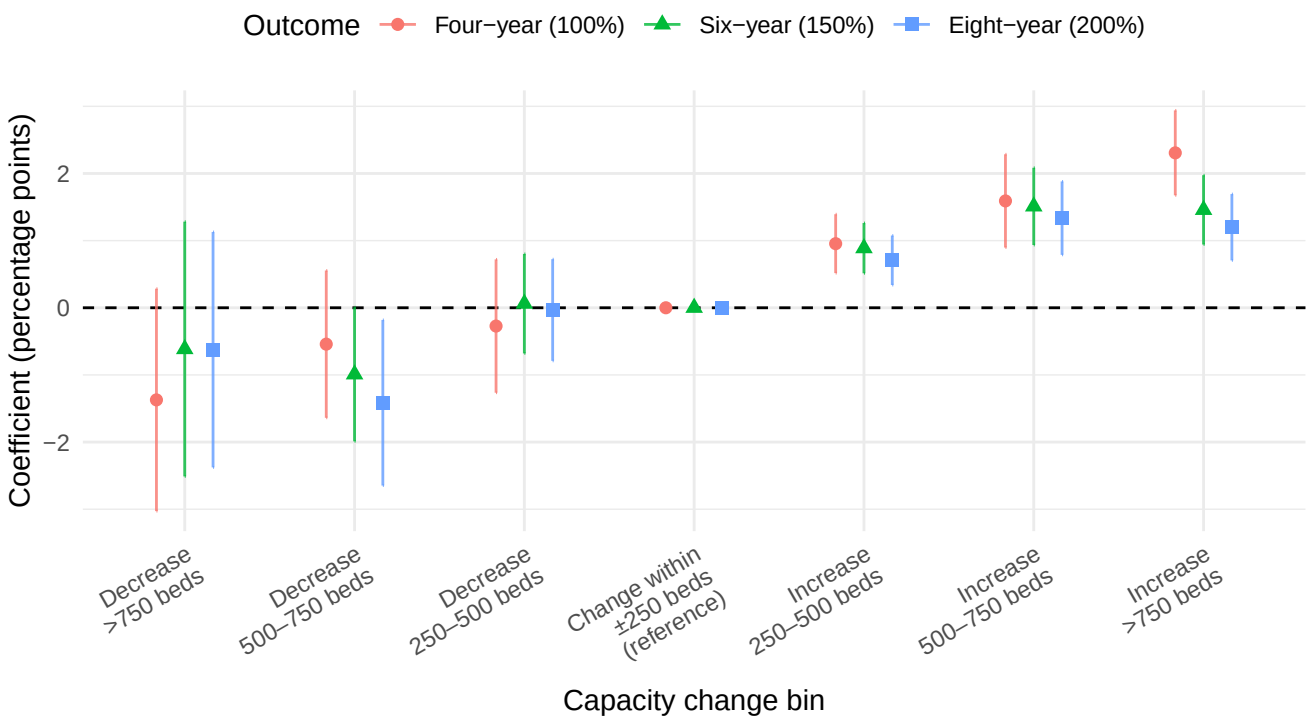


Figure 3: Graduation rate coefficients by capacity-change bin for four-, six-, and eight-year outcomes. All estimates include spending, selectivity, and cohort-size controls. Points are estimated effects; bars are 95% confidence intervals. The omitted category is changes in capacity of at most 250 beds.

otherwise unchanged; I continue to use the average exposure across a student's first four years of enrollment. The six-year rate is also the standard outcome measure used in the IPEDS-based graduation literature, and higher average graduation rates in many other countries make six- and eight-year rates more comparable for international external validity.

As shown in Figure 3, the effects on six- and eight-year graduation rates are positive and statistically significant for all three expansion bins, and similar in magnitude to those on four-year completions for expansions of up to 750 beds. For the largest expansions (over 750 beds) the six- and eight-year estimates are smaller than the four-year estimate, which suggests some timing dynamic in this bin: part of the four-year gain may reflect students completing more quickly rather than additional completions. Overall, however, the effects persist at six and eight years, which indicates that the main results are driven by students who would not otherwise have graduated (an extensive-margin effect), and not only by students graduating more quickly.

7. Mechanisms

Section 2 outlined two channels through which housing expansions may improve student outcomes: a campus residence channel, operating through the direct benefits of on-campus living, and a rent market channel, operating through reduced competition for private rental housing near the campus. The data record how many beds an institution has but not who occupies them, so I cannot observe whether the gains are concentrated among students who move on

campus, as the campus residence channel implies. The analysis instead tests for evidence of the rent market channel.

Two approaches are taken. The first examines heterogeneity in the main effects across subgroups of institutions - for instance, by local rent levels, urban setting, and the socioeconomic composition of the student body. Systematic differences in the expected directions would be consistent with a rent market mechanism, though such heterogeneity is not definitive, since it may also reflect variation in campus residence effects. The second approach provides a more direct test: if an expansion at one institution reduces local rents, students at nearby institutions should also benefit. Such spillovers are inconsistent with a pure campus residence channel and would constitute direct evidence for the rent market mechanism.

7.1. Subgroup Analysis

Carrying out the analysis across subgroups of institutions provides insight into the mechanisms underlying the main results. I examine heterogeneity by socioeconomic composition, local rents, the ratio of on-campus room charges to local rents, urban location, sector, on-campus room charges, tuition, in-state enrollment shares, and whether the institution ever required students to live on campus. These dimensions capture variation in students' exposure to local housing markets and in the extent to which on-campus housing may substitute for private rental supply.

If housing expansions operate primarily through a rent market channel, effects should be strongest where students are more exposed to local rental markets and where off-campus housing is relatively costly. This would point to larger effects among colleges in markets where private rents are high relative to on-campus charges, and among colleges serving more economically disadvantaged students. Local rental market pressure is measured using the ACS per-person cost of sharing a two-bedroom unit in a one-mile ring around the campus (see Appendix B for further details), and the ratio of the on-campus room charge to this per-person rent is used as a further split. Other subgroup splits (urban versus rural location, public versus private sector, in-state enrollment shares, and institutional price levels) provide additional evidence on the institutional settings in which on-campus housing expansions may matter most. Full details on the construction of subgroup variables and the definition of each split are provided in Appendix G.

Figure 4 reports graduation results across these subgroups. Overall, the estimated effects are broadly similar across institutional contexts. There is some variation in magnitudes across panels, but no coherent pattern indicating systematically larger effects in the settings where a rent market channel would be expected to be strongest. The subgroup results therefore do not provide systematic evidence in favor of a rent market channel. Instead, they suggest that the effects of housing expansions are relatively uniform across a broad range of institutional environments.

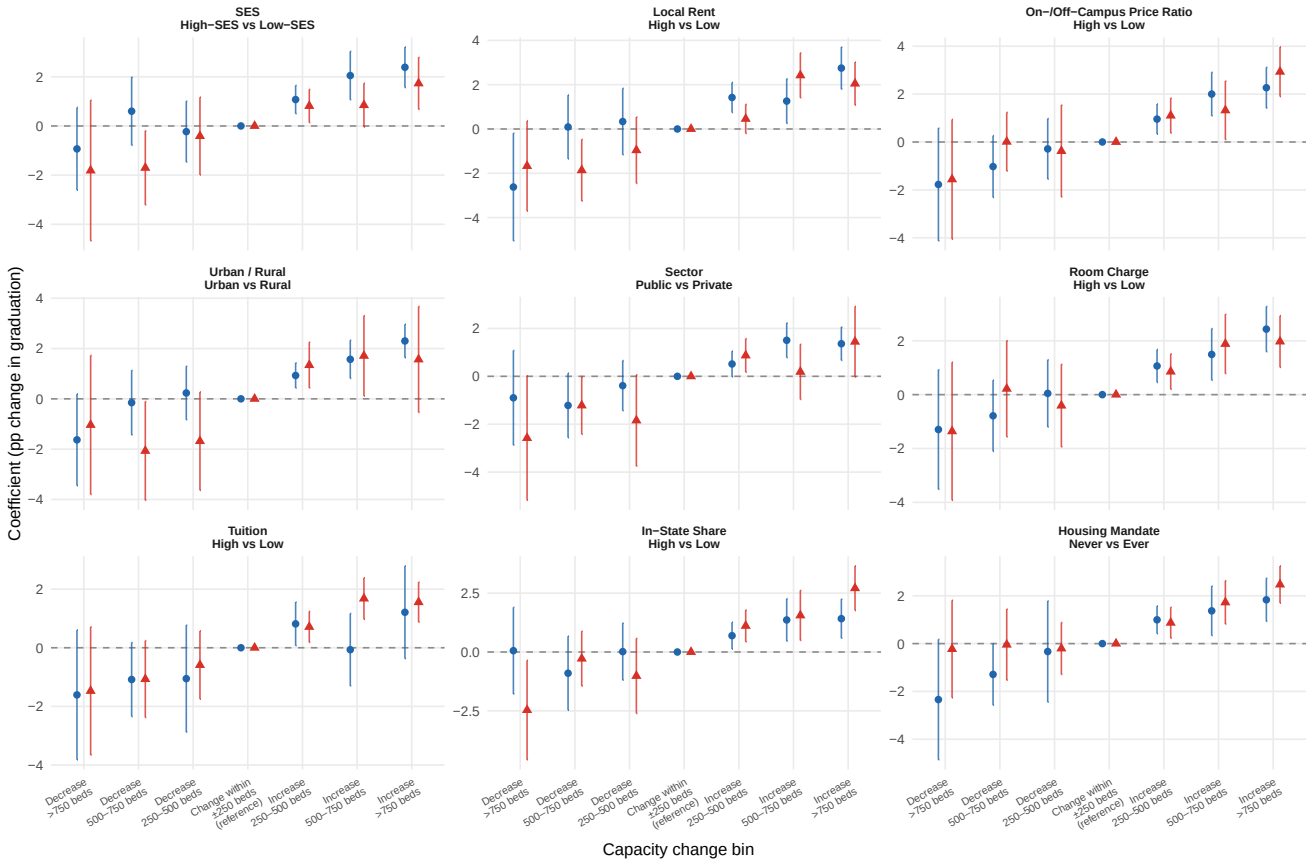


Figure 4: Binned differencing estimates for graduation by subgroup. Within each panel, the first-named subgroup is shown as blue circles, plotted on the left within each bin, and the second-named subgroup as red triangles, plotted on the right. The Housing Mandate panel compares institutions that never required on-campus residence (Never) with those that did at some point (Ever).

7.2. Spillover Effects

The cleanest test of the rent market channel exploits the geographic structure of local housing markets. If expansions operate partly by reducing local rents, then students at neighboring institutions should also benefit when a focal college builds: the new supply eases pressure in a shared off-campus rental market. Such spillover effects are impossible under a pure residence channel, in which the benefit flows exclusively to students occupying the new beds. The absence of positive spillovers would therefore be difficult to reconcile with a rent market channel. To examine this possibility, I construct a measure of each college's exposure to neighbor capacity changes and embed it in the first-difference framework used in the main analysis.

Using each campus's latitude and longitude as reported in the IPEDS institutional directory, I form a geodesic distance matrix across institutions and define, for each college i, the set of neighbors within $R = 20$ miles, $\mathcal{N}_i(R) = \{\, j \neq i : \text{dist}(i, j) \leq R \,\}$.[17]

[17]Neighbor sets are built once from a full distance matrix over campus coordinates (Haversine distance), avoiding per-row radius searches. Exposure sums are leave-one-out by construction (own capacity is excluded).

For each horizon $h \in \{1, 4\}$ (one year for retention; four years for graduation), let $\Delta_h B_{j,c}$ denote the h-year change in dormitory capacity (or, in the case of graduation, in average capacity over four years) at neighbor j. The "leave-one-out" neighbor exposure for college i in year c is:

$$E_{ic}^{(h)}(R) = \sum_{j \in \mathcal{N}_i(R)} \Delta_h B_{jc}, \qquad h \in \{1, 4\}, \; R = 20 \text{ miles}. \tag{5}$$

This is zero when no neighbor changes capacity and positive (negative) when neighbors expand (contract). To mirror the main design, I first translate $E_{ic}^{(h)}(R)$ into mutually exclusive neighbor bins.

Let $\Delta_h Y_{ic}$ be the h-year difference in the outcome (for retention, the one-year change in next-cohort retention; for graduation, the four-year change in the four-year graduation rate). Let $N_{ic}^{(h)}$ collect the neighbor bin indicators, $O_{ic}^{(h)}$ the own capacity bins, and X_{ic} and Z_{ic} the spending and selectivity/size controls, respectively, as in the main models. The binned specification is:

$$\Delta_h Y_{ic} = \alpha + \beta' N_{ic}^{(h)} + \pi' O_{ic}^{(h)} + \gamma' \, \Delta_h \mathbf{X}_{ic} + \theta' \, \Delta_h \mathbf{Z}_{ic} + \Delta_h \varepsilon_{ic}, \qquad h \in \{1, 4\}, \tag{6}$$

To improve precision and reduce estimation noise, I collapse the full bin structure into three groups: colleges experiencing a capacity increase of 500 beds or more, colleges experiencing a capacity decrease of 500 beds or more, and those with smaller changes (within ±500 beds), which serve as the reference category. Standard errors are clustered by five-digit ZIP code, since colleges in the same area are exposed to common local shocks. Of these colleges, 27.1% have no neighbors within 20 miles, 18.3% have exactly one, 11.8% have two, 7.8% have three, and 35.0% have four or more; institutions with no neighbors are then dropped. If the effects observed in the main results occur primarily through the rent market, we would then expect the estimated coefficients for the neighbor bins to be positive; if effects occur primarily through the campus residence channel, these coefficients will be close to zero.

The results in Table 5 provide little evidence that neighboring capacity expansions improve outcomes at nearby institutions. Neighboring expansions have small and statistically insignificant effects on both retention and graduation, while expansions in a college's own capacity remain positive and significant for graduation despite the smaller sample and the coarser treatment definition. The estimates are also precise enough to be informative: the 95% confidence interval for a neighboring expansion of 500 or more beds excludes positive effects on graduation larger than 0.07 percentage points, compared with an own-expansion effect of 1.72 percentage points in the same specification.[18] This pattern is difficult to explain if the main graduation gains

Spillover regressions include the same differenced controls as the main models. Repeating the exercise with alternative radii (10, 15 and 25 miles) yields the same qualitative conclusion: no robust spillover effects, while own expansions continue to raise graduation.

[18] Own and neighbor capacity changes are essentially uncorrelated, so the neighbor estimates are informative about spillovers and are not weakened by collinearity with a college's own expansions.

Table 5: Spillover effects (20 mi): Neighbor capacity changes on retention and graduation

	Retention	Graduation
Capacity Effect		
Neighbor decrease (≥500 beds)	0.196	0.555
	(0.157)	(0.738)
Neighbor increase (≥500 beds)	-0.112	-0.340
	(0.087)	(0.207)
Own decrease (≥500 beds)	-0.089	-1.440**
	(0.339)	(0.625)
Own increase (≥500 beds)	-0.059	1.720***
	(0.138)	(0.303)
Controls		
Spending controls	✓	✓
Selectivity index	✓	✓
Cohort size	✓	✓
Sample size		
	11,109	6,586

Notes:
Each column is a first-difference regression using 20-mile neighbor capacity changes.
Treatment bins: decrease of at least 500 beds and increase of at least 500 beds; omitted: changes of less than 500 beds in absolute value.
SEs clustered by campus ZIP code; *, **, *** indicate significance at 10%, 5%, and 1%.

operate primarily through broad local rent relief, since students at nearby colleges should also benefit from reduced pressure in the shared rental market. It is more consistent with a campus-residence channel, in which the benefits accrue to students who gain access to the additional institutionally controlled beds.

As an additional check, Appendix H returns to the main binned graduation specification and adds direct measures of broader local housing construction at four geographies, from the metropolitan area to a one-mile ring around the campus: the annual dollar value of building permits for residential buildings with five or more units (for the metropolitan area, the county and the permit office around the campus) and the change in the stock of such units within a mile of the campus. Adding these measures of private rental-supply growth leaves the estimated effects of own-campus capacity expansions essentially unchanged, and none of the private-supply measures is economically meaningfully related to graduation outcomes. Overall, these results provide little support for rent-market relief as the primary mechanism. They do not rule out smaller or more localized rental-market effects, but they suggest that the main effects are more likely to operate through access to campus residence.

7.3. Possible Components of the Campus Residence Channel

The evidence outlined above points away from the rent-market channel and towards campus residence itself. With IPEDS data I cannot establish this for certain, and I cannot observe which aspects of campus residence drive the effect. The existing literature can, however, provide suggestive evidence, by establishing which factors are associated with living on campus, and which of these factors have been linked to improved academic outcomes.

Students who live on campus differ from other students in ways that could matter for their outcomes. National surveys find that residents report somewhat more social engagement (Pike

and Kuh 2005) and more collaborative learning and discussion with diverse peers (Graham et al. 2018), and that they participate more in extracurricular activities (Schudde 2011). Data from individual colleges show that where students live shapes whom they befriend, with roommates and dormmates much more likely to interact with and become friends with one another (Marmaros and Sacerdote 2006; Mayer and Puller 2008). European survey evidence also shows that students living in halls commute far less than those living with their parents (Gwosc 2024).

