Graded Walks: How Students' Dorm Locations Affect Grade Performance

David H. Feldman[1,2], Adam D. Jutt[1], Roswell Miller[1]

[1]Department of Economics, William & Mary, 300 James Blair Drive, Williamsburg, 23185, VA, US.

[2]Corresponding Author, dhfeld@wm.edu, 757-221-2372, ORCID: 0000-0001-6514-795X.

Keywords: Distance, Dorm, GPA

Abstract

We evaluate the impact of walking distance from dormitory to classroom on academic performance among first-year students at William & Mary. Leveraging a natural experiment created by the university's randomized dorm assignments, we analyze student-class level data from two cohorts (2016–17 and 2017–18). We find a statistically significant negative impact of distance, both for all classes and for the subset of "first classes of the day," as the distance measure may reflect actual travel more accurately for these classes. The effect is robust to student fixed effects and various class-level controls, and is particularly pronounced in large-enrollment and early morning classes. This is consistent with attendance-related frictions driving the effect. These findings support institutional policies that minimize first-year students' walking distance to class.

Keywords: Walking Distance; Dorm; GPA

Introduction

We use data from two entering classes at William & Mary to evaluate the impact on grade performance of walking distance from the student's residence to their classes. First-year students were assigned to dormitories in a manner that did not allow them to select for proximity to their classes. Additionally, students were assigned many of their classes before arriving at the university, so they had little direct knowledge of the physical campus. This random allocation offers us a natural experiment to identify the causal impact of distance on student performance, controlling for a wide range of other correlates with grades. To our knowledge, we offer the first causal evidence of walking distance on grades in a residential college setting.

The effect of distance is large and negative. We find a significant and meaningful grade penalty for students who live farther from their classes. For students whose class is one standard deviation further than the mean distance, the penalty is roughly 0.1 GPA points (on a 4-point scale). This is enough to move many students from an A to an A- or from a B+ to a B, for example. The estimates are similar when restricting the sample to classes which are the first of the day, a subset of classes for which our distance measure may better capture actual travel. We show that distance interacts with class size. The grade penalty grows with class size, which is consistent with an absenteeism story. On the other hand, the impact of distance on student performance does not seem strongly related to student preparation measured by SAT score.

Distance is a friction with consequences for many economic actions. Nations trade substantially more with countries that are closer than with otherwise similar countries that are further away (Anderson and Wincoop 2004). In health care, the distance to a clinic or hospital is a cost that can influence individual and community health outcomes (Buchmueller et. al 2006). Nutritional outcomes may be affected by the lack of easy access to grocery stores selling healthy foods (Fan et al, 2019).[1]

Distance is also potentially important to higher education outcomes. Some students have little access to nearby broad-access public institutions (Hillman 2016 and Hillman 2019). Students in these "education deserts" may experience worse educational outcomes than similar students whose commuting zones offer a robust set of affordable and accessible public two-year and four-year institutions. Acton, Cortes, and Morales (2024) use a rich data set from Texas to show that students are sensitive to the distance they must travel to public higher education institutions and that the effects are heterogeneous. Higher-income students respond to a lack of available two-year institutions by enrolling at four-year universities. Hispanic, Black, and other economically disadvantaged students respond by foregoing post-secondary education altogether. Distance as a friction is also at work internationally. Alba-Vivar (2023) used a difference-in-difference approach to show that the development of new transportation lines in Lima Peru had significant enrollment effects.

At residential colleges and universities, students who live on campus may have better outcomes, measured by classroom performance and retention, than students who commute into school. Students who live in campus dorms are closer to their classes, so they may have better attendance over the academic term. They may also benefit from having more extensive peer interactions on campus than commuter students who return home after class. Students who live on campus also may access support services more readily than commuters. For many colleges and universities these arguments are persuasive enough that they require incoming first-time-full-time students to live in college housing.

De Araujo and Murray (2010) use survey data from Indiana University - Purdue University Indianapolis (IUPUI) to evaluate the benefits of living on campus for students who are in their second-

[1] The "food desert" literature is complicated because demand (preferences) interacts with supply (distance to stores) in complex ways. People in food deserts face modestly higher prices. The effect is worsened if they are restricted to purchasing within their own census tracts.

year and beyond. These students aren't required to live on campus, so selection issues make identifying a causal impact of campus residency difficult. De Araujo and Murray use two instruments to account for this selection problem – distance from the student's hometown to Indianapolis, and a dummy for whether or not a student was denied on-campus housing because of space limitations. Both are correlated with whether a student lives on campus, but are arguably not related to measures of performance like GPA. They find that living on campus significantly improves semester performance and long-term academic success by .2 to .5 GPA points, or roughly half a letter grade. This result strongly supports required residence policies for incoming students at residential colleges and universities. In follow-up work with the same data, De Araujo and Murray (2010b) show that the main impetus for grade gains come from peer interactions rather than greater use of university support services.

