

# When Inputs Are Outputs: The Case of Graduate Student Instructors

Maxime Tant <sup>1</sup>, Eric Watelain <sup>2</sup>, Maxime Tant <sup>3,\*</sup>

<sup>1</sup> Stanford University and NBER, 520 Galvez Mall, CERAS Building, Room 522, Stanford, CA 94305, USA

<sup>2</sup> Harvard Graduate School of Education and NBER, Gutman Library 465, 6 Appian Way, Cambridge, MA 02138, USA

<sup>3</sup> Harvard Graduate School of Education, Gutman Library 469, 6 Appian Way, Cambridge, MA 02138, USA

\* Corresponding Author. E-mail: eric\_taylor@gse.harvard.edu

**Abstract.** We examine graduate student teaching as an input to two production processes: the education of undergraduates and the development of graduate students themselves. Using fluctuations in full-time faculty availability as an instrument, we find undergraduates are more likely to major in a subject if their first course in the subject was taught by a graduate student, a result opposite of estimates that ignore selection. Additionally, graduate students who teach more frequently graduate earlier and are more likely to subsequently be employed by a college or university.

**Keywords:** Higher education; Graduate student teaching.

## 1. Introduction

There are few cases in economics where inputs into the production function are also outputs in the production function. However, such is the case in the production of doctoral students. Doctoral students are essential inputs to significant outputs in higher education. They provide research support, generate peer effects on each other, and often teach undergraduate students. Doctoral students are also an important output in education. Graduate programs define and often rank themselves by the quality and quantity of doctoral students they produce. Moreover, among those who go on to work in higher education, the productivity of graduate students in their subsequent careers is another important "output" of doctoral programs in that it increases the prominence of the doctoral-granting organization, increases the status of the doctoral adviser, and contributes to the subsequent education of new scholars in the field.

To date, there is very little research on the productivity of graduate students and on the relationships between their graduate activities and subsequent careers, particularly with regards to their teaching responsibilities while enrolled in school. What research there is with regards to the teaching function of graduate students, however, suggests that undergraduates that have a graduate student as their instructor, particularly non-native English speaking graduate students, suffer worse outcomes than comparable students that have faculty instructors (e.g., Borjas, 2000). Moreover, there exists no evidence on how graduate students' teaching experiences contribute to their subsequent careers.<sup>1</sup>

This paper examines graduate student teaching as an input to two production processes: the education of undergraduates and the development of graduate students themselves. As such we attempt to answer two related research questions. First, we quantify the effect of graduate student instructors on the academic outcomes of the undergraduate students they teach. As outcomes we examine students' choice of major, course taking, and credits earned. While graduate students participate in several facets of knowledge production and instruction, this is perhaps the most common avenue for their participation. In this paper, we are distinguishing between the roles of being a teaching assistant (i.e., course support) versus serving as the primary teacher of a course (i.e., the

instructor); our analysis focuses on the latter. About 46% of undergraduate students at four-year colleges take courses that have graduate students as their primary instructors,<sup>2</sup> and in 2000, over 70% of graduate students had some teaching responsibility. At issue is whether using graduate students as instructors relative to other possible types of instructors (faculty or adjuncts) is a good way to allocate resources as measured by the outcomes of undergraduate students.

Second, we examine how teaching experience contributes to graduate students' subsequent academic careers. One of the justifications for using graduate students as instructors is that it provides essential training for their subsequent careers. We measure the relationship between teaching experiences and the subsequent academic outcomes and career choices of the graduate student. We examine the likelihood the graduate student completes their doctorate and is later employed at a college or university.

For our analysis, we use administrative data from Ohio's public universities. To answer the first question, we track the outcomes of undergraduate students who initially enrolled in college in the fall of 1998 or 1999. We identify the first course (e.g. Introductory Economics) that an undergraduate student took in a specific department (e.g. Economics) and observe whether these courses were taught by graduate students as opposed to full-time or part-time faculty. Using exogenous variation in faculty availability, we estimate the causal impact of graduate student instructors on undergraduate students' subsequent course-taking behavior.