These findings are associational and may reflect other differences between students who live on and off campus. Nevertheless, several studies link each of these factors to better outcomes. Randomly assigned roommates affect one another's grades (Sacerdote 2001), and the peer composition of randomly assigned groups affects dropout and achievement (Booij et al. 2017; Thiemann 2022). Observational evidence connects engagement with first-year grades and persistence (Kuh et al. 2008). Longer commutes reduce the number of days students spend on campus (Kobus et al. 2015), and a mandatory-attendance discontinuity shows that attendance raises exam performance (Dobkin et al. 2010). These factors are plausible components of the campus residence effect, but the IPEDS data cannot separate them.

8. Robustness: Differencing Models

8.1. Alternative Bins

The main results from the binned difference models compare periods in which there are small, moderate or large increases and decreases in housing capacity with periods in the same college in which little change occurs. The definitions of "small", "moderate", "large" and "little" are somewhat arbitrary, however. For this reason, in Appendix Section I I repeat this analysis using alternative bins. The results are generally consistent with the main models: capacity changes have negligible effects on first-year retention, and capacity expansions have a positive effect on graduation rates, with the magnitude of the effect increasing monotonically with the size of the expansion. The effects of decreases in capacity are not significant at conventional levels. The main results are therefore not an artifact of the chosen cut-offs.

8.2. Proportional Capacity Changes

Using absolute changes in beds is useful to capture discrete building events. However, absolute thresholds naturally make large institutions more likely to fall into "expansion" bins and leave smaller ones more often in the reference group. This concentrates treatment among larger colleges and complicates comparisons across very different baselines. For example, the effect of a 250-bed increase is likely to vary based on whether prior capacity was 250 or 2,500. As a robustness check, I repeat the binned analysis using proportional rather than absolute capacity changes, under two symmetric binning schemes (5-10% and 10-20% thresholds). Results are consistent with the main findings: retention effects are statistically indistinguishable from zero across both schemes, and capacity expansions deliver positive, significant gains in graduation rates. Full details of the specification and results are reported in Appendix J.

8.3. Start-to-Start Difference Models

Although the main specification aligns treatment with a cohort's full four-year exposure, it is informative to verify that results are not driven by the averaging window itself. A "start-to-start" design (comparing capacity at entry for adjacent cohorts) cleanly compares the housing environment students in different cohorts face when they arrive. This "start-to-start" definition implicitly assumes that the path of capacity in the intervening years (years 2–4 of enrollment) does not affect graduation. Aside from the use of first-year measures of capacity and spending controls, the model is equivalent to the main analysis. The results in Appendix K are consistent with the main findings: increases in entry-year capacity have a statistically significant effect on graduation rates. The slightly smaller magnitudes relative to the main results are expected, as the start-to-start design: (i) does not credit post-entry expansions that help the focal cohort, and (ii) does not net out pre-entry expansions that helped the previous cohort, attenuating the estimates.

8.4. Count Model

The main analysis models graduation rates. If the benefit accrues mainly to the students who occupy the new beds, as the evidence in Section 7 suggests, each additional bed can benefit at most one student, so it is natural to ask how many additional graduates result per additional bed. In this section, I replace the outcome with the four-year change in the number of four-year graduates. Cohort size and total enrollment enter as controls to ensure that any increase in graduate counts is not driven mechanically by enrollment growth; all other aspects of the specification are unchanged.

Results are shown in Appendix L. The coefficient is positive and significant across all specifications, consistent with the rate-based results. The fully specified model (column (4) of Table L.1) implies approximately 0.05 additional four-year graduates per additional average bed, or about 5 additional graduates per 100 added beds. This should be read as a reduced-form capacity effect: the additional beds need not all be occupied by students whose graduation outcomes are affected, and colleges may allocate new capacity across cohorts and student years in different ways. The count estimate therefore captures the number of additional on-time completions generated by expanding capacity, rather than an individual-level effect of receiving an on-campus bed.

9. Difference-in-Differences

9.1. Design

As a further check on the differencing results, I implement a cohort-level stacked difference-in-differences design that exploits large, discrete expansions in housing capacity. Rather than hand-curate events, I apply transparent selection rules to each college's capacity series. An expansion in year t is included as a treatment event if the following simultaneous criteria hold:

1. Large, discrete increase: a one-year increase of at least 300 beds, where the increase is at least 10% of capacity in $t-1$ and at least 2.5% of total enrollment in $t-1$.
2. Stable pre-period: in each of the three transitions $t-4, ..., t-1$, there are no annual proportional changes greater than 5%.
3. No immediate reversals: in $t, ..., t+3$ there are no contractions of at least 150 beds or 5%.
4. One event per institution: if multiple years in one college satisfy conditions 1-3, I retain the first qualifying event.

Robustness to different selection criteria is examined in Appendix Section M. The rules detect 116 events between 2006 and 2019. The estimation stacks contain 63 events for graduation and 70 for retention: events in the last three years of the panel have no control window through $t+3$, events in the first years have no $t-4$, and a college-event is dropped if the outcome or any of the six covariates is missing in either row. The graduation stack is also restricted to cohorts entering by 2015, so that no cohort's four-year window extends into 2020 or later. The 70 treatment events used in estimation are listed and plotted in Appendix A. The intent of this selection is to focus on plausibly exogenous large shocks to on-campus supply, i.e., sharp, isolated expansions that are less likely to be the mechanical by-product of gradual institutional change, and which are large enough to affect student outcomes through on-campus availability and potentially local rent markets. Reassuringly, whether an institution-year satisfies these qualifying-event criteria cannot be predicted from that institution's own lagged trends in enrollment, spending, selectivity, or tuition (Appendix D).

For each treated college event (i, t), I construct a two-period panel for each outcome. For graduation, the comparison is between students entering in the year of the expansion (exposed throughout their four years of study) and those entering four years earlier (unexposed throughout). For retention, the comparison is between students entering in the expansion year (exposed in the first year) and those entering one year earlier (unexposed in the first year). In both cases, later cohorts are excluded. The same stacked difference-in-differences framework is applied to both outcomes.

Controls for an event at t are colleges with flat capacity (annual changes of at most 200 beds and less than 5%) over the eight capacity years $t-4, \ldots, t+3$. Any college with a qualifying expansion in any year is excluded from the control group. The median event has 144 controls in the estimation samples for both graduation and retention.

I follow the stacking approach of Wing et al. (2024). Each event is a sub-experiment containing the treated college and its controls, with two rows per college. Corrective weights give treated rows weight one and a control in event a the weight $(N_a^D/N^D)/(N_a^C/N^C)$, where N_a^D and N_a^C are the numbers of treated and control colleges in event a, and N^D and N^C are their totals across events, so that the regression coefficient equals the average of the event-level two-by-two estimates. The estimating equation is

$$Y_{ice} = \delta(Treated \times Post)_{ice} + \mathbf{X}'_{ice}\gamma + \alpha_{ie} + \lambda_{ce} + \varepsilon_{ice}, \tag{7}$$

where $(Treated \times Post)_{ice} = 1$ for the treated cohort row of college i in event e and 0 otherwise, α_{ie} are college-by-event fixed effects, λ_{ce} are cohort(year)-by-event fixed effects, and $\mathbf{X}_{ice}$ contains the same covariates as in the differencing regressions (cohort size, the selectivity index and spending on scholarships, instruction, academic support and student services). Standard errors are clustered by college. The coefficient δ compares the average four-year graduation rate (or first-to-second-year retention rate) of the first cohort exposed to the capacity jump with that of an earlier cohort at the same college, net of shocks which affect the treated and control colleges equally. Because only two cohort rows are retained per college and event, and every comparison is made within an event, the design is not subject to the heterogeneous-timing bias that affects conventional two-way fixed-effects difference-in-differences with staggered adoption. Identification rests on parallel trends between treated and control colleges within each event window.

To assess the plausibility of parallel trends, I estimate event-study models on the same college-events as the main estimates, keeping only college-events observed at every event time in the window. The reference period is the cohort that serves as the pre-treatment comparison in the main estimate. For graduation this is the cohort entering at $t-4$, the last cohort to complete its four years before the expansion, so cohorts entering at $t-3, \ldots, t-1$ are partially exposed and cannot serve as pre-treatment leads; the leads are the cohorts entering at $t-7, \ldots, t-5$. For retention the reference is $t-1$, with leads at $t-4, \ldots, t-2$. Figure 5 presents the results. The pre-treatment coefficients are small and statistically insignificant for both outcomes, and the joint test that the leads are zero has p = 0.922 for retention (69 events) and p = 0.743 for graduation (30 events), supporting the parallel trends assumption.

I do not use event studies throughout, mainly because of sample size. Estimating the effect at t already requires a treated college's capacity to follow a specific pattern, and control colleges to have flat capacity, over eight years $(t-4, \ldots, t+3)$. Estimating effects at $t+1$ and beyond, or extending the leads, requires further restrictions on both groups, and the sample is depleted with each additional period (for example, the balanced pre-trend sample for graduation retains only 30 of the 63 events). I therefore use event studies only for the pre-trend check, and examine the effect on the partially exposed cohorts, those entering between $t-4$ and t, separately in Section 9.4. Appendix N repeats the event studies in Figure 5 on the unbalanced sample, using all available lead years.

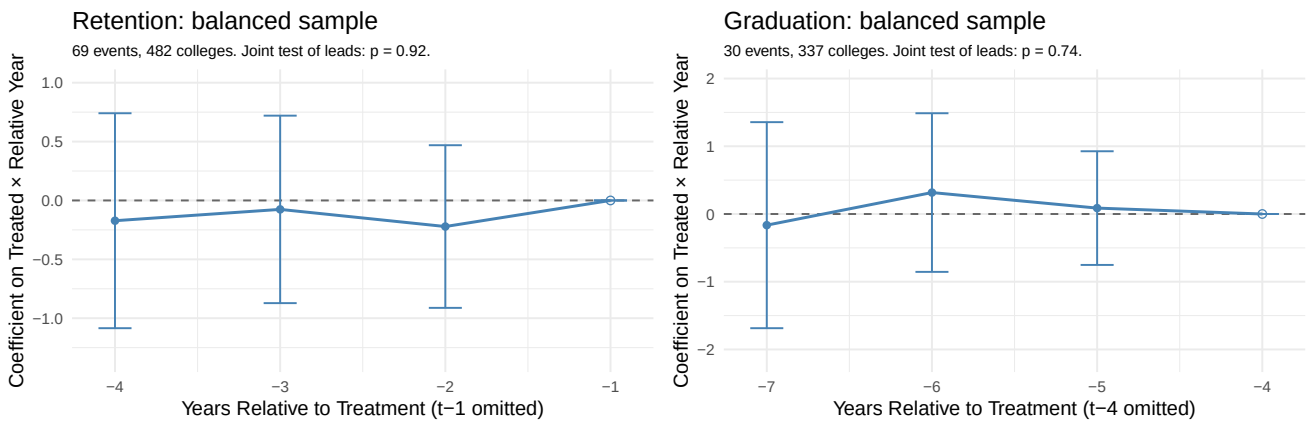


Figure 5: Pre-trend event studies, stacked difference-in-differences. Retention (left) and graduation (right), balanced sample. Bars are 95% confidence intervals.

Table 6: Stacked Difference-in-Differences Estimates

	Retention Models		Graduation Models	
	(1) Retention	(2) Retention + Controls	(3) Graduation	(4) Graduation + Controls
Post × Treated	0.450	0.610*	1.474**	1.584***
	(0.342)	(0.348)	(0.584)	(0.578)
Cohort Size		✓		✓
Selectivity Index		✓		✓
Scholarships		✓		✓
Instruction Spend		✓		✓
Academic Support		✓		✓
Student Services		✓		✓
Event × college and event × year FE	✓	✓	✓	✓
Observations (college-event-years)	21,956	21,956	19,232	19,232
Treatment Events	70	70	63	63
Control college-events	10,908	10,908	9,553	9,553
Unique Control Colleges	422	422	393	393

Notes: Stacked difference-in-differences estimates of the effect of large on-campus housing expansions on first-year retention and four-year graduation rates. Each treated college is compared with all colleges whose bed capacity is flat from four years before to three years after the event. College-by-event and year-by-event fixed effects; observations weighted so that the coefficient is the average of the event-level estimates (Wing, Freedman and Hollingsworth, 2024). Controls as indicated. Standard errors clustered by institution in parentheses.

9.2. Results

Table 6 reports difference-in-differences estimates for retention and graduation.[19] For first-year retention, point estimates are positive but imprecisely estimated: 0.45 without controls and 0.61 with controls, the latter significant at the 10% level only. As in the differencing models, this is consistent with a scenario in which first-year students are typically housed on campus, so upper-year students benefit at the margin in the event of an expansion.

Capacity expansions do have a clear positive effect on four-year graduation rates. After the

[19]The number of observations in the difference-in-differences results table is large because a control college appears in the stack of every event for which it is eligible, so its rows are repeated across many events.

inclusion of all controls, an expansion increases the graduation rate by 1.58 percentage points (1.47 without controls), based on 63 treated college-events and 393 unique control colleges. The point estimate increases slightly after adding covariates, indicating that observable cohort and spending characteristics were, if anything, less favorable in expansion years. The magnitude is similar to the estimated effects of increasing capacity by over 500 beds in the binned differencing models.

While these estimates stem from a relatively small subset of discrete treatment events, they offer a cleaner causal contrast by leveraging sharp capacity jumps and flat-capacity control windows. Despite its narrower scope, the binary design delivers a sharply identified estimate of the causal impact of a major housing expansion. Reassuringly, these results reinforce the main findings of the differencing approach: housing expansions lead to sizable improvements in graduation outcomes, with no detectable average effect on first-to-second-year persistence.

9.3. Treatment Definitions

To ensure that the results are not driven by the particular treatment definition, I re-estimate the design under five alternative sets of treatment thresholds, keeping all other elements fixed. Specifications 1 through 5 range from the most restrictive (15 events in the graduation stack) to the most permissive (149 events), with Specification 3 as the main analysis (63 events); full threshold details are reported alongside the results in Appendix M. The graduation estimates range from 1.02 to 1.97 percentage points and are significant in all but the smallest sample; the result is present whenever the sample is large enough to detect it. I also drop one or two of the three size criteria (beds, share of capacity, share of enrollment) at a time (Appendix M). As a further check on the stability of the main estimate, re-estimating the model after dropping each of the 63 graduation treatment events one at a time yields point estimates ranging from 1.42 to 1.78 percentage points, all of which are significant at the 5% level (Appendix O).

9.4. Exposure to the Expansion

The differencing models in Section 6 measure housing supply for a given cohort as average capacity over the four years of the degree, so a cohort fully exposed to an expansion throughout all four years benefits more than one exposed for only part of its degree. An advantage of the difference-in-differences design is that this exposure gradient can be examined directly. The main estimate compares the first cohort to spend all four years under the expanded capacity (entering at t) to a fully pre-expansion cohort ($t-4$). Cohorts already enrolled when the expansion occurred - those entering at $t-1$, $t-2$, and $t-3$ - received progressively less exposure: one, two, or three of their four degree-years elapsed before additional capacity was available. If the graduation benefit scales with duration of exposure, the effect should attenuate as the treated cohort is shifted backward toward the expansion year. To test this, I re-estimate the graduation difference-in-differences comparing each already-enrolled cohort ($t-1$, $t-2$, $t-3$) against the same fully pre-expansion cohort ($t-4$) used in the main analysis.

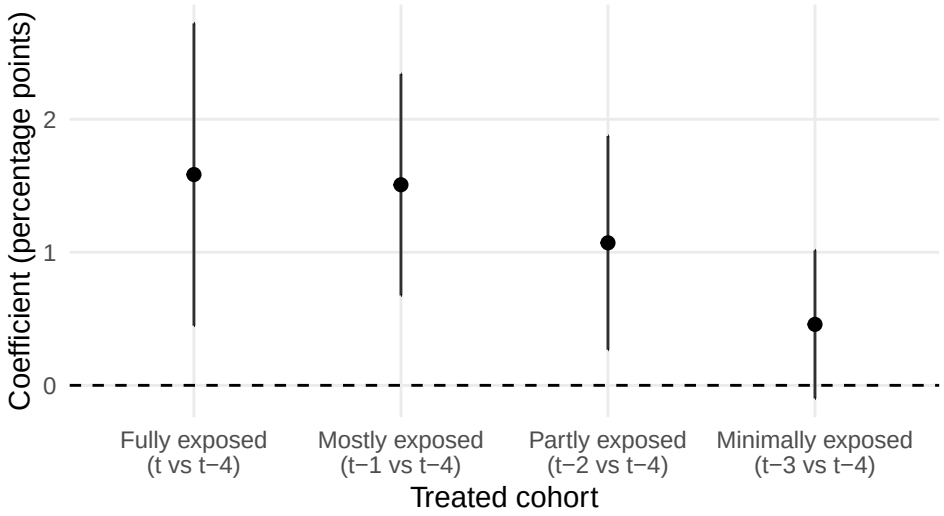


Figure 6: Graduation effect by treated cohort. The first point (t) is the main difference-in-differences estimate (first fully-exposed cohort vs $t-4$); subsequent points shift the treated cohort to students already enrolled one, two, and three years before the expansion. All specifications include controls. Points are estimated effects; bars are 95% confidence intervals.