Residence requirements, however, do not remove all distance effects. First-year students are assigned housing that may be near or far from their classes. Students who are assigned dorms that are farther from classes may miss more classes than students who are located closer. Selection is again a problem if students choose their dormitories. Students who perceive a larger academic benefit from living closer to class may request dormitories that are closer to their classes. The explanatory variable (distance) would then be jointly determined with academic performance, leading to biased and inconsistent estimates of how much dorm location affects grades.

In our data, most of the new first-year students were randomly assigned to residence halls, so students who might perceive an advantage of living closer to class could not select their residence hall on that basis. First-year students are generally not allowed parking privileges on campus, so travel across campus is on foot or by bicycle. We have no data on how students get to class, e.g. whether they walked or biked, but if students who live further away select into bicycle use at higher rates, then our result is an underestimate of the grade penalty for having a longer walk to class.

Data

Our data was provided by the office of institutional research at William & Mary. The data include individual-level information for all W&M freshmen for academic years 2016-17 and 2017-18. For each student, we observe each class they took during the fall and spring semesters, and the grade the student received. We also have information about each class that may be related to average grade performance. This includes the total number of credit hours, the time of day and the days of the week the class met, the total enrollment in the class, the department that offered the class, whether the class was one of two small-section general education courses each first-year student takes, and the academic building where the class met. We also know the assigned residence hall for each student. Knowing the residence hall and academic building allows us to measure the walking distance from the student's dorm to that class. We use Google pedometer to calculate distance for each dorm-academic building pair. Finally, we have access to each student's SAT score as a proxy for ability.

Variable	Obs	Mean	Std. dev.	Min	Max
Grade	27,376	3.414494	0.753567	0	4
Distance (km)	27,376	0.687707	0.26513	0.01	1.64
SAT	27,376	1373.955	118.3339	890	1600
Credits	27,376	3.096983	0.870952	1	5
COLL Class	27,376	0.215627	0.411264	0	1
Enrollment	27,376	59.63329	63.38092	5	283
Start Time (Mins)	27,376	716.5496	147.2834	480	1200

Table 1: Summary Statistics

The summary statistics for the student and class-level variables are given in Table 1. Since observations are on the student-class level, a student's characteristics are counted for each class that the student took. For instance, if a student is enrolled in five classes, that student's SAT score is counted five times in the average for the summary stats. Similarly, the characteristics of a class are counted multiple times: once for each student in the class. Actual enrollment is the only variable for which this structure systematically skews interpretation of the summary stats. Classes with higher enrollment will be counted more times in the average, so the “mean” for class size in our dataset, 60, is significantly larger than the actual average class size at William & Mary. Also, note that the summary statistics exclude zero credit classes. These classes -- typically corequisite “discussion sections” for large introductory courses -- are ungraded, so they cannot be included in our regressions.

Methodology

We regress student grade performance on the distance between their dorm and the academic building where the class takes place. Student performance is captured by a numerical representation of the course grade. An A in the class is recorded as a 4.0, an A- is a 3.7, B+ is 3.3, B is 3, etc. Implicit in this setup is the critical choice to look at individual class grades instead of the full semester GPA. This decision is motivated by a desire to measure the distance friction precisely, without sacrificing valuable information about class characteristics. As we shall see, using these characteristics to account for variation in class grade is critical to uncovering a strong distance effect.

We use two distinct methods to control for student-specific characteristics that may affect grade performance. The first is the student's SAT score. For all its shortcomings, SAT scores are correlated with academic performance and later earnings. See, for instance, Chetty (2023), University of California (2020), and University of California Senate (2020). The second method is the use of student fixed effects, which allows us to control for unobserved and unchanging student characteristics at a much more thorough level. Each student takes eight to ten classes during the academic year, so this fixed effects model gives us a powerful way to minimize omitted variable bias. We offer our main specification in two versions, one with SAT score and one with student fixed effects.[2]

[2] For robustness, we will also consider additional fixed effects, like dorm effects and subject effects. Note that, because dorm does not typically vary within the first year for a given student, we can consider dorm effects with the SAT control, but not in the student fixed effects model.

We also control for a set of class characteristics. These include number of credit hours, start time, enrollment, whether the class is a small-section general education class, and fixed effects for broad curricular category (STEM, humanities, social sciences).[3] We also include year and semester dummies.