To answer the second question, we study students who began their graduate program in 1998 or 1999, and who subsequently taught at least one course. We link university administrative data with data from the Ohio Department of Jobs and Family Services (ODJFS). ODJFS tracks employment outcomes throughout the state, and these data allow us to measure graduate students' subsequent earnings and location of employment, so long as they stay within the state. We use the combined data to relate the number of teaching opportunities graduate students experience to their subsequent academic and professional outcomes (within six years of starting graduate work). Although our data are limited to observing subsequent behavior in Ohio alone, six years after the start of their graduate program, we find 82.5% of graduate students in either our employment data (52.6%) or university administrative data (29.9%, Table 4). Therefore, we believe we are observing subsequent outcomes for the vast majority of students. Moreover, while limitations remain, this is the first large-scale study to shed light on the relationship between teaching and outcomes after graduate school.

Our research questions are fraught with concerns about selection. Graduate students are unlikely to randomly choose to teach a course, and other work suggests that undergraduates may actively shy away from taking courses taught by certain kinds of instructors (Bettinger & Long, 2010). We use multiple strategies to deal with these issues. To address the fact that students may sort non-randomly across courses, we present the results using (a) course-by-department fixed effects, which limits the identifying variation to students who took the same course at the same institution but from different instructor types, and (b) course-by-department-by-term fixed effects, which limits identification to students who took the same course at the same time but from different instructors. Additionally, we estimate the results using an instrumental variables approach that capitalizes on the natural fluctuations from term to term in the faculty available to teach. In terms of whether a graduate student chooses to teach, we use department fixed effects to look within departments. However, there may be residual selection bias and so we caution about the interpretation of the results. Still, using a variety of assumptions about the size and direction of this bias, we still believe these results shed light on this understudied issue.

Our results suggest that graduate students are effective instructors relative to faculty members—at least as judged by the measures of their student's subsequent academic progress we can observe. Undergraduates taught by graduate students in a given subject are more likely to subsequently major in the subject compared to their peers who take the same course from full-time faculty. However, we find no statistically significant differences in the number of subsequent credits earned in the subject. Given that we use random variation in students' exposure to graduate students, we argue that our estimates suggest causal relationships. Second, graduate students who teach more frequently are more likely to complete their doctoral degree in a timely manner and more likely to be employed

subsequently by a college or university. Regardless of whether we interpret these as causal or selection, the results suggest graduate student teaching benefits the sector. If our results only reflect selection, that selection, at least in our sample, identifies and incorporates effective future faculty into undergraduate production. By contrast, if our results reflect causality, then they suggest that undergraduate instructional experiences positively impact the short-run, academic job prospects of graduate students.

In the next section, we discuss existing academic literature on graduate student teaching. In Section 2 we describe our empirical strategies and the Ohio sample and setting. Section 3 presents and discusses the results, and Section 4 concludes.

## 2. Existing literature

The input-output duality of graduate student teaching has long been recognized in American colleges and universities. During the rapid growth of graduate student teaching after World War II, one presenter at the 1966 annual meeting of the American Council on Education stated that “universities, in their need to staff elementary courses, have thoroughly abused...the chief means of training prospective college teachers—the teaching assistantship” (Wise, 1967). Scholars in education, economics, and sociology have contributed to an academic literature on the topic which has largely followed the two research questions at the heart of this paper.

The first set of results has focused on the effectiveness of graduate students as instructors. Ex-ante, it is unclear what we should expect about the quality of graduate student teaching. On the one hand, teaching experience can be a meritocracy of sorts, and department administrators may attempt to place the “best” teaching graduate students in prime teaching slots. Yet some researchers have highlighted the poor working conditions of graduate students (e.g. McLeod & Schwarzbach, 1993). These researchers document how low wages, large workloads, and working without training or guidance may actually generate negative effects (Koehnecke, 1991). In fact, these stresses coupled with the underlying trend toward increased reliance on graduate student instructors have led to recent efforts to strengthen graduate student teacher unions at universities across the nation (Mattson, 2000, Sharnoff, 1993 and Vaughn, 1998).