Figure 6 presents the results. The estimated graduation effects are 1.58, 1.51, 1.07 and 0.46 percentage points for the cohorts entering at t, $t-1$, $t-2$ and $t-3$ respectively, attenuating monotonically as exposure falls; the last has a 95% interval of -0.10 to 1.01. This is consistent with an exposure mechanism: cohorts that spend a larger share of their degree under the expanded on-campus environment benefit more. The small gap between the effects for the t and $t-1$ cohorts (0.08 percentage points) is consistent with the minimal effect of an expansion on first-year students, which is the only difference in exposure between these two cohorts. The gradient is also reassuring for identification: students who had already enrolled when the expansion occurred could not have chosen the institution on the basis of the new capacity, and any concurrent institutional initiatives aimed at improving outcomes would most plausibly target incoming rather than continuing students. The results therefore hold for cohorts that are largely insulated from both student-selection and co-investment confounds.

10. Discussion

The national differencing design yields a coherent story about how on-campus housing supply influences student success, a pattern confirmed by the stacked difference-in-differences in Section 9. In the differencing models, even large absolute swings in reported beds barely move first-to-second-year retention. Similarly, in the difference-in-differences that isolates discrete, clearly timed expansions and compares the first exposed cohort with its immediate predecessor, the retention estimates are positive but imprecise. The natural explanation is that most expansions occur where first-years were already prioritized for housing, leaving little margin for improvements in retention. Although I do not observe within-college housing allocation directly, this interpretation is consistent with institutional housing policies: 86% of the colleges whose expansions enter the graduation difference-in-differences estimates (84% for the retention esti-

mates) state a first-year housing policy, whether a requirement to live on campus, a guarantee, priority, or designated first-year housing (see Appendix A). Furthermore, colleges typically have more than enough beds for their entire entering cohort: the average college in the sample has 2.34 beds per first-year student, and beds exceed first-year cohort size in 92% of college-year observations. It therefore seems likely that first-years are already largely accommodated.

Completion is more responsive than the first-year transition. Across difference models, four-year graduation rates improve meaningfully after sizable expansions: while the average effect per +100 beds is about 0.1 percentage points, large expansions of over 250 beds are associated with improvements of 1-2.3 percentage points depending on the magnitude of the expansion. The difference-in-differences estimates are consistent with these findings; large expansions increase graduation rates by approximately 1.58 percentage points. Because baseline four-year graduation averages about 40%, these effect sizes, though numerically modest, are educationally consequential. For the median treated college, which enrolls about 1,350 first-years, a 1.58 percentage point improvement in on-time completion implies approximately 21 additional degrees per cohort.

Effects are larger for cohorts with longer exposure to housing expansions. They reflect increases in total completion rather than a re-timing of degrees. Overall, the evidence is consistent with an on-campus residence channel rather than a rent market mechanism.

The findings have key implications for college administrators. Large, discrete housing expansions are a completion-enhancing investment. Even in the absence of retention gains, the accumulated improvements in graduation imply dozens of additional degrees per typical cohort at medium-to-large institutions. Where short-run persistence is a priority, institutions may wish to pair capacity policy with targeted supports (orientation, advising, financial aid, or academic support) rather than relying on beds alone, particularly when first-years are largely already housed.

The two channels emphasized in the framework imply that allocation interacts with institutional objectives. Where first-year housing is not universal, prioritizing entrants is the configuration most likely to influence the first-year transition; where the objective is on-time completion, allocating some capacity to later-year students is most likely to influence persistence beyond the first year. The effectiveness of expansions likely depends on how new capacity is allocated across student years, underscoring the importance of institutional housing policies. The following back-of-the-envelope calculations illustrate the cost per additional graduate from both the college's and the state's perspective.

From the college's perspective, residence halls are typically operated as auxiliary enterprises, meaning that operating costs and capital charges are largely covered by room revenue.[20] If

[20]In U.S. higher-education reporting, auxiliary enterprises (e.g., residence halls, dining) are defined as activities managed to be essentially self-supporting from user fees (see IPEDS glossary: https://surveys.nces.ed.gov/ipeds/public/glossary); several state systems specify auxiliaries operate on a self-supported basis (for example,

housing is fully self-financing, the marginal cost to the central academic budget is close to zero. Let $s \geq 0$ denote the net subsidy per bed-year from the central budget (after accounting for room revenue), B^* the number of new beds, Δg the graduation-rate effect in percentage points, and N the size of a typical entering cohort. The annual central-budget cost per additional graduate, once the expansion is in place, is $s \times B^*/[(\Delta g/100) \times N]$.

When $s \approx \$0$, the central-budget cost per additional graduate is correspondingly close to zero. Even when a subsidy is required, the implied cost per graduate remains modest relative to typical higher-education expenditures. Housing is a durable intervention: once built, additional capacity can benefit each entering cohort until the asset is retired. Unlike many student-success programs that require repeated annual funding, a single capital investment can affect completion for many future cohorts.

A separate calculation is relevant from the perspective of the state, which may finance construction through one-off capital grants. If the state provides a grant of size G and the expansion remains in service for L entering cohorts, the gross public capital cost per additional graduate over the asset's life is $G/[L \times (\Delta g/100) \times N]$, ignoring discounting. This is a gross state cost, not a net institutional cost: it does not subtract any future recovery of costs through room revenue. Because housing is a durable investment, the cost is spread over many cohorts, substantially lowering the cost per additional graduate.

To illustrate magnitudes, consider a stylized example of the median treated event in the difference-in-differences sample: a 450-bed expansion at a college enrolling $N = 1,350$ first-year students. The difference-in-differences estimate with controls is a 1.58 percentage point increase in graduation, which implies approximately 21 additional graduates per cohort. This is in line with the count model, which implies 0.05 additional graduates per additional bed, or about 23 for a 450-bed expansion. The difference-in-differences estimate is close to the binned differencing estimate of roughly 1.6 percentage points for expansions of 500-750 beds and above the estimate of roughly 1 percentage point for 250-500 beds, the bin containing the median event; the latter would imply about 14 additional graduates per cohort.

From the college's perspective, the annual central-budget cost per additional graduate depends on the net subsidy per bed-year, s. It is \$0 when $s = \$0$, rising to \$21,429 when $s = \$1,000$ and \$53,571 when $s = \$2,500$.

From the state's perspective, I take an illustrative capital cost of \$75,000 per bed,[21] so a

the University System of Georgia: https://www.usg.edu/business_procedures_manual/section15/manual/C1492).

[21] This is approximately the median cost per bed of the housing-only projects, among the expansions in this study, for which a project cost and bed count could be found. For example, Laurel Village at Colorado State University (615 beds, opened 2014) had a project cost of \$46.2 million (see https://building.colostate.edu/project/laurel-village/ and https://housing.colostate.edu/housing/halls/laurel-village/), and Evergreen Hall at Southern Illinois University Edwardsville (about 500 beds, opened 2007) cost \$39.3 million (see https://www.siue.edu/about/points-of-pride/completed-construction.shtml), or about \$75,000 and \$79,000 per bed respectively. These examples are broadly representative of such projects.

450-bed hall costs about $33,750,000. If the expansion remains in service for 30 cohorts, it generates approximately $30 \times 21 = 630$ additional graduates. A capital grant covering the full cost would therefore amount to roughly $33,750,000 / 630 \approx $53,571 per additional graduate over the asset's life, before discounting. These calculations are illustrative: they depend on a single point estimate, assume a stable treatment effect across cohorts, and ignore equilibrium responses. Nevertheless, these figures give a sense of the order of magnitude of the cost per additional graduate, which is lower still when room revenue covers operating expenses and capital costs are spread over many cohorts.

The above calculations capture only the fiscal costs of expansion. They do not account for fiscal returns to the state from additional degree completion. Quasi-experimental evidence suggests that degree completion results in substantial labor-market gains for marginal students. Zimmerman (2014) finds earnings gains of 22 percent between 8 and 14 years after high-school completion, while Denning et al. (2019) note that in their setting bachelor's degree receipt is associated with roughly $32,000 in additional annual earnings. Taking a deliberately conservative calibration of $30,000 in additional annual earnings and a 20 percent effective tax rate implies about $6,000 in additional annual tax revenue per additional graduate. Under the calibration above, one newly graduating exposed cohort would therefore generate about $21 \times \$6,000 = \$126,000$ in additional annual tax revenue while that cohort remains in the labor market; as successive exposed cohorts graduate, these annual tax revenues accumulate across cohorts. These fiscal returns are only one component of the broader social return, since most of the earnings gains accrue privately to graduates.

The policy implication is therefore twofold. First, for colleges that operate housing on an auxiliary basis, expansion may be close to self-financing if room revenue covers operating and capital costs. Second, even if a state subsidizes construction and does not capture rents directly, housing can still be a comparatively inexpensive way of generating additional graduates because one capital intervention can benefit many future cohorts. These figures are illustrative and institution-specific, but they suggest that housing expansions can be a cost-effective route to higher completion both as a low-net-cost institutional investment and as a durable public capital investment.

However, as the evidence points to the campus residence channel as the primary driver of improved outcomes, an important question for policy is what specific aspects of on-campus living generate these gains. While residence halls can generate positive returns for institutions, they require substantial upfront capital investment and may be difficult to finance. Expanding capacity therefore may not always be feasible. If the benefits derive less from the sheer number of beds and more from the environment they facilitate (peer networks, proximity to services, or on-campus support), then investments that replicate or enhance these features, even without large construction projects, could deliver comparable benefits at lower upfront cost.

Because dormitories are often profit-generating enterprises for colleges, two additional college-

side mechanisms could also explain the effects. First, rental-revenue incentives: after an expansion, institutions may have stronger incentives to retain students who pay for institutional housing, leading to institutional practices that raise persistence or completion. Second, performance justification: institutions may intensify completion efforts to demonstrate returns on capital spending (and maintain donor/political support), mechanically lifting graduation rates. Future research with data on precise grades, and not just coarse measures of retention and graduation, could seek to examine whether colleges appear to manipulate grades (for instance through changes in grade dispersion or pass-rate thresholds) following expansions.

11. Conclusion

This paper shows that expanding on-campus housing capacity improves student success. In a differencing design spanning 17 years and 1,208 colleges, large capacity increases have no average effect on first-year retention but raise four-year graduation by about 1-2.3 percentage points, with effects increasing in the size of the expansion. A rule-based difference-in-differences that isolates discrete expansions at large colleges finds an approximately 1.58-percentage-point gain in graduation and no precise evidence of retention effects.

Two additional results provide further insight into how housing expansions affect completion. Comparing cohorts with differing lengths of exposure to an expansion suggests that longer exposure is associated with larger graduation gains. Extending the outcome window to six- and eight-year graduation rates indicates that these effects reflect increases in overall completion rather than simply a re-timing of degrees.

Evidence on the underlying channel points away from a rent market mechanism. Subgroup analyses do not reveal systematic differences in effects across institutional characteristics, including local rent conditions or student socioeconomic composition. Consistent with this, there is no evidence of spillover effects on neighboring colleges. These findings suggest that the effects are unlikely to operate through broader changes in local housing markets and are more consistent with improvements in access to on-campus housing.

Large, discrete expansions are completion-enhancing investments; where first-years are already largely housed, institutions seeking immediate persistence gains should focus on targeted supports, while broader expansions can raise on-time graduation. Overall, relaxing on-campus housing constraints stabilizes enrollment through to degree. Future work linking capacity expansions to student-level survey data would help identify which aspects of campus living drive the observed effects and enable analysis of more detailed academic outcomes (such as GPA) and non-academic measures (including student well-being).

Acknowledgements

I thank Benjamin Elsner, David Madden, Aedín Doris, Kevin Denny, Sarah Parlane, Nora Strecker, Kla Kouadio, and particularly Paul Devereux for helpful feedback and comments. I

would also like to thank conference participants at the CESifo/ifo Junior Workshop on Economics of Education, the Leuven Economics of Education Research Conference, the Irish Economics Association Annual Conference, the EDGE Jamboree, and the Irish Postgraduate and Early Career Economist Conference, as well as seminar participants at University College Dublin and the Young Irish Economists Seminar Series.

Data Availability

The IPEDS, ACS, and BPS data used in this paper are publicly available from the respective sources cited in Section 4. The analysis code and constructed datasets are available from the corresponding author upon request.

Funding

This research did not receive any specific grant from funding agencies in the public, commercial, or not-for-profit sectors.

Declaration of Competing Interests

The author declares that they have no known competing financial interests or personal relationships that could have appeared to influence the work reported in this paper.

References

ACT Education Corp. 2018. *ACT / SAT Concordance*. https://www.act.org/content/act/en/products-and-services/the-act/scores/act-sat-concordance.html.

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): 136–63. https://doi.org/10.1257/app.1.1.136.

Araujo, Pedro de, and James Murray. 2010. *Estimating the Effects of Dormitory Living on Student Performance*. https://dx.doi.org/10.2139/ssrn.1555892.

Arcidiacono, Peter. 2004. "Ability Sorting and the Returns to College Major." *Journal of Econometrics* 121 (1): 343–75. https://doi.org/10.1016/j.jeconom.2003.10.010.

Arcidiacono, Peter, Patrick Bayer, and Aurel Hizmo. 2010. "Beyond Signaling and Human Capital: Education and the Revelation of Ability." *American Economic Journal: Applied Economics* 2 (4): 76–104. https://doi.org/10.1257/app.2.4.76.

Blimling, Gregory S. 1989. "A Meta-Analysis of the Influence of College Residence Halls on Academic Performance." *Journal of College Student Development*.

Booij, Adam S., Edwin Leuven, and Hessel Oosterbeek. 2017. "Ability Peer Effects in University: Evidence from a Randomized Experiment." *The Review of Economic Studies* 84: 547–78. https://doi.org/10.1093/restud/rdw045.

Bozick, Robert. 2007. "Making It Through the First Year of College: The Role of Students' Economic Resources, Employment, and Living Arrangements." *Sociology of Education* 80

(3): 261–85. https://doi.org/10.1177/003804070708000304.

Cameron, A. Colin, Jonah B. Gelbach, and Douglas L. Miller. 2008. "Bootstrap-Based Improvements for Inference with Clustered Errors." *The Review of Economics and Statistics* 90 (3): 414–27. https://doi.org/10.1162/rest.90.3.414.

Denning, Jeffrey T., Benjamin M. Marx, and Lesley J. Turner. 2019. "ProPelled: The Effects of Grants on Graduation, Earnings, and Welfare." *American Economic Journal: Applied Economics* 11 (3): 193–224. https://doi.org/10.1257/app.20180100.

Dobkin, Carlos, Ricard Gil, and Justin Marion. 2010. "Skipping Class in College and Exam Performance: Evidence from a Regression Discontinuity Classroom Experiment." *Economics of Education Review* 29 (4): 566–75. https://doi.org/10.1016/j.econedurev.2009.09.004.

Graham, Polly A., Sarah Socorro Hurtado, and Robert M. Gonyea. 2018. "The Benefits of Living on Campus: Do Residence Halls Provide Distinctive Environments of Engagement?" *Journal of Student Affairs Research and Practice* 55 (3): 255–69. https://doi.org/10.1080/19496591.2018.1474752.

Guzman, Matthew. 2021. "The Effects of on-Campus Housing Requirements on Student Outcomes." https://web.archive.org/web/20250804154546/https://www.matthew-guzman.com/wp-content/uploads/2021/12/20211207-Guzman-EAD942-Final-Paper.pdf.

Gwosc, Christoph. 2024. "Students' Housing Situation." In *Social and Economic Conditions of Student Life in Europe: EUROSTUDENT 8 Synopsis of Indicators 2021–2024*. Wbv Media. https://doi.org/10.3278/6001920ew009.

Hauschildt, Kristina, Christoph Gwosc, Hendrik Schirmer, Sylvia Mandl, and Cordelia Menz. 2024. *Social and Economic Conditions of Student Life in Europe: Eurostudent 8 Synopsis of Indicators 2021-2024*. Wbv Media GmbH & Company KG. https://doi.org/10.3278/6001920ew009.

Kobus, Martijn B. W., Jos N. Van Ommeren, and Piet Rietveld. 2015. "Student Commute Time, University Presence and Academic Achievement." *Regional Science and Urban Economics* 52: 129–40. https://doi.org/10.1016/j.regsciurbeco.2015.03.001.

Kuh, George D., Ty M. Cruce, Rick Shoup, Jillian Kinzie, and Robert M. Gonyea. 2008. "Unmasking the Effects of Student Engagement on First-Year College Grades and Persistence." *The Journal of Higher Education* 79 (5): 540–63. https://doi.org/10.1080/00221546.2008.11772116.