Each of those variables is plausibly related to grade performance. Classes worth one credit, for instance, often are graded very leniently. For first-year students, examples of one-credit classes include some lab sections and individual or group music lessons. Start time and class size may play a role through an attendance channel. Early morning classes and high enrollment classes might be subject to higher rates of absenteeism that produce less learning. Small general education classes are known to be graded more leniently than other classes, on average, and attendance is usually part of the grade. The class category controls account for any significant grading differences between broad academic areas. STEM classes, for instance, may be “harder" overall than humanities classes because more of them are mandatory introductory courses that winnow the pool of students who want particular academic pathways into programs where seats are limited.[4]

Labs in particular deserve a bit more attention. For many introductory STEM courses, e.g. General Chemistry 1, Organic Chemistry 1, and Intro to Molecules, Cells, and Development, students register for the main, three-credit class and a co-requisite, one-credit lab. Students receive independent grades for the course and the lab, which are coded under different course registration numbers (CRNs) in our dataset. We elect to keep these labs in our dataset as distinct courses.[5] For a few courses, most commonly Calculus 1 and 2 in the dataset, “lab” sections are coded under the same CRN as the lecture section for a collective four-credit unified grade. These “labs” typically involve students working on

[3] Though for robustness we also run the main model with the finer set of department-level fixed effects.
[4] See Achen and Courant (2009) for a discussion of grades as a complex equilibration mechanism in an academic market.
[5] Results are nearly identical if one chooses to drop labs from the set.

practice problems for an hour a week with a teaching assistant. In these cases, we do not consider the labs as distinct courses.

Selection Into Distance

To make causal claims we must make a strong case that distance to class is random and not subject to selection. If students can select into classes closer to their dorms or select into dorms that are closer to class, and their selection criteria are correlated with their performance, the coefficient on distance in most of the models could be biased. We take up this issue of selection now.

During the 2016-17 and 2017-18 academic years, first-year students were assigned to dormitory spaces in two distinct ways. Roughly ten percent of first-year students were classified as "Monroe Scholars,"[6] and these students were mostly assigned to the two honors dorms on campus, Monroe and Willis, though they could opt to live elsewhere. The remaining 90 percent of the incoming class were asked if they preferred larger dorms (140+ students) or smaller dorms (less than 140 students). They did not select specific dorms and were not guaranteed that they would receive their preferred size choice. Thus, within this set of non-honors students, selection into dorms closer to or further from academic buildings in a manner correlated with performance seems highly unlikely.

Dorm	Avg Dist (km)	Std. dev.	Freq.
Barrett	0.59333334	0.118772	6
Botetourt	0.80549308	0.229067	6,652
Brown	0.71533179	0.300621	1,296
GGV	0.83886826	0.109771	7,325
Hunt	0.69104301	0.250131	930
Jefferson	0.38352108	0.226889	3,249
Lemon	0.20333333	0.21362	3
Monroe	0.46494662	0.319736	2,529
Reves	0.63584616	0.231367	65
Willis	0.47854465	0.251192	907
Yates	0.64626643	0.14323	4,414
Total	0.68770675	0.26513	7,376

[6] We will primarily refer to these students as "honors students."

Table 2: Average distance to class by dorm

Honors students, however, pose a problem. Monroe Hall and Willis Hall both have low average distance to class relative to other residence halls (see Table 2). Accordingly, we drop those students from the data for our initial regressions. We do offer one model using the students housed in Monroe and Willis, in which we both control for residence in these dorms and add an interaction of honors dorm residence (as a binary variable) and distance, to see if the distance effect holds for this group of relatively homogeneous students. Note that not all students in Monroe and Willis are honors students. If an honors student and non-honors student decided to live together, they could do so in one of the honors dorms.

There is the separate possibility that students, once knowing their residence hall, could select into classes in buildings near said hall. This seems unlikely as well. First-year students are motivated to find spots in required classes, classes in their intended concentration, or possibly classes perceived as easier. Selection on distance seems unlikely to drive many decisions. Additionally, students register for some of their classes before arriving on campus, which further decreases the chances that they would employ distance-conscious selection. Once students are more familiar with campus life and locations, students might use distance as an additional criteria in class selection. We test for this by evaluating if the impact of distance on grade performance diminishes in the second semester, and find no evidence of that, as will be discussed.