Most of the empirical literature of graduate student instruction quality has focused on the effects of foreign graduate students on undergraduate students’ outcomes. Borjas (2000) studied graduate student instruction in economics at a large public institution in the United States. In that setting, undergraduates taught by foreign-born teaching assistants received lower final grades by about 0.2 grade points relative to US-born graduate student instructors. However, the negative effect was muted for foreign-born undergraduates and when the graduate student was given higher class-preparedness ratings by the undergraduates. Borjas’ (2000) results are consistent with preceding findings in economics instruction (Watts & Lynch, 1989), though in other subjects the limited evidence is mixed (Jacobs and Friedman, 1988 and Norris, 1991). Subsequently, Fleisher, Hashimoto, and Weinberg (2002) found no differential, and possibly positive, effects of foreign graduate student instructors when they received training and completed at least one year of their own graduate study prior to taking-on teaching responsibilities.

Beyond comparisons of foreign and native graduate students, empirical evidence on instructional effectiveness in the college setting is scarce. This is in stark contrast to the American K-12 setting, for which there are numerous studies estimating the range of outcomes students with different teachers have and how this variance is related to teacher characteristics such as training level and certification. There are two studies that document the variation in instructor effectiveness. For example, using variation in instructor quality at the US Air Force Academy and benefitting from the random assignment of cadets to course sections, Carrell and West (2010) estimate the standard deviation in instructor effectiveness as approximately 0.052 student standard deviations on a common end-of-course exam. Hoffman and Oreopoulos (2009) find similar estimates for common exam courses (0.055 standard deviations), but the standard deviation is larger for all courses. These estimates are somewhat smaller than the evidence on variability in K-12 teacher effectiveness (Hanushek & Rivkin, 2010).

Beyond the general variance of outcomes across instructors, the above two studies also look at the effects of particular instructor characteristics. Carrell and West (2010) find that students assigned less-experienced instructors and instructors without a terminal degree (two correlated characteristics) score higher in the short-run but perform less-well in subsequent classes that build on content from the first. Hoffman and Oreopoulos (2009) find that part-time instructors, which include both adjuncts and graduate students in their sample, have no differential effect on undergraduate outcomes compared to full-time faculty. Another study focuses on the relative effects of full-time versus part-time faculty. The authors find small positive effects for adjuncts compared to full-time faculty in terms of students' subsequent interest in the subjects; however, the effects vary by subject (Bettinger & Long, 2010). However, other work using the same data find that the percentage of instructors a freshman student has that are adjuncts during their first year is negatively related to college persistence (Bettinger & Long, 2006).

Our primary outcome of interest is students' subsequent engagement in a course as measured by decisions to major, decisions to take additional courses, and the total number of subsequent classes that students took. Our data limit us from being able to track performance in the current or any subsequent course. However prior literature on teaching quality in K-12 frequently equates good teaching with subsequent engagement in courses (e.g. Reeve and Jang, 2006 and Skinner and Belmont, 1993), and similar connections exist in higher education (e.g. Coates, 2005). While subsequent engagement may be the result of other graduate student characteristics (i.e. generating interest through entertainment, ease of grading or other mechanisms), the prior literature certainly finds a correlation with overall teaching quality.

A second set of papers focuses on the role of graduate student teaching as an input to their own course of study and development. Researchers have long-argued that graduate student teaching experience is essential to training future professors (e.g. Knotts and Main, 1999, Meyers and Loreto Prieto, 2000, Prieto and Altmaier, 1994, Slevin, 1992 and Smith, 2001). Most existing empirical evidence measures proximate outcomes for doctoral students, including degree completion and time to completion, which are a present concern for American higher education (Bowen & Rudenstine, 1992). A few studies have found that graduate students who receive teaching assistantships, as compared to other forms of financial aid, are less likely to complete their degree and those who do finish take longer to do so (Ehrenberg and Mavros, 1995 and Seagram et al., 1998). These effects may differ from discipline to discipline, and students with teaching assistantships do better than those who have no institutional support (Bair & Haworth, 2004).