Lopez Turley, Ruth N., and Geoffrey Wodtke. 2010. "College Residence and Academic Performance: Who Benefits from Living on Campus?" *Urban Education* 45 (4): 506–32. https://doi.org/10.1177/0042085910372351.

Marmaros, D., and B. Sacerdote. 2006. "How Do Friendships Form?" *The Quarterly Journal of Economics* 121 (1): 79–119. https://doi.org/10.1093/qje/121.1.79.

Mayer, Adalbert, and Steven L. Puller. 2008. "The Old Boy (and Girl) Network: Social

Network Formation on University Campuses." *Journal of Public Economics* 92 (1-2): 329–47. https://doi.org/10.1016/j.jpubeco.2007.09.001.

National Center for Education Statistics. 2014. *National Postsecondary Student Aid Study: 2012 Undergraduate (NPSAS:12)*. Https://nces.ed.gov/surveys/npsas/.

National Center for Education Statistics. 2026. *Integrated Postsecondary Education Data System (IPEDS)*. U.S. Department of Education. https://nces.ed.gov/ipeds/.

Oreopoulos, Philip, and Kjell G. Salvanes. 2011. "Priceless: The Nonpecuniary Benefits of Schooling." *Journal of Economic Perspectives* 25 (1): 159–84. https://doi.org/10.1257/jep.25.1.159.

Perez-Felkner, Lara, Christine Baker-Smith, and Sara Goldrick-Rab. 2022. *Does Rent-Free Community Housing Make a Difference in Higher Education Outcomes?* Hope Center for College, Community, and Justice.

Pike, Gary R., and George D. Kuh. 2005. "First- and Second-Generation College Students: A Comparison of Their Engagement and Intellectual Development." *The Journal of Higher Education* 76 (3): 276–300. https://doi.org/10.1080/00221546.2005.11772283.

Reynolds, C. Lockwood. 2020. "The Effect of Dormitory Residence During College on Student Outcomes." *Journal of Human Capital* 14 (2): 249–89. https://doi.org/10.1086/709534.

Sacerdote, Bruce. 2001. "Peer Effects with Random Assignment: Results for Dartmouth Roommates." *The Quarterly Journal of Economics* 116 (2): 681–704. https://doi.org/10.1162/00335530151144131.

Schudde, Lauren T. 2011. "The Causal Effect of Campus Residency on College Student Retention." *The Review of Higher Education* 34 (4): 581–610. https://dx.doi.org/10.1353/rhe.2011.0023.

Stinebrickner, Ralph, and Todd R. Stinebrickner. 2006. "What Can Be Learned about Peer Effects Using College Roommates? Evidence from New Survey Data and Students from Disadvantaged Backgrounds." *Journal of Public Economics* 90 (8): 1435–54. https://doi.org/10.1016/j.jpubeco.2006.03.002.

Thiemann, Petra. 2022. "The Persistent Effects of Short-Term Peer Groups on Performance: Evidence from a Natural Experiment in Higher Education." *Management Science* 68 (2): 1131–48. https://doi.org/10.1287/mnsc.2021.3993.

Thompson, Jane, Virginia Samiratedu, and John Rafter. 1993. "The Effects of on-Campus Residence on First-Time College Students." *NASPA Journal* 31 (1): 41–47. https://doi.org/10.1080/00220973.1993.11072336.

U.S. Census Bureau. 2020. *American Community Survey 5-Year Estimates, 2015–2019*. U.S. Census Bureau. https://www.census.gov/programs-surveys/acs.

U.S. Census Bureau. 2021. *Building Permits Survey*. U.S. Census Bureau. https://www.census.gov/construction/bps/.

Webber, Douglas A., and Ronald G. Ehrenberg. 2010. "Do Expenditures Other Than In-

structional Expenditures Affect Graduation and Persistence Rates in American Higher Education?" *Economics of Education Review* 29 (6): 947–58. https://doi.org/10.1016/j.econedurev.2010.04.006.

Wing, Coady, Seth M. Freedman, and Alex Hollingsworth. 2024. *Stacked Difference-in-Differences*. Working Paper 32054. National Bureau of Economic Research. https://doi.org/10.3386/w32054.

Zimmerman, Seth D. 2014. "The Returns to College Admission for Academically Marginal Students." *Journal of Labor Economics* 32 (4): 711–54. https://doi.org/10.1086/676661.

Zong, Chen, and Alan Davis. 2024. "Modeling University Retention and Graduation Rates Using IPEDS." *Journal of College Student Retention: Research, Theory & Practice* 26 (2): 311–33. https://doi.org/10.1177/15210251221074379.

A. Individual Treatment Events

Table A.1 lists the individual treatment events used in the difference-in-differences analysis. "Capacity t" refers to capacity in the treatment year, while $t-4$ and $t-1$ represent the pre-treatment capacity for graduation and retention outcomes respectively. Capacity in $t+3$ is shown to verify that each expansion was sustained rather than temporary. The figures that follow the table plot the expansion paths in each treated college; the red dashed lines represent year t (the year of treatment) and year $t+3$ (the year in which the cohort enrolling in year t graduated), while the green dotted lines represent year $t-1$ (the pre-treatment year for retention models) and year $t-4$ (the pre-treatment year for graduation models).

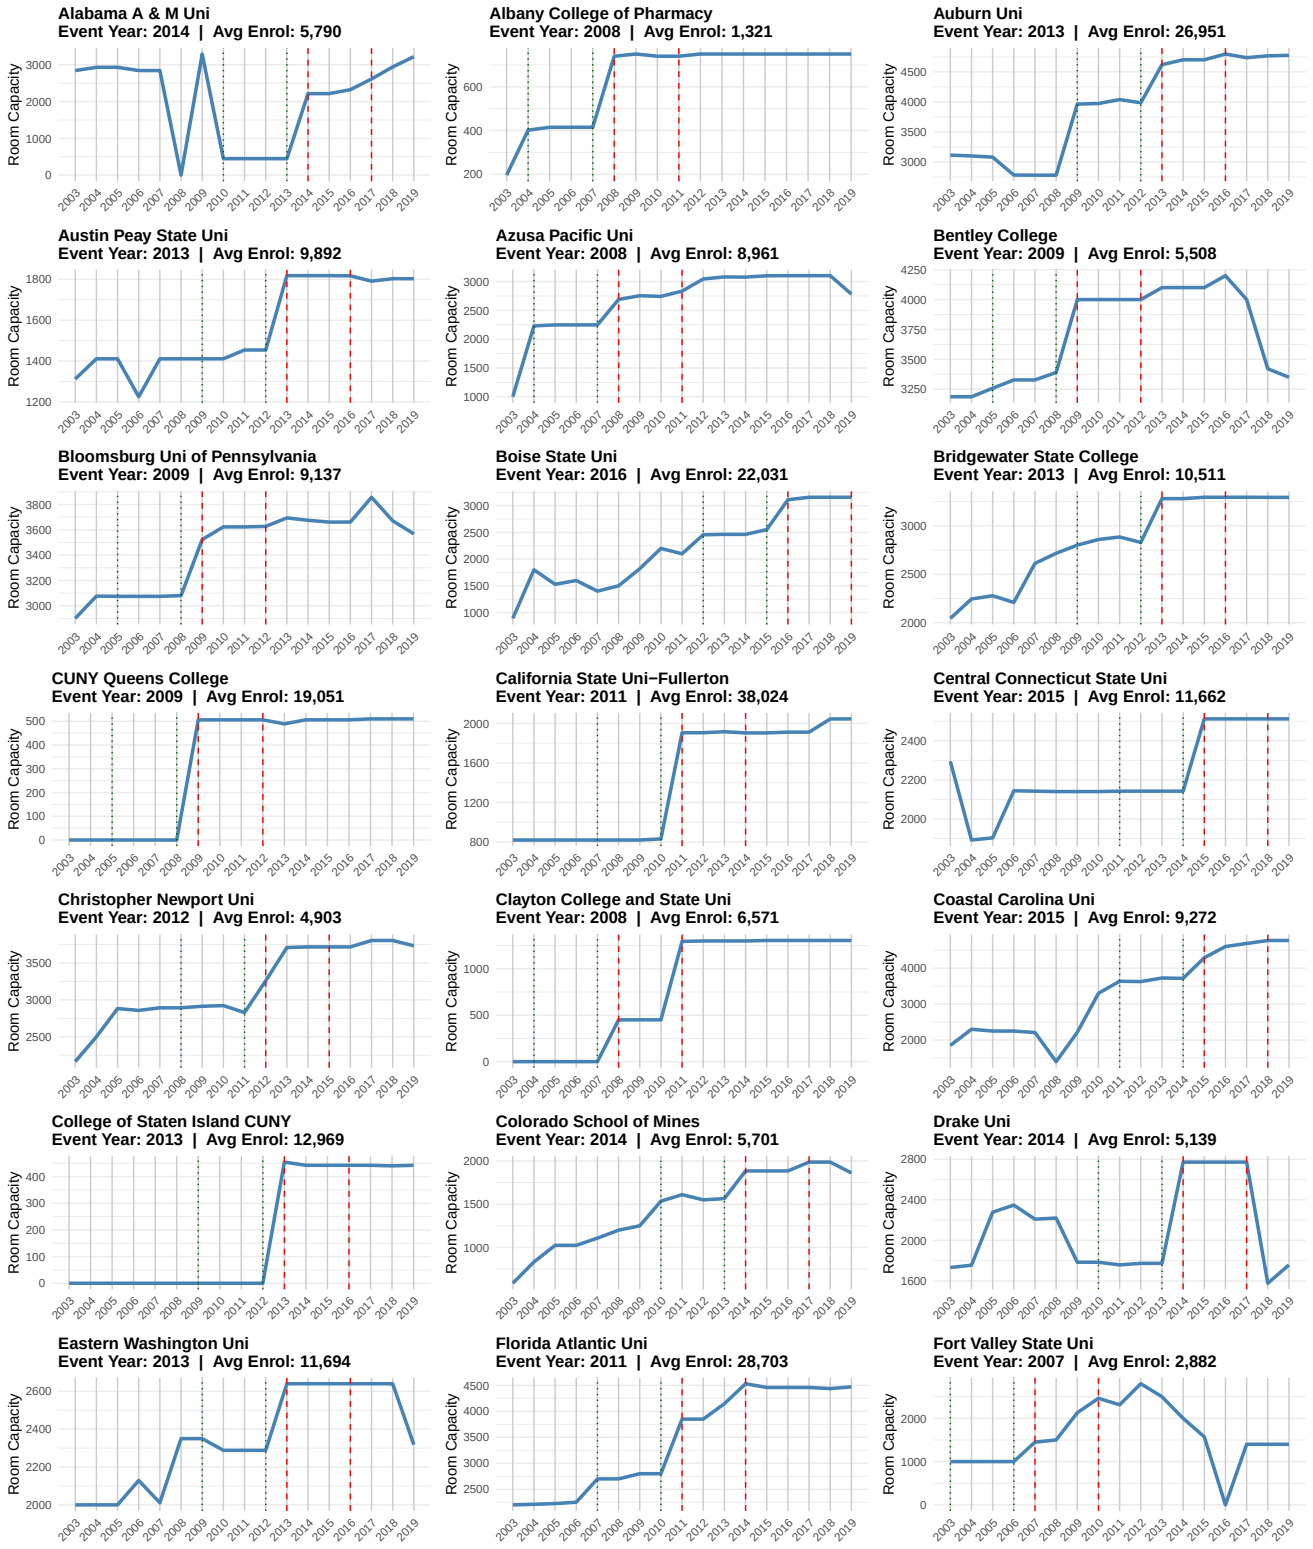
Table A.1: Capacity around treatment events (treated colleges in the estimation stacks).

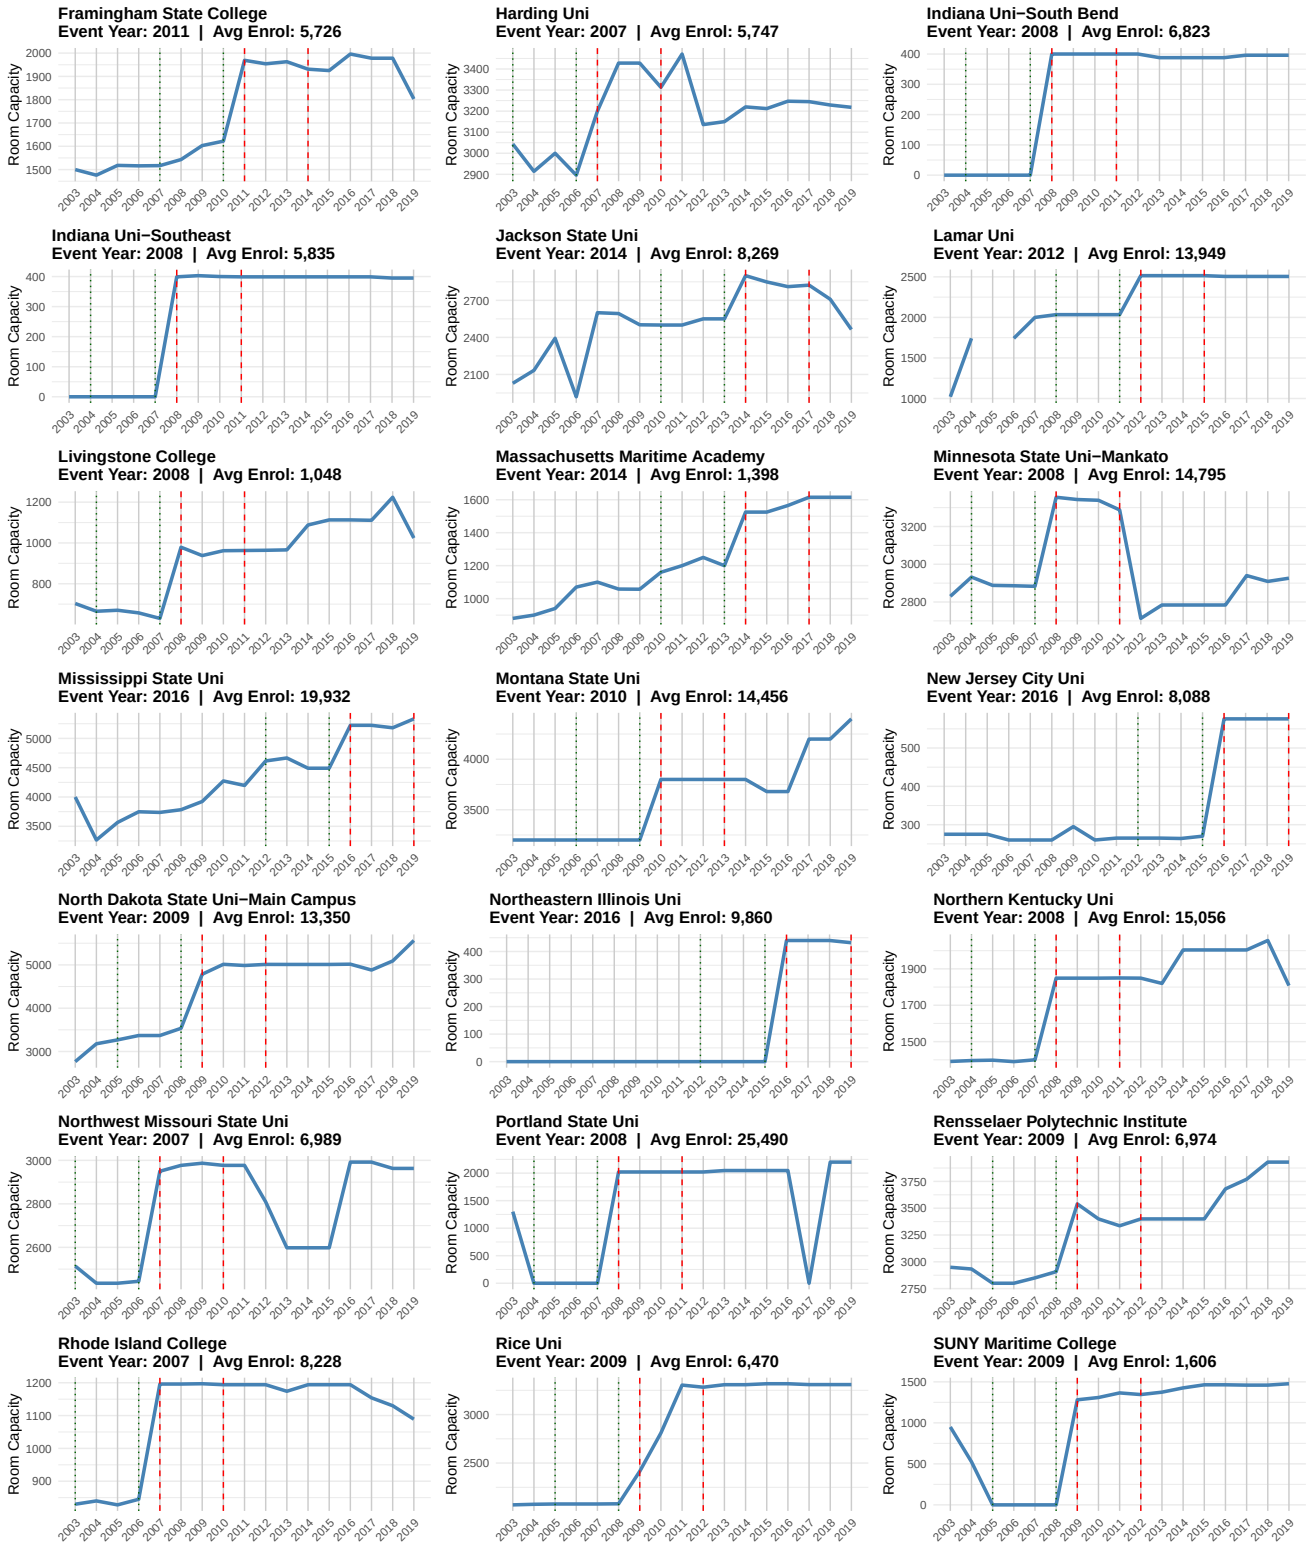
Institution	Event year	Capacity t-4	Capacity t-1	Capacity t	Capacity t+3	Graduation sample
Alabama A and M Uni†	2014	448	448	2216	2614	Yes
Albany College of Pharmacy†	2008	402	415	740	740	Yes
Auburn Uni	2013	3964	3986	4620	4798	Yes
Austin Peay State Uni†	2013	1411	1454	1818	1817	Yes
Azusa Pacific Uni†	2008	2231	2248	2690	2834	Yes
Bentley College†	2009	3258	3390	4000	4000	Yes
Bloomsburg Uni of Pennsylvania†	2009	3074	3080	3523	3627	Yes
Boise State Uni†	2016	2456	2553	3112	3158	No
Bridgewater State College†	2013	2801	2829	3281	3295	Yes
CUNY Queens College	2009	0	0	506	506	Yes
California State Uni-Fullerton†	2011	821	832	1906	1904	Yes
Central Connecticut State Uni	2015	2142	2142	2512	2512	Yes
Christopher Newport Uni†	2012	2893	2829	3259	3718	Yes
Clayton College and State Uni†	2008	0	0	450	1295	Yes
Coastal Carolina Uni†	2015	3635	3715	4288	4768	Yes
College of Staten Island CUNY	2013	0	0	454	442	Yes
Colorado School of Mines†	2014	1535	1565	1884	1986	Yes
Drake Uni†	2014	1784	1773	2772	2772	Yes
Eastern Washington Uni†	2013	2350	2288	2640	2640	Yes
Florida Atlantic Uni†	2011	2700	2800	3850	4531	Yes
Fort Valley State Uni†	2007	1000	1000	1450	2463	Yes
Framingham State College†	2011	1517	1622	1969	1931	Yes
Harding Uni†	2007	3044	2896	3200	3313	Yes
Indiana Uni-South Bend	2008	0	0	400	400	Yes
Indiana Uni-Southeast	2008	0	0	399	399	Yes
Jackson State Uni†	2014	2500	2550	2900	2823	Yes
Lamar Uni†	2012	2034	2034	2515	2515	Yes
Livingstone College†	2008	665	630	979	963	Yes
Massachusetts Maritime Academy†	2014	1160	1200	1525	1615	Yes
Minnesota State Uni-Mankato	2008	2932	2883	3355	3288	Yes
Mississippi State Uni†	2016	4615	4491	5223	5333	No
Montana State Uni†	2010	3200	3200	3800	3800	Yes