SAT Bin	Avg Dist (km)	Std. dev.	Freq.
Under 1000	0.77406593	0.219074	91
1000-1100	0.6716145	0.232708	607
1100-1200	0.69405567	0.23867	1,832
1200-1300	0.69735485	0.248532	3,686
1300-1400	0.70724591	0.252653	7,222
1400-1500	0.69376586	0.269872	9,695
1500-1600	0.62993165	0.295237	4,243
Total	0.68770675	0.26513	7,376

Table 3: Average distance to class by SAT bin

Data on student SAT scores offers another bit of evidence that selection is not an issue. Table 3 shows that there is no meaningful relationship between SAT score and the average distance to class. Average distance to class drops for SAT group 1500-1600, but this is because honors students tend to have high GPAs and are, as discussed, primarily housed in dorms which tend to be close to classes. Recall that we drop these students from the main regressions. Distance is also slightly larger on average for those with SAT below 1000, but there are very few such students, so the difference is likely pure noise.[7] So, there is no evidence of any meaningful selection into distance based on SAT score. However, distance could be correlated with non-cognitive, performance-relevant skills like drive and organization. This is another advantage of having models which utilize student fixed effects to control for unobserved student characteristics that are related to grade performance.

Random dorm selection establishes a low risk of selection bias in our restricted sample. After the first semester, however, students have more information available to them. They have experienced a semester of walking to class and may have a sense that distance affects their performance. Hence, we may see students with high first semester distances and/or students with low first semester grades re-sort into shorter distances for the second semester. This would bias our results toward zero, as low-GPA students moving toward lower distance will obscure any “actual” negative relationship between distance and GPA. If we find evidence that re-sorting occurred, we should drop second semester observations from our sample and only consider fall classes. If we find evidence that there was little to no such resorting, we can pool all four semesters.

The most direct way to test whether there is a different relationship between distance and performance in the second semester, and thus re-sorting, is to run a regression with an interaction of

[7] Dropping these students has essentially no effect on results

distance and a semester dummy. In this model, the coefficient on distance represents the first semester distance effect, while the sum of the coefficient on distance and the coefficient on the interaction term represents the second semester effect. If this interaction term is significant, we have evidence that distance affects grades differently across the two semesters. A more careful explanation of the control variables in this model can be found in the description of model 2 in the results section. The distance coefficients are found in Table 5 (model 7), and they suggest that there was no significant, intentional re-sorting in the second semester. The coefficient on the interaction term is highly insignificant, so we can treat the semesters as similar and pool the data.

Results and Analysis

	(1)	(2)	(3)
	SAT Control	Class Controls	Class Controls
		(SAT)	(Student FE)
Distance	-0.0298	-0.114***	-0.121***
	(0.0189)	(0.0183)	(0.0220)
SAT	0.00167***	0.00162***	
	(0.0000482)	(0.0000459)	
2017		-0.0443***	
		(0.00924)	
Spring Semester		0.0354***	0.0244***
		(0.00915)	(0.00752)
Credits		-0.168***	-0.156***
		(0.00620)	(0.00580)
COLL class		0.310***	0.302***
		(0.0143)	(0.0122)
Enrollment		-0.00154***	-0.00165***
		(0.0000875)	(0.0000771)
Start Time		0.000166***	0.000171***
		(0.0000326)	(0.0000284)
Other		0.0571***	0.0609***
		(0.0214)	(0.0199)
STEM		-0.163***	-0.186***
		(0.0121)	(0.0116)
Social Sciences		-0.174***	-0.187***
		(0.0122)	(0.0110)
Effect at Mean Dist	-0.02	-0.078	-0.083
Effect 1SD Above Mean	-0.028	-0.108	-0.114
Adjusted R-squared	0.064	0.150	0.432
Observations	23940	23940	23940

Table 4: Standard errors in parenthesis. *p<0.10 **p<0.05 ***p<0.01

Our main results are found in Table 4. We begin with a simple model using distance and one control for student ability (SAT score) and then add more controls and restrictions in subsequent models to explain more variation. In the first model, distance has the predicted negative impact on grade performance, but the estimate is weak and imprecise. Much stronger results appear in model 2 after controlling for class characteristics (start time, enrollment, an indicator for whether the class is part of the college's first-year general education curriculum, credit hours, and curriculum category). This increases the point estimate to -0.114, a result significant at the 1% level, and increases R-squared to 0.15.

Model 3 is similar to model 2, but the SAT control is replaced with a student fixed effect. Because SAT does not vary across observations for a given student, we cannot use both of these controls simultaneously. Ultimately, both methods seek to address the issue of student heterogeneity. Since SAT scores are an imperfect measure of student ability and these test scores do not capture non-cognitive traits that correlate with grade performance, the fixed effects method is a more general way to control for student heterogeneity. On the other hand, since we only have eight to ten observations per student, much of the variation across the sample will be absorbed by the student fixed effects, perhaps reducing the power of the distance variable to capture relationships. This concern does not prove to be the case. In model 3, the distance coefficient remains significant at the 1% level. This provides more evidence that distance has a negative causal effect on performance, as we have controlled for omitted student-specific characteristics that are potentially related to distance and course outcomes. The point estimate is slightly larger in this model (-0.121) than in model 2.