By contrast, there is little empirical evidence on how teaching during one's time as a graduate student affects long-run outcomes, including subsequent labor market prospects, research productivity, and other academic and professional outcomes. In this paper, we begin to address this hole in the empirical literature using data which track graduate students' employment after degree completion.

### 3. Empirical strategy

#### 3.1. Outcomes for undergraduate students

Our first objective is to identify the effects of taking a class taught by a graduate student on undergraduates' future course-taking behavior and choice of major. We focus on variation in the instructors of undergraduate students' initial, often introductory, course in a given subject because that early experience may affect long-term interest and success in the subject area.

Eq. (1) describes the core of our approach: predict different outcomes,  $y_{ijk(t+1)}$ , for undergraduate student  $i$  in subject  $k$  at campus  $j$  at time  $t+1$  as a function of whether they were taught by a graduate student in their introduction to the subject at time  $t$ .

Operationally,  $GradStuijkt$ , the explanatory variable of interest, will be equal to one if the student took her first course in the subject from a graduate student and zero if she took the course from someone other than a graduate student. If she took multiple courses in a given subject during the first term of exposure  $GradStuijkt$  will equal the proportion of courses taught by a graduate student.

We use this structure to study three different outcomes: (i) whether student  $i$  took any additional courses in subject  $k$  beyond time  $t$ , (ii) the total number of credit hours taken in subject  $k$  beyond time  $t$ , and (iii) whether student  $i$  majored in subject  $k$ . Since undergraduate students' course and major choices often lay the groundwork for future careers, graduate student instruction may affect long-run labor market outcomes for both individuals and at a more macro-level, through these nearer-term decisions. These mechanisms could, for instance, change the composition of available types of labor if undergraduates sort among disciplines as a result of exposure to high- or low-quality graduate student teaching.

In addition to the explanatory variable of interest, we control for a vector of student characteristics,  $X_i$ , which includes demographics (gender, race/ethnicity), initial ability (measured by ACT score), state of residence, and year of college entry (an indicator for being in the fall 1999 versus 1998 cohort). We also account for each student's expressed interest in and exposure to subject  $k$  with the vector  $Z_{ik}$ , which captures whether the subject is the student's intended major and the number of credits hours taken in the subject during the first term of exposure. Our specification also includes fixed effects for the term during which first exposure occurred,  $\pi_t$ , the subject,  $\tau_k$ , and campus,  $\sigma_j$ . Finally, we also include a control for exposure to adjunct instructors in initial courses,  $Adjunct_{ijkt}$ , which we define in the same way as  $GradStuijkt$ . Separating out  $Adjunct_{ijkt}$  clarifies the interpretation of our coefficient of interest so that  $\beta$  only measures the effect of graduate student instructors relative to full-time faculty. <sup>3</sup>

Because students are observed in multiple subjects, we address within-student correlation by clustering the standard errors at the student level. Additionally, since the treatment of interest (which type of instructor the student is exposed to) varies by section, we also report standard errors clustered at the section-by-course level.

To interpret our estimate of  $\beta$  causally requires the assumption that undergraduates taught by graduate students are not systematically different from other undergraduates, at least conditional on the included variables in Eq. (1). A likely violation is the non-random way in which students and instructors are sorted into course sections. Such sorting may occur across courses within departments, across terms, or across sections of a course taught in the same term. For example, students may have prior beliefs about the relative effectiveness of graduate student instructors and so may try purposely to avoid them and enroll in course taught by full-time faculty members. Thus we report estimates which develop our core approach (Eq. (1)) in three ways.

First, we present results which add course-by-department fixed effects; we define departments as the interaction of subject and campus. This limits the identifying variation to students who took the same course at the same institution, but from different instructor types. The restriction addresses differences in the distribution of instructor types across departments and differences in how departments assign courses to different instructor types. Still, student decisions about which term to take a course may be influenced by expectations of future instruction type options. Thus, second, we report estimates with course-by-department-by-term fixed effects, limiting identification to students who took the same course at the same time but from different instructors. To the extent students are randomly assigned to course sections in a given term, these course-by-department-by-term fixed effects estimates will be unbiased. But even this assumption is potentially too strong. For example, full-time faculty may be more likely to teach honors sections or daytime sections with regular or late-hours sections left to graduate students. Students planning to major in a given subject may actively seek-out full-time faculty per se or honors sections.