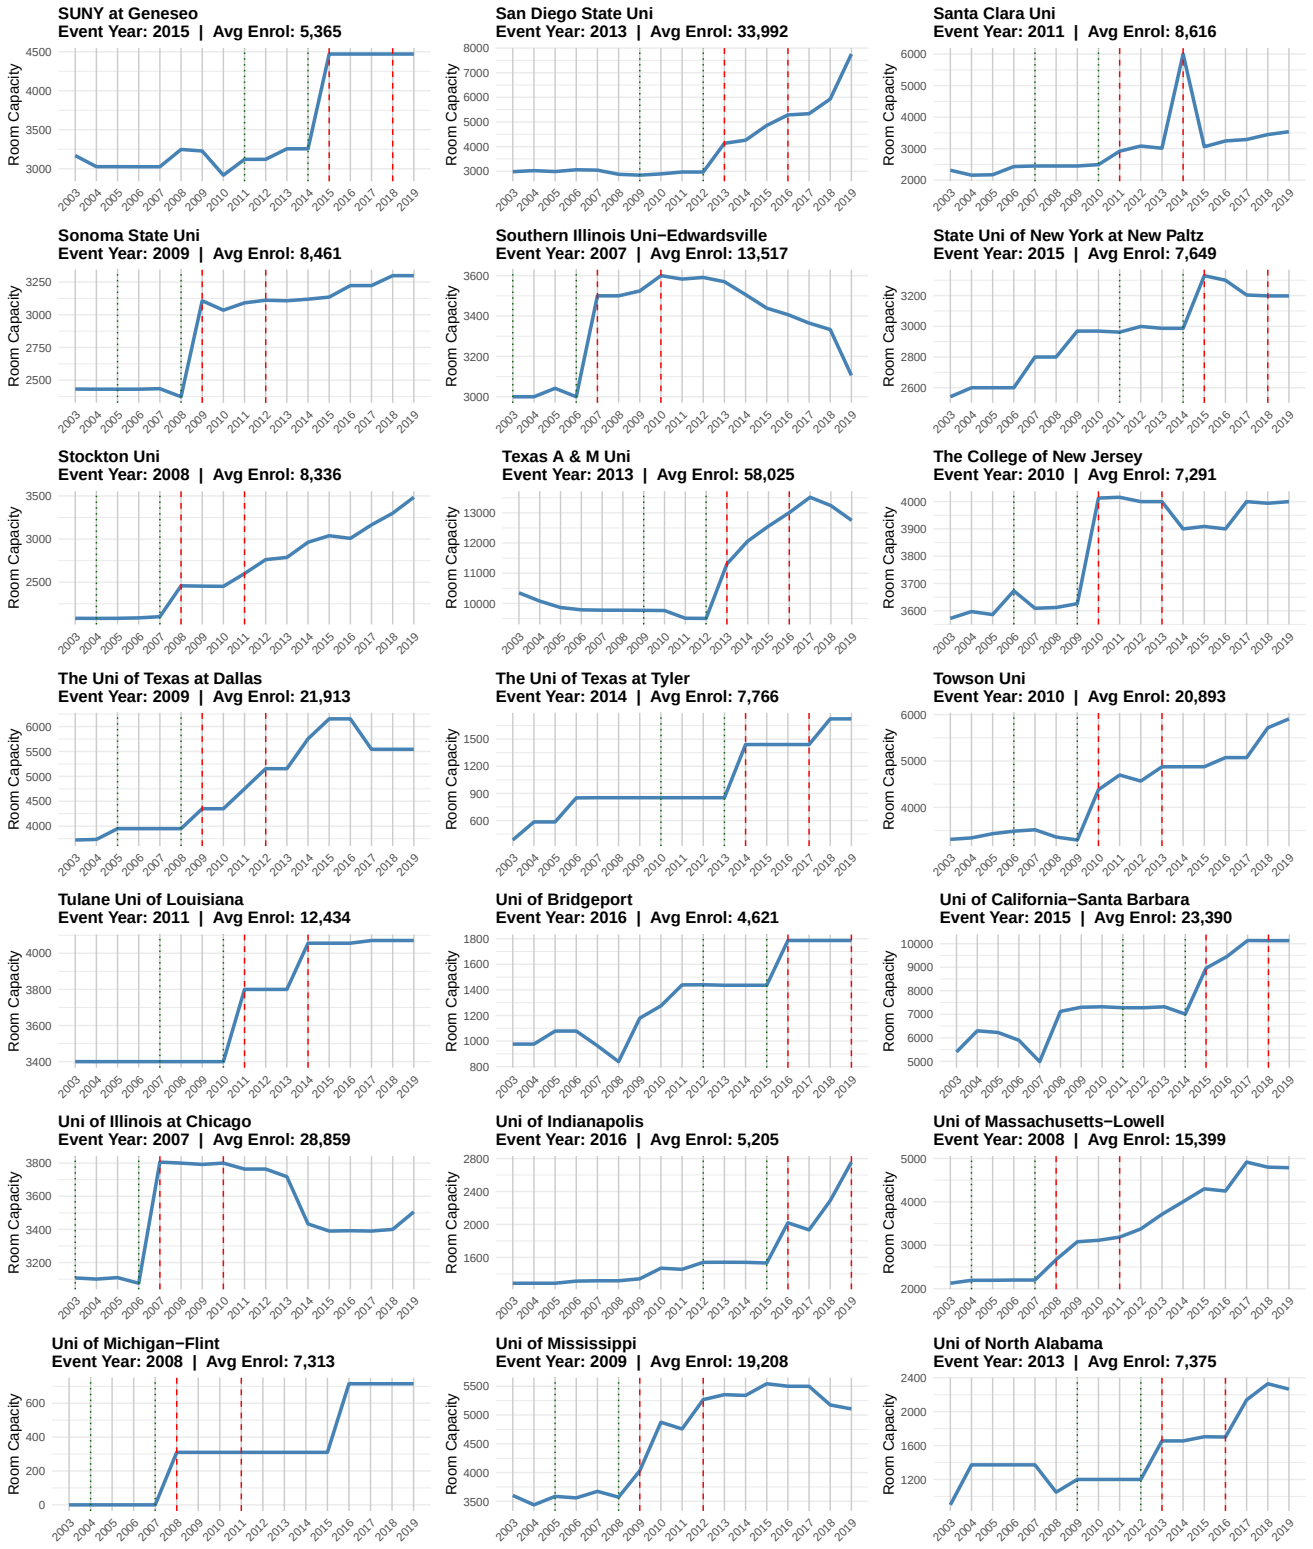
Table A.1: Capacity around treatment events (treated colleges in the estimation stacks). *(continued)*

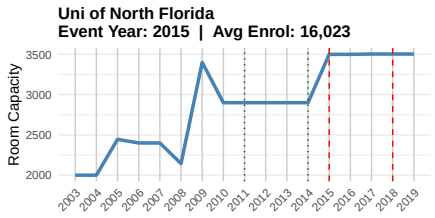
Institution	Event year	Capacity t-4	Capacity t-1	Capacity t	Capacity t+3	Graduation sample
New Jersey City Uni†	2016	265	270	576	576	No
North Dakota State Uni-Main Campus†	2009	3268	3535	4784	5011	Yes
Northeastern Illinois Uni	2016	0	0	440	432	No
Northern Kentucky Uni†	2008	1396	1400	1849	1850	No
Northwest Missouri State Uni†	2007	2514	2444	2950	2977	Yes
Portland State Uni†	2008	0	0	2021	2021	Yes
Rensselaer Polytechnic Institute†	2009	2800	2908	3540	3400	Yes
Rhode Island College†	2007	830	845	1196	1194	Yes
Rice Uni†	2009	2078	2080	2415	3282	Yes
SUNY Maritime College†	2009	0	0	1281	1345	Yes
SUNY at Geneseo†	2015	3120	3255	4471	4471	Yes
San Diego State Uni†	2013	2841	2969	4135	5285	Yes
Santa Clara Uni†	2011	2450	2494	2919	6002	Yes
Sonoma State Uni†	2009	2431	2373	3107	3111	Yes
Southern Illinois Uni-Edwardsville†	2007	3000	3000	3500	3600	Yes
State Uni of New York at New Paltz†	2015	2962	2987	3330	3198	Yes
Stockton Uni†	2008	2080	2100	2458	2600	Yes
Texas A and M Uni†	2013	9772	9505	11308	13008	Yes
The College of New Jersey†	2010	3673	3626	4013	4000	Yes
The Uni of Texas at Dallas†	2009	3950	3950	4350	5154	Yes
The Uni of Texas at Tyler†	2014	852	852	1440	1440	Yes
Towson Uni†	2010	3487	3296	4381	4875	Yes
Tulane Uni of Louisiana†	2011	3400	3400	3800	4056	Yes
Uni of Bridgeport†	2016	1440	1436	1786	1786	No
Uni of California-Santa Barbara†	2015	7279	7010	8959	10129	Yes
Uni of Illinois at Chicago	2007	3107	3075	3806	3800	Yes
Uni of Indianapolis	2016	1540	1532	2020	2757	No
Uni of Massachusetts-Lowell†	2008	2193	2199	2670	3188	Yes
Uni of Michigan-Flint	2008	0	0	310	310	Yes
Uni of Mississippi†	2009	3586	3573	4031	5268	Yes
Uni of North Alabama†	2013	1200	1200	1655	1701	Yes
Uni of North Florida†	2015	2900	2900	3500	3504	Yes
Uni of Portland†	2009	1421	1421	1790	1841	Yes
Uni of Utah†	2008	3500	3500	4300	4526	Yes
Uni of Wisconsin-Whitewater†	2010	3415	3370	3773	4435	Yes
West Chester Uni of Pennsylvania†	2009	3458	3538	4321	4850	Yes
Worcester State College†	2014	1167	1177	1577	1577	Yes
Xavier Uni†	2009	1815	1809	2123	2356	Yes

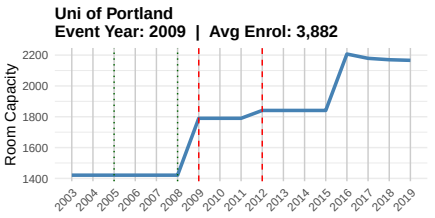
Notes: † The college has a first-year housing policy: a requirement to live on campus, a guarantee of a bed, priority in allocation, or designated first-year residence halls. Colleges without a dagger have no such policy or the policy could not be determined. Classifications are drawn from university housing pages, catalogues and news coverage; for 53 of the 70 colleges the source is dated 2003–2019 or shows that the policy was long-standing, and for the rest the classification reflects the current policy.

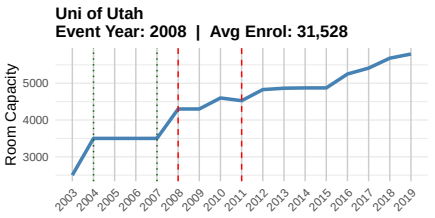

Alabama A & M Uni
Event Year: 2014 | Avg Enrol: 5,790
Albany College of Pharmacy
Event Year: 2008 | Avg Enrol: 1,321
Auburn Uni
Event Year: 2013 | Avg Enrol: 26,951
Austin Peay State Uni
Event Year: 2013 | Avg Enrol: 9,892
Azusa Pacific Uni
Event Year: 2008 | Avg Enrol: 8,961
Bentley College
Event Year: 2009 | Avg Enrol: 5,508
Bloomsburg Uni of Pennsylvania
Event Year: 2009 | Avg Enrol: 9,137
Boise State Uni
Event Year: 2016 | Avg Enrol: 22,031
Bridgewater State College
Event Year: 2013 | Avg Enrol: 10,511
CUNY Queens College
Event Year: 2009 | Avg Enrol: 19,051
California State Uni−Fullerton
Event Year: 2011 | Avg Enrol: 38,024
Central Connecticut State Uni
Event Year: 2015 | Avg Enrol: 11,662
Christopher Newport Uni
Event Year: 2012 | Avg Enrol: 4,903
Clayton College and State Uni
Event Year: 2008 | Avg Enrol: 6,571
Coastal Carolina Uni
Event Year: 2015 | Avg Enrol: 9,272
College of Staten Island CUNY
Event Year: 2013 | Avg Enrol: 12,969
Colorado School of Mines
Event Year: 2014 | Avg Enrol: 5,701
Drake Uni
Event Year: 2014 | Avg Enrol: 5,139
Eastern Washington Uni
Event Year: 2013 | Avg Enrol: 11,694
Florida Atlantic Uni
Event Year: 2011 | Avg Enrol: 28,703
Fort Valley State Uni
Event Year: 2007 | Avg Enrol: 2,882
Room Capacity

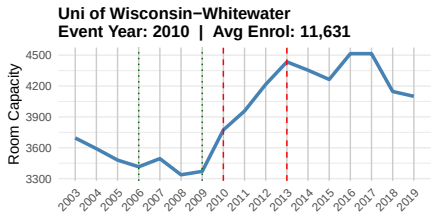

Framingham State College
Event Year: 2011 | Avg Enrol: 5,726
Harding Uni
Event Year: 2007 | Avg Enrol: 5,747
Indiana Uni−South Bend
Event Year: 2008 | Avg Enrol: 6,823
Indiana Uni−Southeast
Event Year: 2008 | Avg Enrol: 5,835
Jackson State Uni
Event Year: 2014 | Avg Enrol: 8,269
Lamar Uni
Event Year: 2012 | Avg Enrol: 13,949
Livingstone College
Event Year: 2008 | Avg Enrol: 1,048
Massachusetts Maritime Academy
Event Year: 2014 | Avg Enrol: 1,398
Minnesota State Uni−Mankato
Event Year: 2008 | Avg Enrol: 14,795
Mississippi State Uni
Event Year: 2016 | Avg Enrol: 19,932
Montana State Uni
Event Year: 2010 | Avg Enrol: 14,456
New Jersey City Uni
Event Year: 2016 | Avg Enrol: 8,088
North Dakota State Uni−Main Campus
Event Year: 2009 | Avg Enrol: 13,350
Northeastern Illinois Uni
Event Year: 2016 | Avg Enrol: 9,860
Northern Kentucky Uni
Event Year: 2008 | Avg Enrol: 15,056
Northwest Missouri State Uni
Event Year: 2007 | Avg Enrol: 6,989
Portland State Uni
Event Year: 2008 | Avg Enrol: 25,490
Rensselaer Polytechnic Institute
Event Year: 2009 | Avg Enrol: 6,974
Rhode Island College
Event Year: 2007 | Avg Enrol: 8,228
Rice Uni
Event Year: 2009 | Avg Enrol: 6,470
SUNY Maritime College
Event Year: 2009 | Avg Enrol: 1,606
Room Capacity