For the models with a robust set of controls (models 2 and 3), we find significant coefficients greater than -0.1 in magnitude. These represent the expected grade penalty for walking an extra kilometer from dorm to a given class. Average distance to class is around 0.682 km, so the line in the table labeled “Effect at Mean Dist” rescales the coefficient to reflect expected grade penalty at this mean distance. “Effect 1SD Above Mean" reflects expected penalty at dist 0.943 km, as this is one standard deviation above the mean for distance.

As a rather extreme robustness check, not included in the table, we can replace all of the course controls in model 3 with a course fixed effect, such that the only regressors in the model are the distance variable, student fixed effects, and course fixed effects. The results are qualitatively similar. The significance of the relationship between distance and grade, however, falls to the 10% level, as the standard error increases, and the point estimate falls in magnitude. This loss of significance is not surprising, as the model contains over 5,000 fixed effects, meaning fewer than five observations per parameter. The trade-off is the loss of efficient, grade-relevant course characteristic controls in exchange for isolating variation in distance within each specific course. The key result is that the coefficient remains negative and meaningfully different than zero, which supports the overall finding of a robust relationship between distance and student performance.

There are other, less dramatic effects worth considering for robustness. For the model with SAT controlling for student differences (model 2), we consider a model which adds academic department fixed effects, a model which adds dorm fixed effects, and a model which adds both. For the model with student fixed effects (model 3), we consider adding department effects but cannot consider dorm effects (as for vanishingly few students does dorm vary across observations). The intuition for considering department is the possibility that two subjects (e.g. kinesiology and physics) coded as the same category (STEM) are of different average difficulties (physics courses tend to be harder) and different average distance (physics classes are housed in Boswell Hall, which is in a distant corner of campus). Dorm may be relevant because some dorms are closer to academic building than others, and because there is heterogeneity across dorms in terms of amenities and environment, which may be relevant to academic performance. The results for all four additional models prove similar to our main specifications: for all, we find a negative distance effect significant at the 1% level.[8]

[8] Both sets of fixed effects have major shortcomings. While both are relevant to distance and to performance individually, there is no reason to expect the relationship with distance and the relationship with difficulty to be correlated. Additionally, department fixed effects require many more variables than class category (STEM, humanities, social science, other) controls, and the extra variables do not significantly improve our understanding of variation in grade performance.

	(4)	(5)	(6)	(7)
	Enrollment Int.	Honors Included (w/ Int.)	SAT Int.	Semester Int.
Distance in Meters	-0.0446*	-0.113***	-0.145	-0.0962***
	(0.0246)	(0.0182)	(0.256)	(0.0256)
Enrollment#Dist	-0.00130***			
	(0.000394)			
Honors#Dist		0.102***		
		(0.0382)		
SAT#Dist			0.0000230	
			(0.000184)	
2nd Sem#Dist				-0.0347
				(0.0358)
Adjusted R-squared	0.151	0.161	0.150	0.150
Observations	23940	27376	23940	23940

Table 5: Standard errors in parenthesis. *p<0.10 **p<0.05 ***p<0.01

We believe that high distance increases the cost associated with class attendance and thus increases the frequency of absences and tardiness. One way to investigate whether attendance effects are driving the strength of these results is to consider class enrollment size more carefully in the model. This brings us to Table 5, which contains models that include an interaction term. The various interactions considered are enrollment, honors dorm residence, SAT, and semester. All of these models will offer more insight into the nature and scope of the distance effect. We begin with enrollment (model 4). We expect that students will be more likely to skip a class that has high enrollment than a class with low enrollment. Students are more anonymous in large introductory classes, so there is less of a psychological penalty for missing class. If we interact distance and enrollment, while still controlling for both, we expect a negative interaction coefficient, i.e. a negative partial effect of enrollment on the overall distance effect. This is what we observe. We see each additional student in class associated with a -0.0013 greater effect of distance on GPA, significant at the one percent level.

Model 5 is identical to our main specification in model 2, except that we include students in the two honors dorms, controlling for honors dorm status and interacting distance with honors dorm status. Because we cannot claim that this group of students is randomly selected, we include the interaction term to reflect that the relationship between distance and performance may not be the same for this group as for the rest of the student body. The reason we include this specification at all, given the lack of random selection, is that it offers the potential to increase the sample and indicate whether a similar effect holds for a group whose academic abilities are clumped at the right tail of the distribution. In this model, the main distance effect is nearly identical to that in the main specification (-.113). However, the honors interaction term is statistically significant and positive, almost as large (.102) as the non-interacted distance coefficient. This indicates that the negative effects of distance on performance may be significantly smaller -- or non-existent -- for the group of students living in honors dorms (the vast majority of whom, but not all of whom, are honors students). Honors students may attend class more consistently and be less likely to miss class because of the commute. Additionally, this could intersect with the peer effects literature; perhaps the environment of an honors dorm promotes class attendance.