To address the remaining selection we employ a third strategy: an instrumental variables approach to limit identification to variation in instructor type that is uncorrelated with prior, unobserved student interest in the subject. Specifically, we instrument for instructor type, identified by the variables  $GradStuijkt$  and  $Adjunct_{ijkt}$  in Eq. (1), using each department's term-by-term variation in the proportion of courses taught by professors of different rank. Departments face natural fluctuations from term to term in the faculty available to teach as professors are hired, retire, take sabbaticals, and replace courses with research grant funds. Departments also experience temporary fluctuations in the demand for particular courses that translates to more course sections. Since graduate students are one

solution to short-run staffing fluctuations, we hypothesize that variation in the proportion of full-time professors from term to term will predict undergraduates' exposure to graduate student instruction. Likewise, this instrument should predict exposure to part-time instructors, and we therefore instrument for adjuncts as well.<sup>4</sup>

The first-stage, described by Eq. (2), models the probability of exposure to graduate student instruction as a function of the department-specific variation from term to term in the proportion of courses taught by full-time professors relative to a department steady state.

First, for each department  $k$  at campus  $j$ , we calculate the number of assistant, associate and full professors teaching during time  $t$ , where  $t$  is the term of a specific school year.<sup>5</sup> To address differences in scale across departments, we normalize these counts by the typical number of full-time professors teaching in the department during the same term.<sup>6</sup> These normalized counts are captured in the variables  $AsstProfjkt$ ,  $AssocProfjkt$ , and  $FullProfjkt$  in Eq. (2). Second, we include campus and department fixed effects,  $\sigma_j$  and  $\tau_k$  respectively, so that we are identifying based on deviations from a given department's average reliance on different types of faculty from term to term.

As reported below, deviations in the proportion of faculty teaching in a given term are highly correlated with the probability that an undergraduate took her initial course from a graduate student. However, for our instrumental variables estimates to be unbiased, this distribution of teaching across ranks should not affect undergraduate outcomes directly. Consider, for example, a department in long-run decline where the number of full-time faculty shrinks in successive time periods and replacing them is unlikely to take place. To address short-run course commitments, the department would likely turn to adjunct and graduate student instructors. In this scenario, students might be reluctant to continue their studies in a collapsing department, even if instructor type has no independent relationship with undergraduate students' future academic trajectories. We are not particularly concerned about this threat, however, as prior analysis (Bettinger & Long, 2010) suggests that such declining departments, representing just three percent of all departments across all schools, are rare in our analysis sample.

### 3.2. Outcomes for graduate student instructors

Our second objective is to explore the relationship between individual graduate students' teaching experiences and their own academic and professional outcomes. One way to interpret this analysis is causal: that the teaching experiences could affect the subsequent outcomes of the graduate students. However, given that graduate students do not randomly choose whether to teach, we acknowledge the likely selection involved. Therefore, we pursue this analysis in part to gain insight into the selection of graduate student instructors from among the large pool of graduate students in a department. The limited evidence on graduate student teaching, notably Borjas (2000), suggests that understanding which graduates are teaching is critical to understanding the effects of their teaching on undergraduate students.

We examine this relationship by estimating the probability that graduate students complete a doctoral degree, and separately, gain employment at another college or university in Ohio within six years of beginning their graduate program as a function of their teaching experience during graduate school. Eq. (3) describes our approach.

The key predictor variable,  $TermsTaughtijk$ , is the number of terms graduate student  $i$  taught an undergraduate class (within six years of starting graduate school). Our specification also includes fixed effects for department-by-campus  $\tau_{kj}$ . Standard errors are clustered at the department-by-campus level.

We estimate Eq. (3) using our sample of graduate students who began their graduate programs in 1998 or 1999, and subsequently taught at least one undergraduate class.<sup>7</sup> Thus  $TermsTaughtijk$  cannot be zero, and the correct interpretation of  $\beta$  is how each additional term teaching, conditional on teaching in one term, predicts the outcome of interest. The data available for this study do not include graduate students with zero teaching experience, thus we cannot estimate the margin of teaching versus no teaching experience.