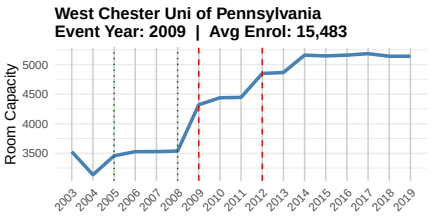

SUNY at Geneseo
Event Year: 2015 | Avg Enrol: 5,365
San Diego State Uni
Event Year: 2013 | Avg Enrol: 33,992
Santa Clara Uni
Event Year: 2011 | Avg Enrol: 8,616
Sonoma State Uni
Event Year: 2009 | Avg Enrol: 8,461
Southern Illinois Uni−Edwardsville
Event Year: 2007 | Avg Enrol: 13,517
State Uni of New York at New Paltz
Event Year: 2015 | Avg Enrol: 7,649
Stockton Uni
Event Year: 2008 | Avg Enrol: 8,336
Texas A & M Uni
Event Year: 2013 | Avg Enrol: 58,025
The College of New Jersey
Event Year: 2010 | Avg Enrol: 7,291
The Uni of Texas at Dallas
Event Year: 2009 | Avg Enrol: 21,913
The Uni of Texas at Tyler
Event Year: 2014 | Avg Enrol: 7,766
Towson Uni
Event Year: 2010 | Avg Enrol: 20,893
Tulane Uni of Louisiana
Event Year: 2011 | Avg Enrol: 12,434
Uni of Bridgeport
Event Year: 2016 | Avg Enrol: 4,621
Uni of California−Santa Barbara
Event Year: 2015 | Avg Enrol: 23,390
Uni of Illinois at Chicago
Event Year: 2007 | Avg Enrol: 28,859
Uni of Indianapolis
Event Year: 2016 | Avg Enrol: 5,205
Uni of Massachusetts−Lowell
Event Year: 2008 | Avg Enrol: 15,399
Uni of Michigan−Flint
Event Year: 2008 | Avg Enrol: 7,313
Uni of Mississippi
Event Year: 2009 | Avg Enrol: 19,208
Uni of North Alabama
Event Year: 2013 | Avg Enrol: 7,375
Room Capacity

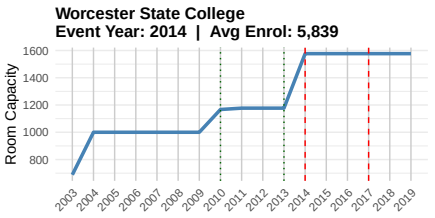

Uni of North Florida
Event Year: 2015 | Avg Enrol: 16,023
Room Capacity

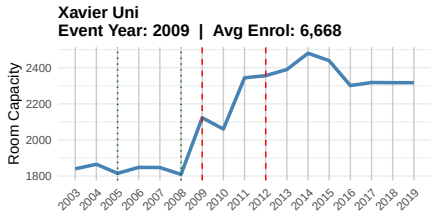

Uni of Portland
Event Year: 2009 | Avg Enrol: 3,882
Room Capacity


Uni of Utah
Event Year: 2008 | Avg Enrol: 31,528
Room Capacity


Uni of Wisconsin−Whitewater
Event Year: 2010 | Avg Enrol: 11,631
Room Capacity


West Chester Uni of Pennsylvania
Event Year: 2009 | Avg Enrol: 15,483
Room Capacity


Worcester State College
Event Year: 2014 | Avg Enrol: 5,839
Room Capacity


Xavier Uni
Event Year: 2009 | Avg Enrol: 6,668
Room Capacity

B. Local Rental Market Measurement

A natural question is whether on-campus housing is financially attractive relative to the local private rental market. If campus residence is cheaper than the outside option, expansions may improve outcomes partly by lowering students' housing costs; if it is more expensive, the benefits are more likely to reflect the non-price value of living on campus. Local rental market conditions are measured using the American Community Survey (ACS) 2019 five-year estimates (2015-2019) for the area around each campus. Three rent variables are used: median gross rent for all unit sizes, median gross rent for one-bedroom units, and median gross rent for two-bedroom units. All are median gross rent (the monthly cash rent plus any utilities paid by the tenant) for occupied units paying cash rent. The per-person cost of sharing a two-bedroom unit is the primary benchmark used throughout, as it most closely approximates the realistic off-campus alternative for most students.

Rents are measured at the smallest geography for which each is published. Median gross rent for all unit sizes is available for block groups (statistical areas of roughly 600 to 3,000 residents), but a single block group is imprecisely estimated (median relative standard error about 10%), and rents by number of bedrooms are published only for census tracts and larger areas. I therefore measure all-unit rent as the mean over block groups within one mile of the campus, weighting each block group by the share of its area inside the circle and by the inverse of the sampling variance of its estimate, which reduces the median relative standard error to about 2%. The two-bedroom rent per person scales each block group's all-unit rent by the ratio of two-bedroom to all-unit rent in its census tract, averages over the same one-mile circle, and divides by two. The one-bedroom rent is the median for the census tract containing the campus. Campus locations are the latitude and longitude reported in the IPEDS directory files. The two-bedroom measure is available for 99% of institutions with an on-campus room charge. Measuring rents at the larger PUMA level gives very similar results (correlation 0.87 across colleges).

Table B.1 compares two measures of student housing costs: (i) IPEDS on-campus room charges (converted to monthly equivalent by dividing by nine), and (ii) ACS median gross rent around the campus across different unit configurations. Over 2015–2019, the median on-campus monthly charge is approximately $674. Against the per-person cost of sharing a two-bedroom unit ($447/month), on-campus is more expensive at roughly 85% of institutions.

On this benchmark, campus residence does not appear to be a cheaper housing option for most students. The fact that students nevertheless demand on-campus housing suggests that it provides non-price benefits, such as proximity to campus, access to peers and services, reduced search costs, or a more structured residential environment.

Table B.1: On-campus room charges compared to ACS local rents, by institution type

Institution type	N	On-campus ($/mo)	ACS all-unit ($/mo)	ACS 1-bed ($/mo)	ACS 2-bed (per person, $/mo)	Ratio: on-campus / 2-bed per person	% on-campus < 2-bed per person
Private	600	684	903	740	463	1.39	16%
Public	431	653	824	675	429	1.40	15%
All	1031	674	863	709	447	1.40	15%

Notes: Entries are medians across institutions, 2015–2019. On-campus monthly charge = annual room charge divided by nine (nine-month academic year). ACS 2-bed pp = per-person cost of sharing a two-bedroom unit. Ratio = on-campus monthly / ACS 2-bed pp; values > 1 indicate on-campus is more expensive. ACS rents are 2019 five-year estimates for a one-mile ring around the campus (all-unit and two-bedroom per person) and the campus census tract (one-bedroom).

C. Compositional Stability around Expansion Events

The main graduation regression controls for cohort selectivity directly, and the already-enrolled analysis (Section 9.4) provides indirect evidence that graduation gains are not purely compositional. As a more direct check, I estimate event studies of the composition of the incoming cohort around the 70 expansion events used in the difference-in-differences analysis in Section 9 and listed in Appendix A. The design is the stacked difference-in-differences of the main estimates, with college-by-event and year-by-event fixed effects and no covariates.

Nine measures of the entering cohort are used: the share receiving Pell grants, the selectivity index, the shares who are female, Black, Hispanic, White and international, the share who graduated from high school in the past 12 months, and the share from the college's own state. Each outcome is divided by its standard deviation across institutions in the pre-expansion period (event times $t-6$ to $t-1$), so coefficients are in standard deviations of the outcome and comparable across panels. Race/ethnicity and gender of first-time students are available from 2008, and the residence measures (recent high-school graduates and in-state share) are published in full only in alternate years, so these panels rest on fewer events (the number is shown in each panel title).

Figure C.1 shows the six years before the expansion, the expansion year and the two years after. In the year of the expansion (t), the cohort directly relevant to the main analysis, none of the nine coefficients is significant at the 5% level. Of the 45 pre-expansion coefficients, one is significant at the 5% level: the selectivity index at $t-6$. Among the 18 coefficients for the two cohorts after the expansion, one is significant at the 5% level: the selectivity index at $t+2$. Selectivity is therefore the one measure with some weak indication of drift: it is somewhat lower six years before the expansion and somewhat higher two years after, relative to the year before. One significant coefficient among 18 is about what would be expected by chance. In any case, this does not bear on the main estimate, which compares the cohort entering at t with its comparison cohort and controls for selectivity directly. For the other eight measures the figure gives no indication of a change in the composition of the entering cohort, either before or at the expansion.

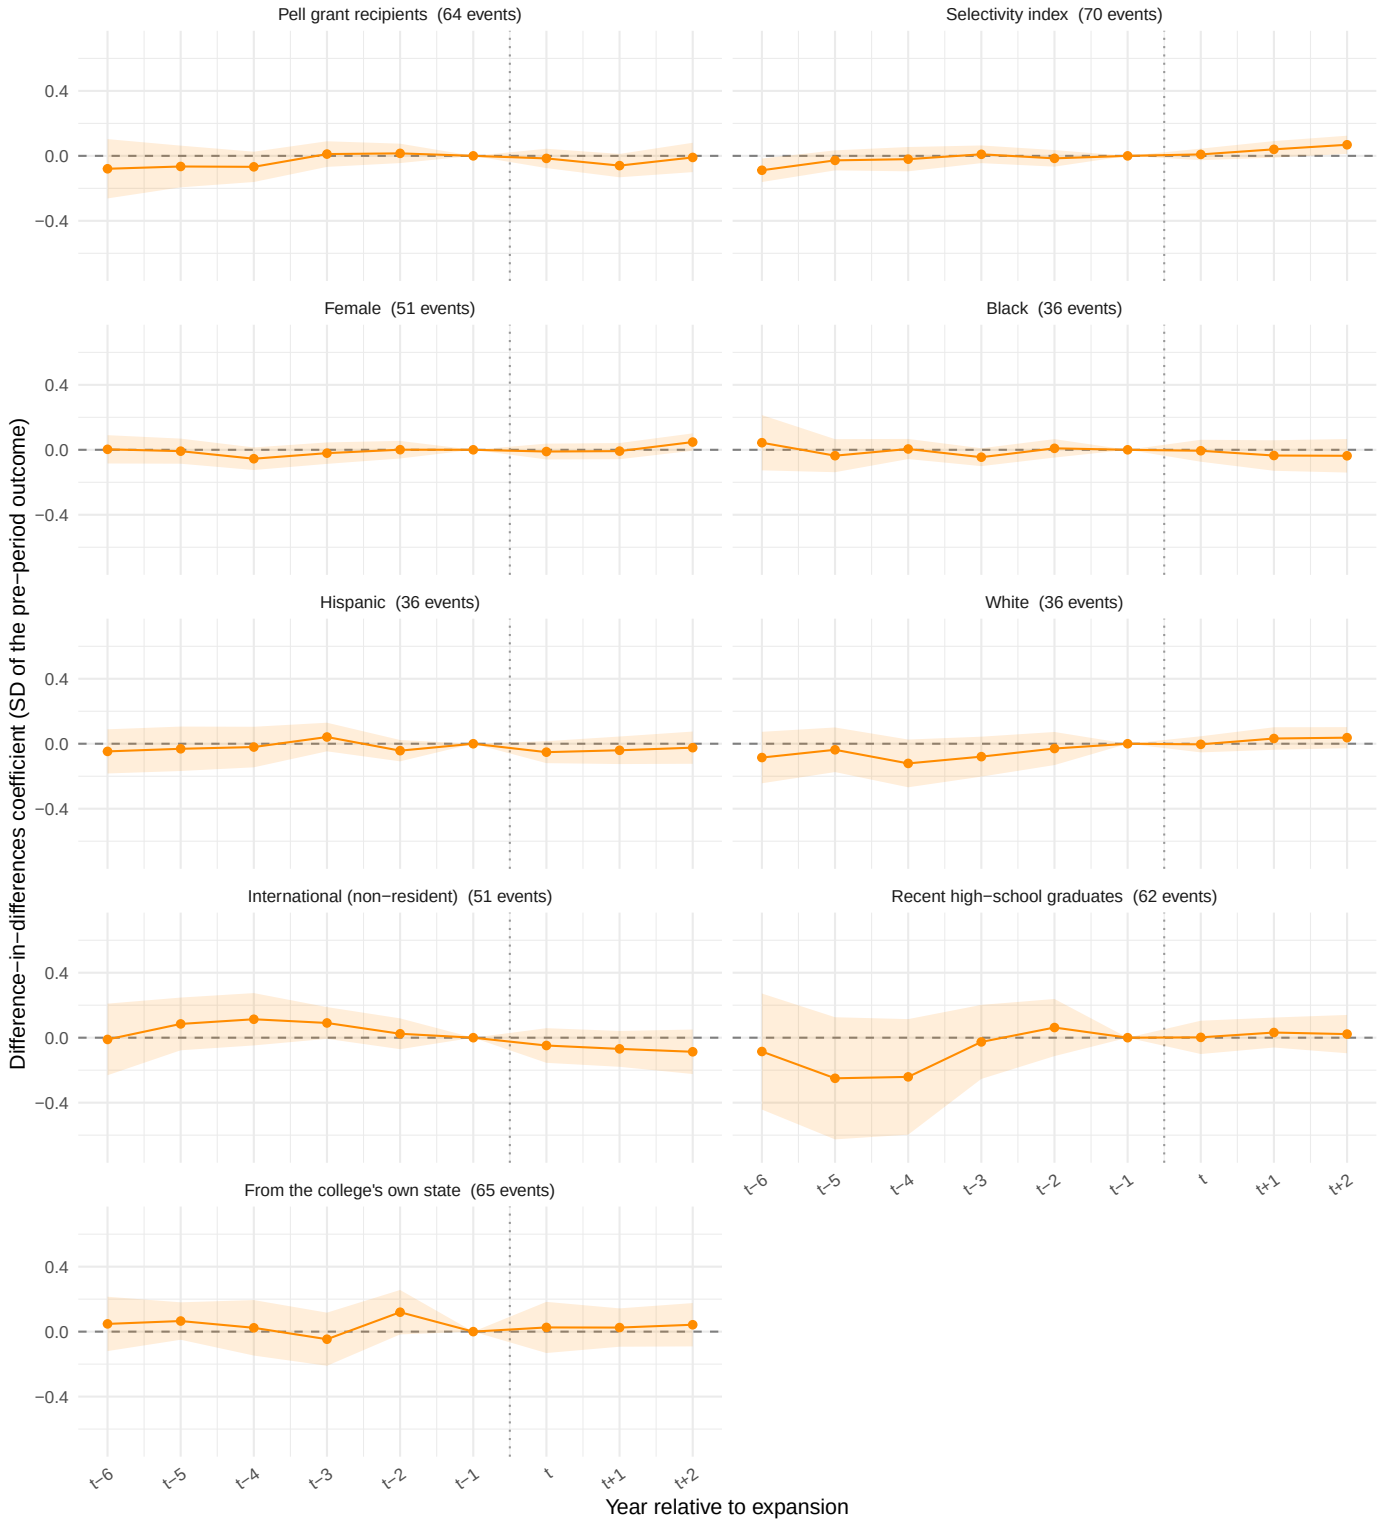


Figure C.1: Event studies of the composition of the incoming (first-time) cohort around housing expansion events: stacked difference-in-differences for the retention events, with college-by-event and year-by-event fixed effects. Each outcome is divided by its pooled pre-period standard deviation across institutions, so coefficients are in standard deviations; all panels share one y-axis. Year -1 is the omitted reference; shaded bands are 95% confidence intervals. The number of events behind each panel is given in its title.

D. Predictability of Expansion Timing

Identification in both the differencing and the difference-in-differences designs relies on the exact timing of new beds coming on stream being unpredictable. This appendix asks whether, with the data available, the timing of an expansion can be predicted from what the college could observe before the year in which it opens.