We can also test more generally whether the distance effect differs across the ability distribution. Students with weaker predictors for academic success might be affected more strongly by distance. Students with better high school preparation may have better developed non-cognitive skills that insulate them from the potential negative effects of longer distances. Model 6 tests that hypothesis directly by adding to model 2 an interaction between distance and SAT score. The interaction is essentially zero. We cannot say with precision that students with weaker predictors of academic success are at a higher risk of struggling academically as a result of being assigned to a residence space farther from their classes. It is interesting that there appears to be a much weaker distance effect for honors students, but not for high SAT students generally. One possibility deals with the fact that honors students were not selected just via test scores, but also based on extracurriculars and leadership experiences. This may hint that drive and other non-cognitive skills are more critical than the cognitive skills captured by the SAT at insulating one from the negative effects of distance.

First Class Restriction

The distance measure we use may not precisely capture a student's walking behavior over the course of an entire day. Our implicit assumption is that a student travels directly from their dorm to their class for every class during the day. Especially in cases where a student has back-to-back classes, this is an imperfect assumption that likely raises measurement error in the distance variable. Students might occasionally have to walk significantly different distances than the distance measure implies, such as in situations with back to back classes. We can increase the precision of our distance measure by considering only each student's first class each day. We cannot know how a student gets from their first class to their second class in a given day, but we can be more confident that a student goes to their first class from their dorm. This approach may improve the reliability of our estimate (signal-to-total variance ratio) of the effect of distance, avoiding attenuation bias.

In this section, we will consider a similar set of models of the distance effect, but with the sample restricted to only those observations corresponding to classes which are the first class of the day for a given student each day they have it.[9] Thus, we are trading off the benefit of a potentially more accurate distance measure against the costs of losing the majority of our observations. This tradeoff could lead to higher point estimates of the impact of distance on grades: If the full sample contains many observations where our measured distance is uncorrelated with “actual distance” -- given by the path the student takes to class -- the point estimate on distance in the full sample regression is biased toward zero. However, while we may expect an increase in point-estimate magnitude, it is not clear what the net

[9] Alternatively, one could define the set of first classes to include all those classes which are first at least one day of the week. However, the stronger requirement is a better definition, because it ensures a more meaningful restriction from the full sample. Additionally, consider a class which takes place MWF, but is only the first class on Fridays. This class seems to resemble never-first classes better than it resembles always-first classes. Ultimately, results are essentially the same regardless of the definition chosen.

effects of this tradeoff will be on our standard errors. Additionally, we note that even this restricted sample may be subject to error in the distance measure. For some students, the first class of the day occurs late in the morning or in the afternoon, after the student has moved about the campus for other reasons (a trip to the library or dining hall). In later models we will further refine our distance measure by restricting the sample to classes that start at 9:30 in the morning or earlier.

	(8)	(9)	(10)	(11)	(12)
	SAT Control	Full Controls	Full Controls	Early Classes Only	Early Classes Only
		(SAT)	(Student FE)	(SAT)	(Student FE)
Distance in Meters	-0.0996***	-0.0849***	-0.0968**	-0.0912**	-0.161**
	(0.0334)	(0.0329)	(0.0450)	(0.0450)	(0.0816)
Adjusted R-squared	0.060	0.126	0.434	0.129	0.432
Observations	7991	7991	7991	4514	4514

Table 6: Standard errors in parenthesis. *p<0.10 **p<0.05 ***p<0.01

Table 6 holds results for the "first class" models without interaction terms. When we restrict our data to the first classes of the day, the distance effect is negative and statistically significant with only an SAT control (model 8). Models 9 and 10 represent our results using the full set of controls. Model 9 uses SAT to capture student-level differences, while model 10 uses student fixed effects instead. The point estimates of the effects of distance on grade performance are similar to their counterparts in the full data (models 2 and 3), though slightly lower in magnitude. The similarity of results for these models across the full sample and the first class of the day subsample again suggests that this is a robust result. The standard errors are generally larger for this first class of the day subsample. This does not necessarily suggest a weaker relationship with distance for first classes, but is instead more likely due to the significantly reduced sample size.[10]

[10] As with the full sample specifications, here we also consider a model with both student and course-level fixed effects. Unlike with the full sample specification, with the first class restriction the distance coefficient is insignificant. This is unsurprising, as the model contains over 5,000 fixed effects with less than twice that many observations.

Models 11 and 12 further restrict the data to classes that begin at or before 9:30 in the morning, with the former using SAT as the measure of individual effects and the latter using student fixed effects. We mentioned a desire to increase precision in the distance measure as a motivation for this restriction, but there is an additional reason we might expect a larger effect in these models: if high distance increases the propensity to miss class or arrive late, the effect may be more pronounced in early morning classes. Model 11 produces a point estimate very similar to models run before implementing the time of day restriction, while model 12 produces a much larger point estimate than any seen prior (-0.161). This last model provides evidence that the effect of distance on performance is exacerbated in the mornings.

It is worth considering more carefully why the estimated effect is so much larger when we use student fixed effects than when we use SAT as a control for individual student differences. The gap suggests that there are student characteristics not captured by the SAT score that are correlated with distance and with performance in the same direction, pushing the estimate upward. This could happen if students with higher non-cognitive skills -- such as good sleep and study habits, being well-organized, etc. -- are more willing to register to undertake early morning, high-distance classes than students with weaker non-cognitive skills. Though the gaps were smaller, we observe differences of the same sign (student fixed effects uncovering a more negative distance effect than SAT control) when comparing model 2 to model 3 and when comparing model 9 to model 10. In line with the same story, perhaps students with weaker non-cognitive predictors of success are more likely to select classes specifically because they are close to their dorm than are their more driven peers.

That is not a trivial finding. One might have expected that those with higher non-cognitive skills (e.g. self-awareness) would be more cognizant of the costs and potential grade penalty associated with distance, and accordingly select into closer classes at higher rates than those lacking such forethought. In this case, the gap would likely be the opposite of what we observed; we would have seen larger effect estimates for SAT control than student FE. So, the gaps we observe are interesting. Strong students appear more willing to brave adverse class locations and class times, not more likely to try to avoid them.

	(13)	(14)	(15)
	Enrollment Int.	Honors Included (w/ Int.)	SAT Int.
Distance in Meters	0.00112	-0.0890***	-0.877**
	(0.0453)	(0.0328)	(0.446)
Enrollment#Dist	-0.00170**		
	(0.000740)		
Honors#Dist		0.00896	
		(0.0732)	
SAT#Dist			0.000581*
			(0.000322)
Adjusted R-squared	0.126	0.134	0.126
Observations	7991	9088	7991

Table 7: Standard errors in parenthesis. *p<0.10 **p<0.05 ***p<0.01

Table 7 provides the interaction term results. There are two major differences between these results and the parallel results with full sample. The first is in model 14, where the honors interaction is insignificant and near zero, while in model 6 the honors interaction was significant at the 1% level. This may suggest that the distance effect is not significantly weaker for honors students, but it is more likely a reflection of the smaller sample size, as we only have around 1000 honors student-class observations in this model. The other difference is that coefficients are more significant in the SAT interaction model here (model 15) than they were in model 8. However, we still only have tenuous evidence that the distance effect is smaller for high-SAT students.

Conclusions and Policy Implications

For first-time full-time students in two cohorts of entering students at William & Mary we have found a consistently significant and meaningful negative relationship between the distance from a student's dormitory to class and the grade they earn in that class. The effect is robust to many controls that are arguably drivers of student performance, including student fixed effects, and the effect is stronger both for classes with a large enrollment and for classes that begin at or before 9:30 in the morning. This is a particularly important finding for colleges and universities, as they can control where first-year students live on campus.

Although we do not have direct evidence on attendance, the larger effects for early classes and for classes with larger enrollments are consistent with an attendance channel. This is because the temptation to skip class created by distance will be acted upon more frequently when other class characteristics also support skipping. The anonymity of larger classes supports skipping, as does the cost associated with waking up early. Missing class can have a negative effect on grades for several reasons. Most directly, it can affect participation grades in certain classes. More importantly, though, missing in-person courses often reduces learning, since time spent in class may be a meaningful portion of the time devoted to learning for many students.

Absenteeism is a well-known negative on student performance. Romer (1993) studied attendance and performance in economics classes across three college types (private mid-size, public-large, and liberal arts) and confirmed that absenteeism is both prevalent and that it significantly decreases learning in a broad set of courses and in a restricted sample of similar students that minimized selection issues. Durden and Ellis (1995) directly estimate the consequences of absences in principles of economics classes and find a nonlinear relationship between absences and performance. A small number of missed class has little discernable effect, but past a threshold number of absences the penalty rises. Flipping the attendance story around, Marburger (2006) shows that mandatory attendance policies both raise attendance and significantly improve exam performance.