For each design I regress an indicator for the capacity change on lagged institutional characteristics and on lagged outcomes, with college and year fixed effects, clustering standard errors by college:

$$\text{Event}_{it} = \sum_{k} \beta_k \, X_{k,i,t-1} + \sum_{m} \delta_m \, Y_{m,i,<t} + \alpha_i + \gamma_t + \varepsilon_{it} \tag{D.1}$$

The predictors X are changes in enrollment, the selectivity index, each of the four spending categories and tuition. They are measured in three ways: as the one-year change from $t-2$ to $t-1$; as the four-year change from $t-5$ to $t-1$ (the "four-year trend"), which allows for slowly building developments that a one-year change would miss; and as the four annual changes from $t-5$ to $t-1$ entered separately, which lets each year of the build-up have its own coefficient (28 predictors). The outcomes Y are changes in the retention rate and in the four-year graduation rate, each measured over one and over four years. Only rates that could be observed at the time of the expansion are used: retention up to $t-1$, and the graduation rates of cohorts that entered in $t-5$ or earlier, which are realised by the end of $t-2$. For the difference-in-differences design, the event indicator marks the college-years of the treatment events used in the estimates. For the differencing design, it marks a capacity increase of 250 or more beds in the year, the same threshold as the smallest expansion bin in the main binned results.

Each cell reports a joint test that all coefficients in the specification, predictors and lagged outcomes together, are zero. Because the outcome is rare and each test involves many simultaneous restrictions, p-values come from a wild cluster bootstrap (Cameron et al. 2008) with 999 replications, clustered by college, rather than from the asymptotic chi-square approximation. A permutation check, which reshuffles each college's event indicator across its own years, gives the same conclusions.

No joint test rejects at conventional levels. The predictors also have very little explanatory power: they account for at most 1.8% of the variation in the event indicator within colleges and years, even when 28 annual changes are entered separately.[22] This pattern supports the identification strategy; at least with the data available, it is not possible to predict the exact year in which housing comes on stream from the college's own recent trajectory in selectivity, spending, tuition, enrollment, or retention and graduation, so the year of an expansion looks as good as random with respect to these observables.

[22] Adding Pell share and racial and age composition to the set of predictors leaves the conclusion unchanged but reduces the sample because these variables are less completely reported.

Table D.1: Can the timing of an expansion be predicted? Joint tests and explanatory power

Predictors	Regressors	Events	Obs.	Within R^2 (%)	p
Difference-in-differences events used in the estimates					
One-year changes	11	50	3,765	0.3	0.85
Four-year trends	11	50	3,724	0.4	0.87
Four annual changes	32	50	3,687	1.8	0.15
Expansion of 250 or more beds (differencing design)					
One-year changes	11	614	8,062	0.2	0.53
Four-year trends	11	606	7,920	0.2	0.45
Four annual changes	32	604	7,849	1.1	0.17

Linear probability models with college and year fixed effects; standard errors clustered by college. Regressors is the number of variables tested jointly: the lagged institutional characteristics plus four lagged retention and graduation changes. p is the wild cluster bootstrap p-value for the joint test that all are zero. Within R^2 is the share of the variation in the event indicator, within college and year, explained by the regressors. The sample is smaller than in the main analysis because the lagged outcomes reach back up to nine years.

E. Continuous Differencing Models - All Specifications

Table E.1: Effect of Housing Capacity on Retention - First Differences (Per +100 Beds)

	(1)	(2)	(3)	(4)	(5)
Capacity Effect					
Per +100 beds (1-year difference)	-0.007	-0.008	0.001	0.000	-0.002
	(0.006)	(0.006)	(0.005)	(0.006)	(0.006)
Controls					
Spending controls		✓		✓	✓
Selectivity index			✓	✓	✓
Cohort size			✓	✓	
Total enrollment					✓
Sample size					
	15,369	15,369	15,369	15,369	15,369

Notes:
Each column reports first-difference regressions using one-year changes in housing capacity. The coefficient is the percentage-point change in first-year retention per +100 beds. Robust standard errors clustered by institution; *, **, *** indicate significance at 10%, 5%, and 1%.

Table E.2: Effect of Housing Capacity on Graduation - Four-Year Differences (Per +100 Beds Average Exposure)

	(1)	(2)	(3)	(4)	(5)
Capacity Effect					
Per +100 average beds (four-year exposure)	0.107***	0.107***	0.116***	0.119***	0.101***
	(0.025)	(0.026)	(0.030)	(0.030)	(0.027)
Controls					
Spending controls		✓		✓	✓
Selectivity index			✓	✓	✓
Cohort size			✓	✓	
Total enrollment					✓
Sample size					
	9,099	9,099	9,099	9,099	9,099

Notes:
Each column reports first-difference regressions with four-year cumulative housing exposure. The coefficient is the percentage-point change in four-year graduation per +100 average beds. Robust standard errors clustered by institution; *, **, *** indicate significance at 10%, 5%, and 1%.

F. Binned Differencing Models - Full Results

Table F.1: Binned Differencing Regressions - Retention

	(1)	(2)	(3)	(4)	(5)
Capacity Effect					
Decrease >750 beds	-0.031	-0.087	-0.053	-0.135	-0.092
	(0.314)	(0.317)	(0.314)	(0.316)	(0.318)
Decrease 500–750 beds	0.173	0.137	0.120	0.067	0.048
	(0.531)	(0.531)	(0.539)	(0.540)	(0.542)
Decrease 250–500 beds	0.264	0.244	0.206	0.177	0.208
	(0.200)	(0.201)	(0.199)	(0.200)	(0.202)
Increase 250–500 beds	-0.178*	-0.201*	-0.041	-0.070	-0.115
	(0.105)	(0.106)	(0.110)	(0.110)	(0.115)
Increase 500–750 beds	-0.242	-0.262	-0.143	-0.174	-0.159
	(0.192)	(0.192)	(0.192)	(0.193)	(0.194)
Increase >750 beds	-0.134	-0.180	0.114	0.056	-0.032
	(0.134)	(0.136)	(0.140)	(0.141)	(0.154)
Controls					
Spending controls		✓		✓	✓
Selectivity index			✓	✓	✓
Cohort size			✓	✓	
Total enrollment					✓
Sample size					
	15,369	15,369	15,369	15,369	15,369

Notes:

First-difference regressions using 1-year changes in housing capacity. Omitted: changes of at most 250 beds in absolute value. Coefficients are percentage-point changes in first-year retention. Robust SEs clustered by institution; *, **, *** denote significance at 10%, 5%, and 1%.

Table F.2: Binned Differencing Regressions - Graduation

	(1)	(2)	(3)	(4)	(5)
Capacity Effect					
Decrease >750 beds	-1.203	-1.016	-1.633**	-1.372	-1.144
	(0.834)	(0.865)	(0.809)	(0.844)	(0.865)
Decrease 500–750 beds	-0.389	-0.281	-0.707	-0.542	-0.403
	(0.598)	(0.606)	(0.549)	(0.559)	(0.553)
Decrease 250–500 beds	-0.060	-0.032	-0.330	-0.272	-0.178
	(0.499)	(0.499)	(0.504)	(0.506)	(0.505)
Increase 250–500 beds	0.815***	0.869***	0.902***	0.954***	0.830***
	(0.230)	(0.230)	(0.225)	(0.223)	(0.223)
Increase 500–750 beds	1.338***	1.431***	1.486***	1.591***	1.412***
	(0.361)	(0.361)	(0.355)	(0.354)	(0.353)
Increase >750 beds	1.937***	2.038***	2.183***	2.307***	1.990***
	(0.303)	(0.334)	(0.306)	(0.322)	(0.321)
Controls					
Spending controls		✓		✓	✓
Selectivity index			✓	✓	✓
Cohort size			✓	✓	
Total enrollment					✓
Sample size					
	9,099	9,099	9,099	9,099	9,099

Notes:

First-difference regressions with four-year cumulative exposure. Omitted category: changes of at most 250 beds in absolute value. Coefficients are percentage-point changes in four-year graduation. Robust SEs clustered by institution; *, **, *** denote significance at 10%, 5%, and 1%.

G. Subgroup Analysis

This section describes the subgroup analysis used to explore potential mechanisms. The subgroup splits are intended to capture differences in students' exposure to local housing markets, the relative cost of living on and off campus, and institutional characteristics that may shape how housing expansions affect student outcomes.

If housing expansions operate primarily through a rent market channel, effects should be strongest in settings where students depend more heavily on private rental markets and where off-campus housing is relatively costly. This motivates subgroup splits based on local ACS rents around the campus, the ratio of on-campus room charges to those rents, and the socioeconomic composition of the student body. A similar logic applies to urban institutions, where students may be more embedded in local rental markets, and to institutions with different in-state enrollment shares, since these may differ in the extent to which students can commute from the parental home rather than rent locally.

The analysis also considers institutional price and type. Colleges are split by on-campus room charges and tuition levels to assess whether the effects of capacity expansions differ across higher- and lower-cost institutions. I also distinguish between public and private colleges, as well as between institutions that ever require first-year students to live on campus and those that never do so. The latter split addresses the concern that retention or completion effects may differ in settings where on-campus residence is partly determined by institutional policy.

The subgroup indicators are defined at the institution level using means over the sample period. Colleges are classified as low-SES if their average share of Pell Grant recipients is above the sample median, and high-SES otherwise. Colleges are classified as being in high-rent markets if their local median per-person rent (ACS 2019 five-year, per-person cost of sharing a two-bedroom unit in a one-mile ring around the campus; see Appendix B) exceeds the sample median. Colleges are similarly classified as having a high or low ratio of the monthly on-campus room charge (the annual charge divided by nine) to that local per-person rent, or high or low on-campus room charges (the typical annual room charge reported to IPEDS, excluding board), according to whether the relevant institution-level mean lies above or below the sample median. The same median split is applied to tuition and in-state enrollment shares. Colleges are classified as urban if they are located in a city or suburb, and rural otherwise.[23] Public-private status is based on the institution's modal control category over the sample period. Finally, the "Never" subgroup consists of institutions that never require first-year students to live on campus in any year of the sample period.

The large-small split is reported separately in this appendix, as large and small colleges

[23]Colleges are defined as "urban" if, based on their address, they are located in a city (large, midsize or small) or a suburb (large, midsize or small) under a methodology developed by the U.S. Census Bureau's Population Division in 2005. The remaining colleges are defined as rural, and are located in a town (fringe, distant or remote) or a rural area (fringe, distant or remote).

exhibit very different distributions of absolute capacity changes. This heterogeneity is estimated using proportional capacity-change bins rather than absolute bins, producing more balanced treatment groups. Results are reported in Table G.1.

Table G.1: Effect of Proportional Changes in Capacity on Graduation by Institution Size

	(1) Large (5–10%)	(2) Small (5–10%)
Capacity Effect		
Large increase in four-year cumulative exposure	0.618**	0.650*
	(0.243)	(0.355)
Moderate increase in four-year cumulative exposure	0.982***	-0.220
	(0.268)	(0.484)
Moderate decrease in four-year cumulative exposure	-0.168	0.487
	(0.388)	(0.901)
Large decrease in four-year cumulative exposure	-0.921*	-1.233*
	(0.472)	(0.654)
Controls		
Spending controls	✓	✓
Selectivity index	✓	✓
Cohort size	✓	✓
Total enrollment		
Sample size		
	4,768	4,233

Notes: Each column uses first-difference regressions with four-year cumulative capacity exposure. Capacity changes are binned; the omitted category is an absolute change of less than 5%. Coefficients represent percentage-point changes in graduation relative to this reference. Robust standard errors clustered by institution. *, **, *** denote 10%, 5%, and 1%.

H. Private Housing Supply

If the graduation effects of on-campus housing expansions operate mainly through reduced pressure in the off-campus rental market, then independent increases in local private rental supply should also be associated with higher graduation rates. I test this by adding measures of local private construction to the main binned graduation specification, using four approaches that measure supply at progressively finer geographies. The first three use the Census Bureau's Building Permits Survey (BPS), which reports the annual dollar value of permits for residential buildings with five or more units, at the metropolitan (Core-Based Statistical Area, CBSA), county and permit-office (city, town or township) levels. Five-or-more-unit permits capture the type of large apartment construction most relevant to student rental markets.[24] The fourth measure uses the housing stock itself: the change in the number of housing units in buildings with five or more units within one mile of the campus.

I match institutions to CBSAs using county Federal Information Processing Standards (FIPS) codes from the Census county-to-CBSA delineation file, obtaining matches for 1,141 of 1,208 institutions. I use the mean annual permit valuation in the four years before enrollment ($t-4$ to $t-1$), the baseline window of the capacity measure, as the supply measure.[25] Since permit valuation is already a flow variable (annual permits issued, not a stock of housing), no further differencing is applied; the variable is expressed in hundreds of millions of dollars. The sample is restricted to institution-years with CBSA-level BPS coverage and cohort years for which the full eight-year permit window is observed, leaving cohorts from 2007-2016.

The county and permit-office measures are constructed in the same way, with the same four-year lag, from the county and place files of the Building Permits Survey and the latitude and longitude of each campus; the permit office is the one whose jurisdiction contains the campus. The housing-stock measure is the change between the American Community Survey (ACS) five-year estimates ending in the year before enrollment and three years after it. This measure is available for the 2011-2015 cohorts.

Table H.1 adds each measure to the main binned graduation regression. The capacity bin estimates change by at most 0.03 percentage points in any column. The samples differ by measure, so this is the change relative to the same specification without the supply measure, estimated on that column's sample, and the coefficients in Table H.1 should not be compared across columns

[24] I use annual BPS permit valuations for 2003-2019: Table 3av, "Annual Permits by Metropolitan Area: Valuation," for the metropolitan areas, and the annual county and place files for the county and permit-office measures. Each measure is a flow of permitted construction, not a housing-stock measure, and includes all permitted construction of buildings with five or more units within the geography. It should therefore be interpreted as a proxy for broader local rental-supply growth, not as a clean measure of private student rental supply.

[25] The conclusion does not depend on this choice. I re-estimate the model using cumulated permit valuation in the four years before the cohort's window, in the cohort's own window, and over both windows, each shifted back by zero to three years to allow for construction lags; single-year levels at $t-4$ and $t-6$; and total residential permits instead of five-or-more-unit permits. The conclusions are consistent across all 32 specifications.

Table H.1: Private Housing Supply at Different Geographies and Graduation Rate

	Metro area (CBSA)	County	Place	Housing stock
Capacity Effect				
Decrease >750 beds	-1.169	-1.394*	-1.775**	-1.297
	(0.952)	(0.838)	(0.849)	(1.107)
Decrease 500–750 beds	0.045	-0.577	-0.589	-0.119
	(0.680)	(0.558)	(0.598)	(0.815)
Decrease 250–500 beds	-0.040	-0.290	-0.432	-0.227
	(0.625)	(0.506)	(0.567)	(0.714)
Increase 250–500 beds	1.024***	0.951***	0.939***	1.201***
	(0.241)	(0.222)	(0.245)	(0.281)
Increase 500–750 beds	1.553***	1.573***	1.514***	1.806***
	(0.407)	(0.352)	(0.395)	(0.453)
Increase >750 beds	2.321***	2.292***	2.356***	2.281***
	(0.370)	(0.325)	(0.336)	(0.416)
Private supply	-0.010	-0.067*	0.008	0.250
	(0.019)	(0.036)	(0.068)	(0.210)
Controls				
Spending controls	✓	✓	✓	✓
Selectivity index	✓	✓	✓	✓
Cohort size	✓	✓	✓	✓
Sample size				
	6,651	9,087	7,891	4,986

Notes:
All columns add one measure of private housing supply to the graduation binned differencing specification. Metro area, county and place: mean annual valuation ($100m) of permits for buildings with five or more units over years t-4 to t-1 (Building Permits Survey), for the metropolitan area, the county or the permit office whose jurisdiction contains the campus. Housing stock: change, in thousands, in housing units in structures with five or more units within one mile of the campus between the ACS 5-year estimates ending in c-1 and c+3.
The samples differ by measure. Robust SEs clustered by institution. *, **, *** denote 10%, 5%, 1%.

or with the full-sample estimates in the main binned results. None of the supply measures is positively and significantly related to graduation: the coefficient is small and insignificant for the metro area, the place and the housing stock, and is marginally negative at the county level ($p = 0.066$). Local private rental supply is not positively related to graduation at any of these geographies, so these results provide no support for the rent market channel.

I. Alternative Bins for Difference Models

Table I.1: Effects of Housing Capacity on Retention and Graduation (Large Bins)

	Ret (1)	Ret (2)	Grad (1)	Grad (2)
Capacity Effect				
Large increase in capacity (>800)	0.059	-0.042	2.258***	1.949***
	(0.146)	(0.157)	(0.325)	(0.325)
Medium increase in capacity (400–800)	-0.113	-0.098	1.305***	1.134***
	(0.129)	(0.138)	(0.272)	(0.272)
Medium decrease in capacity (-800 to -400)	0.586**	0.607**	-0.932*	-0.823*
	(0.286)	(0.289)	(0.501)	(0.498)
Large decrease in capacity (<-800)	-0.282	-0.226	-1.152	-0.931
	(0.313)	(0.312)	(0.893)	(0.916)
Controls				
Spending controls	✓	✓	✓	✓
Selectivity index	✓	✓	✓	✓
Cohort size	✓		✓	
Total enrollment		✓		✓
Sample size				
	15,369	15,369	9,099	9,099

Notes:
Retention uses 1-year first differences; graduation uses four-year cumulative housing exposure. Coefficients are percentage-point changes. Robust SEs clustered by institution;
*, **, *** denote significance at 10%, 5%, and 1%.