As a different mechanism, longer distances could also impact grades by reducing the time available to students to work on assignments and classwork on account of commute times, but given the small range of distances in our sample and the relationships with other variables, we think class attendance is a more important channel.

Generalizing from one institution should be done with due caution. The average William & Mary student has high SAT scores and a relatively small variance of overall academic ability. Walking distances are also small. Given these characteristics, the size of the distance effect seems large. At institutions that enroll a larger fraction of students with a weaker college preparatory background, and where walking distances are substantially larger, the effect of distance could be considerably more pronounced.

There are several potential policy ramifications of this finding. First, colleges could be more intentional about which students are placed in which dorms. Because there is some evidence, albeit weak, that honors students and to a lesser degree high-SAT students are partially or fully insulated from the negative effects of distance, institutions could choose to designate their farthest residence hall or halls as the honors dorm(s). Second, colleges could assign priority registration to students in distant dorms, so that they have a greater opportunity to avoid early classes. Finally, universities planning to construct new buildings or otherwise expand their campus should prioritize placing new dormitories as close to academic buildings as possible, perhaps relegating other buildings like dining halls or rec centers to the periphery.

Institutional policies pertaining to live-on-campus requirements and students with cars are also relevant to this discussion, but -- because the dataset only considers first-year students on campus, with no information on vehicles -- we cannot speak to those policies specifically. Accordingly, future research should focus on conducting similar analysis with datasets that do contain that information. That is, research should focus specifically on 1) a comparison of the distance effect for on-campus and off-campus students, and 2) the effect that vehicles like bikes and cars have at mediating the distance effect for on and off-campus students, respectively.

Declarations

The authors have no relevant financial or non-financial interests to disclose.

The authors have no competing interests to declare that are relevant to the content of this article.

All authors certify that they have no affiliations with or involvement in any organization or entity with any financial interest or non-financial interest in the subject matter or materials discussed in this manuscript.

The authors have no financial or proprietary interests in any material discussed in this article.

References

Alba-Vivar, F. (2023). *Opportunity bound: Transport and access to college in a megacity* [Working paper]. https://fabiolaalbavivar.com/research

Achen, A. C., & Courant, P. N. (2009). What are grades made of? *Journal of Economic Perspectives, 23*(3), 77–92. https://doi.org/10.1257/jep.23.3.77

Acton, R. K., Cortes, K. E., & Morales, C. (2024, October). *Distance to opportunity: Higher education deserts and college enrollment choices* (IZA Discussion Paper No. 17394; also NBER Working Paper No. 33085). IZA – Institute of Labor Economics. https://doi.org/10.3386/w33085

Anderson, J. E., & van Wincoop, E. (2004). Trade costs. *Journal of Economic Literature, 42*(3), 691–751.

Buchmueller, T., Jacobson, M., & Wold, C. (2006). How far to the hospital? The effect of hospital closures on access to care. *Journal of Health Economics, 25*(4), 740–761.

Chetty, R., Deming, D. J., & Friedman, J. N. (2023). *Diversifying society's leaders? The determinants and causal effects of admission to highly selective private colleges* (NBER Working Paper No. 31492). National Bureau of Economic Research. https://doi.org/10.3386/w31492

de Araujo, P., & Murray, J. (2010). Channels for improved performance from living on campus. *American Journal of Business Education, 3*(12), 57–64.

de Araujo, P., & Murray, J. (2010). *Estimating the effects of dormitory living on student performance* [Working paper]. SSRN. https://doi.org/10.2139/ssrn.1615612

Durden, G. C., & Ellis, L. V. (1995). The effects of attendance on student learning in principles of economics. *American Economic Association Papers and Proceedings, 85*(2), 343–346.

Fan, L., Baylis, K., Gunderson, C., Ver Ploeg, M. (2018). Does a Nutritious Diet Cost More in Food Deserts? *Agricultural Economics*, 49(5), 587-597.

Hillman, N. W. (2016). Geography of college opportunity: The case of education deserts. *American Educational Research Journal, 53*(4), 987–1021.

Hillman, N. W. (2019). *Place matters: A closer look at education deserts*. Third Way. https://www.thirdway.org/report/place-matters-a-closer-look-at-education-deserts

Marburger, D. R. (2006). Does Mandatory Attendance Improve Student Performance?, *The Journal of Economic Education*, 37(2), 148-155.

Romer, D., (1993). *Do Students Go to Class? Should They?*, Journal of Economic Perspectives, 7(3), 167-174.

University of California Academic Senate. (2020). *Report of the UC Academic Council Standardized Testing Task Force (STTF)*. University of California.

University of California Office of the President. (2020). *Relationship of the SAT/ACT to college performance at the University of California*. University of California.