J. Proportional Capacity Changes

The proportional differencing operators are defined as

$$\Delta_1^p B_{ic} = \frac{\mathrm{B}_{ic} - \mathrm{B}_{i,c-1}}{\mathrm{B}_{i,c-1}} \tag{J.1}$$

in the retention models, and

$$\Delta_4^p \overline{B}_{ic} = \frac{\overline{B}_{ic} - \overline{B}_{i,c-4}}{\overline{B}_{i,c-4}} \tag{J.2}$$

in the graduation models.

I use two alternative symmetric binning schemes. For retention:

1. 5-10% thresholds: reference group if $|\Delta_1^p B_{ic}| < 0.05$, moderate increase (decrease) if $0.05 \leq \Delta_1^p B_{ic} \leq 0.10$ ($-0.10 \leq \Delta_1^p B_{ic} \leq -0.05$), large increase (decrease) if $\Delta_1^p B_{ic} > 0.10$ ($\Delta_1^p B_{ic} < -0.10$).
2. 10-20% thresholds: reference group if $|\Delta_1^p B_{ic}| < 0.10$, moderate increase (decrease) if $0.10 \leq \Delta_1^p B_{ic} \leq 0.20$ ($-0.20 \leq \Delta_1^p B_{ic} \leq -0.10$), large increase (decrease) if $\Delta_1^p B_{ic} > 0.20$ ($\Delta_1^p B_{ic} < -0.20$).

The same binning system is used for graduation models, but applied to $\Delta_4^p \overline{B}_{ic}$. The use of proportional changes does introduce the risk of results being skewed by spurious "large" percentage changes due to low capacity in $c-1$. Two safeguards are used to reduce this risk. 1) Observations where $B_{c-1} \leq 50$ (or, in the case of graduation models, $\overline{B}_{i,c-4} \leq 50$) are excluded from the analysis. 2) Observations with $|B_{ic} - B_{i,c-1}| < 100$ (or, in the case of graduation models, $|\overline{B}_{ic} - \overline{B}_{i,c-4}| < 100$) are placed in the reference group regardless of $\Delta_1^p B_{ic}$ (or $\Delta_4^p B_{ic}$).

The estimates for the effect on retention, as in the main analysis, are statistically indistinguishable from zero at the 5% level. For graduation, proportional increases in capacity deliver positive, statistically significant gains under the 5–10% scheme: point estimates are roughly 0.5–1.0 percentage points. Proportional decreases are statistically indistinguishable from zero. Under the 10–20% scheme (which has a larger reference group), only moderate increases remain significant at the 5% level, with an effect of roughly 0.4–0.5 percentage points. All other bins are imprecisely estimated (the effect of large expansions is positive, as expected and consistent with the results throughout other specifications, but imprecisely estimated). The proportional analysis supports the main finding that capacity expansions raise on-time completion, though magnitudes attenuate when the reference window widens. The smaller magnitudes are consistent with two explanations. First, proportional measures remain more sensitive to variation in denominators. Second, Safeguard 2 mechanically attenuates the estimates by placing some genuinely large proportional expansions (those with absolute changes < 100 beds) into the reference category.

Table J.1: Effect of Proportional Changes in Capacity on Retention

	(1) 5–10%	(2) 5–10%	(3) 10–20%	(4) 10–20%
Capacity Effect				
Large increase in capacity	-0.096	-0.124	-0.067	-0.100
	(0.103)	(0.104)	(0.164)	(0.166)
Moderate increase in capacity	0.093	0.074	-0.134	-0.156
	(0.096)	(0.100)	(0.126)	(0.128)
Moderate decrease in capacity	0.126	0.127	-0.086	-0.095
	(0.180)	(0.180)	(0.246)	(0.249)
Large decrease in capacity	0.198	0.201	0.696	0.723*
	(0.219)	(0.220)	(0.438)	(0.439)
Controls				
Spending controls	✓	✓	✓	✓
Selectivity index	✓	✓	✓	✓
Cohort size	✓		✓	
Total enrollment		✓		✓
Sample size				
	15,145	15,145	15,145	15,145

Notes: Each column reports coefficients from first-difference regressions. Reference category is small or no change in housing capacity (absolute change of less than 5 percent in columns 1 and 2, and absolute change of less than 10 percent in columns 3 and 4). Robust standard errors clustered by institution in parentheses. *, **, *** indicate significance at 10%, 5%, and 1%, respectively.

Table J.2: Effect of Proportional Changes in Capacity on Graduation

	(1) 5–10%	(2) 5–10%	(3) 10–20%	(4) 10–20%
Capacity Effect				
Large increase in four-year cumulative exposure	0.637***	0.521**	0.434*	0.311
	(0.203)	(0.202)	(0.260)	(0.261)
Moderate increase in four-year cumulative exposure	1.017***	0.951***	0.482**	0.397*
	(0.242)	(0.240)	(0.223)	(0.222)
Moderate decrease in four-year cumulative exposure	0.453	0.532	-0.564	-0.469
	(0.372)	(0.376)	(0.433)	(0.433)
Large decrease in four-year cumulative exposure	-0.612	-0.527	-1.167*	-1.101
	(0.382)	(0.382)	(0.690)	(0.693)
Controls				
Spending controls	✓	✓	✓	✓
Selectivity index	✓	✓	✓	✓
Cohort size	✓		✓	
Total enrollment		✓		✓
Sample size				
	9,001	9,001	9,001	9,001

Notes: Each column reports coefficients from first-difference regressions. Reference category is small or no change in housing capacity (absolute change of less than 5 percent in columns 1 and 2, and absolute change of less than 10 percent in columns 3 and 4). Robust standard errors clustered by institution in parentheses. *, **, *** indicate significance at 10%, 5%, and 1%, respectively.

K. Start-to-start Difference Models

Table K.1: Binned Differencing - Graduation (Start-to-Start)

	Cohort size	Total enrollment
Capacity Effect		
Decrease >750 beds	-0.128	-0.090
	(0.576)	(0.573)
Decrease 500–750 beds	-0.413	-0.326
	(0.448)	(0.446)
Decrease 250–500 beds	1.044***	1.079***
	(0.351)	(0.350)
Increase 250–500 beds	0.668***	0.594**
	(0.243)	(0.242)
Increase 500–750 beds	1.335***	1.185***
	(0.290)	(0.290)
Increase >750 beds	2.039***	1.797***
	(0.296)	(0.303)
Controls		
Spending controls	✓	✓
Selectivity index	✓	✓
Cohort size	✓	
Total enrollment		✓
Sample size		
	9,104	9,104

Notes:

First-difference regressions using four-year start-to-start changes in housing capacity.

Omitted: changes of at most 250 beds in absolute value. Coefficients are percentage-point changes in four-year graduation.

Robust SEs clustered by institution; *, **, *** indicate significance at 10%, 5%, and 1%.

L. Graduate Count Model

Table L.1: Effect of Housing Capacity on Graduate Count — Four-Year Differences (Per +1 Average Bed)

	(1)	(2)	(3)	(4)	(5)
Capacity Effect					
Per +1 average bed (four-year exposure)	0.161***	0.114***	0.077***	0.051***	0.093***
	(0.024)	(0.024)	(0.014)	(0.014)	(0.021)
Controls					
Spending controls		✓		✓	✓
Selectivity index			✓	✓	✓
Cohort size			✓	✓	
Total enrollment					✓
Sample size					
	9,099	9,099	9,099	9,099	9,099

Notes:

Outcome is the four-year change in the number of graduates. Coefficients should be interpreted as additional graduates per +1 average bed over the cohort window. Cohort size and total enrollment enter as controls (not denominators) so that gains are not driven mechanically by enrollment increases. Robust SEs clustered by institution; *, **, *** indicate 10%, 5%, and 1%.

M. Difference-in-Differences: Alternative Treatment Thresholds

Table M.1: Robustness of the graduation DiD to treatment detection thresholds

	Graduation				
	(1) Most restrictive	(2) Slightly restrictive	(3) Main	(4) Slightly permissive	(5) Most permissive
Post × Treated	1.020	1.970***	1.584***	1.701***	1.850***
	(0.891)	(0.649)	(0.578)	(0.503)	(0.458)
Controls:					
Spending controls	✓	✓	✓	✓	✓
Selectivity index	✓	✓	✓	✓	✓
Cohort size	✓	✓	✓	✓	✓
Event × college fixed effects	✓	✓	✓	✓	✓
Event × year fixed effects	✓	✓	✓	✓	✓
Observations	4,840	13,708	19,232	27,124	42,580
Treatment events	15	44	63	90	149
Specification:					
Change in beds (abs)	500	400	300	250	200
Change in beds (percent)	10%	10%	10%	8%	5%
Change in beds (student ratio)	5%	2.5%	2.5%	2%	none

Notes: Each column reports the coefficient from a stacked difference-in-differences regression with event fixed effects under an alternative treatment definition. Robust standard errors clustered by institution in parentheses. *, **, *** indicate significance at 10%, 5%, and 1%, respectively.

Table M.1 re-estimates the graduation effect under five size definitions, from 15 events in the most restrictive to 149 in the most permissive; the estimates range from 1.02 to 1.97. Table M.2 drops one or two of the three size criteria at a time.

With controls, the graduation estimate is between 1.58 and 1.72 percentage points in every column that applies the bed criterion (columns 1, 4, 5 and 7), where the median detected event adds 450–480 beds. It is smaller when the bed criterion is dropped: 1.01 and 1.02 with the percentage-of-capacity criterion alone or together with the enrollment criterion, and 0.58 (not statistically significant) with the enrollment criterion alone. These definitions detect more events, but the median event adds only 150–250 beds. This is consistent with the main finding that effects increase with the size of the expansion.

Table M.2: Stacked Difference-in-Differences Estimates by Size Criteria Used to Define Treatment Events

	One criterion			Two criteria			All three
	(1)	(2)	(3)	(4)	(5)	(6)	(7) Main
Graduation (with controls)							
Post × Treated	1.617***	1.009**	0.576	1.670***	1.722***	1.016**	1.584***
	(0.515)	(0.455)	(0.416)	(0.495)	(0.563)	(0.505)	(0.578)
Treatment events	109	179	247	77	75	148	63
Observations	32,314	48,966	48,970	23,092	22,718	41,290	19,232
Retention (with controls)							
Post × Treated	0.565**	0.177	0.153	0.686**	0.587*	0.375	0.610*
	(0.257)	(0.301)	(0.230)	(0.310)	(0.313)	(0.301)	(0.348)
Treatment events	120	197	267	84	84	165	70
Observations	36,748	54,886	54,552	25,860	26,274	47,156	21,956
Size criteria applied							
Beds added ≥ 300	✓			✓	✓		✓
Increase $\geq 10\%$ of capacity in $t-1$		✓		✓		✓	✓
Increase $\geq 2.5\%$ of enrollment in $t-1$			✓		✓	✓	✓

Notes: Each column defines a treatment event using only the size criteria marked with a tick (thresholds as in the main specification). Standard errors clustered by institution in parentheses. *, **, *** indicate significance at 10%, 5%, and 1%, respectively.

N. Pre-Trends: All Available Event Windows

The pre-trend figure in the main text keeps only college-events observed at every event time in the window (the balanced sample). This is preferable as the coefficients are directly comparable and any slope in them is a trend in the same colleges. However, the tradeoff is sample size. Figure N.1 repeats the exercise using every lead year that exists for each college-event, so that each coefficient is estimated from all events with data at that event time; the label on each point gives the number of events informing it. The conclusions are unchanged.

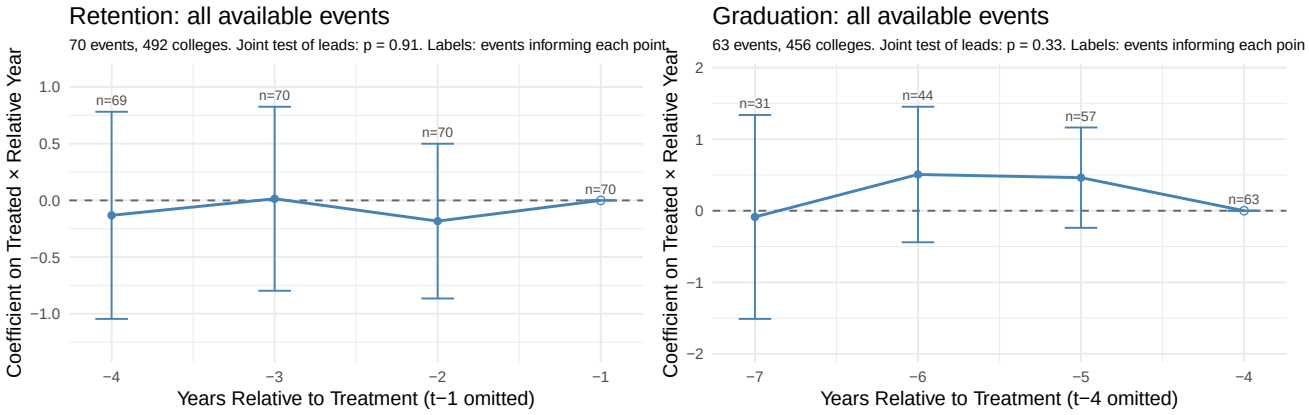


Figure N.1: Pre-trend event studies, stacked difference-in-differences, using all available lead years. Retention (left) and graduation (right). Bars are 95% confidence intervals; labels give the number of events informing each coefficient.

O. Leave-One-Out Difference-in-Differences

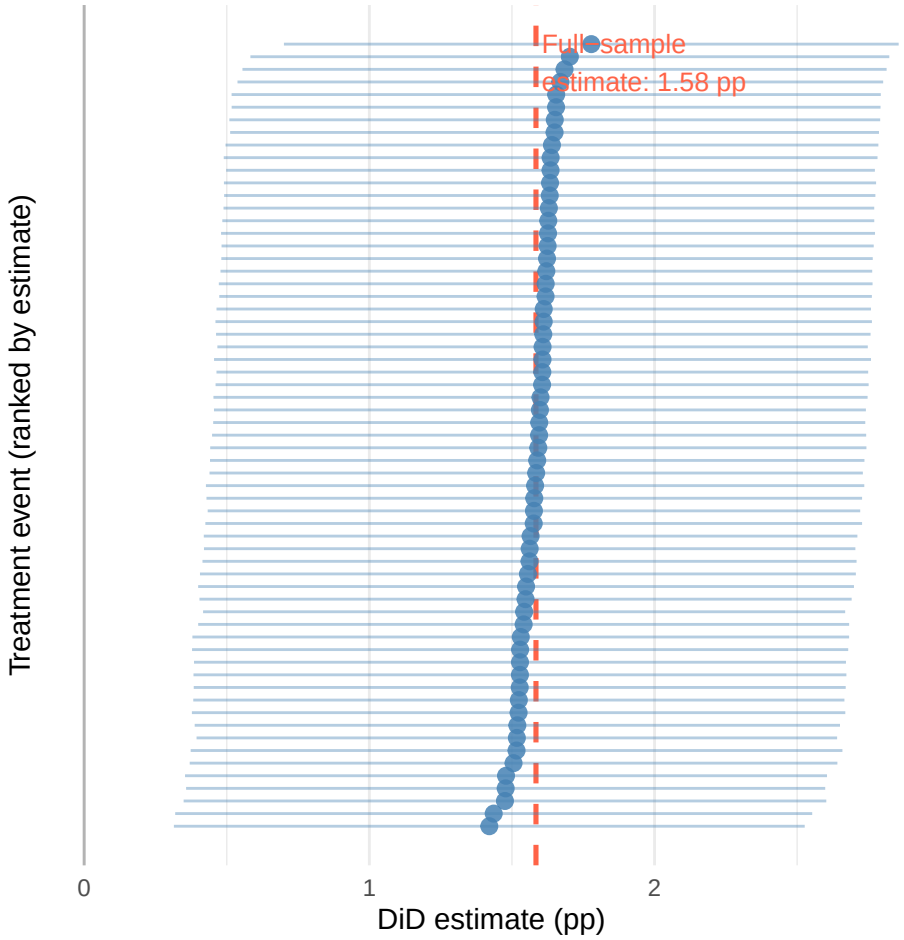


Figure O.1: Leave-one-out sensitivity of the graduation difference-in-differences estimate. Each point re-estimates the main specification after removing one treatment event (the treated college and its controls); the dashed line is the full-sample estimate. Bars are 95% confidence intervals.

Figure O.1 reports leave-one-out graduation difference-in-differences estimates. Each point is the treatment effect re-estimated after dropping one of the 63 treatment events. Estimates are sorted by value. Point estimates range from 1.42 to 1.78 pp and 63 of 63 remain significant at the 5% level.