Selection is again a primary concern for any casual interpretation of estimated effects on graduate student outcomes. Notably, when selecting an instructor from among a pool of graduate students,

departments may be more likely to select graduate students who are planning to seek a faculty job after graduation; alternatively, they may try to protect those students from teaching so that they can focus on building a research portfolio. From the perspective of the graduate student, the ability to secure teaching opportunities may be a function of interest, motivation, and other latent characteristics that are correlated with academic and professional success. While fixed effects reduce some sources of bias by identifying on differences between students in the same department, at this time we do not have additional empirical strategies to address residual selection bias. As a result, we caution against interpreting these estimates causally; rather, we believe these results help clarify the implications of our estimates for undergraduate students.

### 3.3. Data and study sample

To estimate the specifications described above, we use administrative data provided by the Ohio Board of Regents (OBR). The undergraduates in our sample are first-time, full-time, traditional-age (18 to 20 years) freshman who enrolled at one of twelve public, four-year colleges and universities in Ohio during the fall of 1998 or the fall of 1999. The graduate students in our sample were enrolled at one of the same twelve institutions and taught undergraduates from one or both of these two cohorts at least once during the undergraduates' college careers.

For undergraduates the data include demographic details, test scores, major, college courses taken, and the instructor(s) for each course. The OBR data track students across campuses within the Ohio public higher education system, both two- and four-year institutions, which enables us to observe students even when they move from one school to another, or withdraw from any one course. Thus, for example, we can link a student's first biology class at Ohio State with a biology major later declared at the University of Cincinnati. Similarly we can track students who took their first biology course from a graduate student, even if the student withdrew from the course and did not receive a grade. Unfortunately, we do not observe outcomes at private institutions and institutions outside Ohio, and as a result, we cannot distinguish students who transfer to a school outside the Ohio public system from students who drop out of college entirely. Previous work by Bettinger (2004) suggests this is not likely to impact our substantive findings since the rate of transfer to out-of-state or private institutions is quite small among our student sample.

Additionally, for most undergraduates the data also include information taken from forms students complete when sitting for the ACT exam. Particularly important for our identification strategy, these ACT data include the student's intended major stated before they matriculate. Given the value of knowing students' major intentions, we limit our undergraduate analysis sample to individuals who took the ACT. In our data, 89% of four-year college students in Ohio take the ACT.<sup>8</sup>

Table 1 reports summary statistics for these undergraduates and their course taking behavior. Consistent with national trends, the sample is more female and less black than Ohio's college-aged population, as shown in the top panel. They are also spread across Ohio's selective and non-selective public institutions which span the state's urban, suburban, and rural areas. Additionally, the proportion of Ohio high school graduates who immediately enroll in college is near the national average (Mortenson, 2002).

The bottom panel of Table 1 describes students' experiences during their first course in a given subject and subsequent behavior in that subject.<sup>9</sup> Across all courses, undergraduates are taught by graduate students in about one-fifth of their courses in a subject, and undergraduates take, on average, five additional credits in the same subject within a year of taking an introductory course.

The right hand column of Table 1 focuses on our undergraduate analytic sample. In our analysis of undergraduate outcomes we focus on undergraduates' first course in a given subject, and further limit the sample to only courses where at least 10%, but no more than 90%, of sections are taught by graduate student instructors. We also exclude remedial or developmental courses because of their unique design and objectives. This restricted sample of undergraduates is not very different in composition from the full student sample shown on the left of Table 1; however, the restricted sample is drawn slightly more from the selective institutions which are more likely to have graduate programs.

Table 2 compares, by subject, the proportion of undergraduates who took their first course from a graduate student instructor calculated for the full sample and for the sample of courses with variation in instructor type. The restriction affects inclusion of social science and science courses more than some others, such as English and math/statistics, and particularly affects professional departments.

Table 3 characterizes the instructors to which the undergraduates in our student sample were exposed. The first two columns of Table 3 describe the faculty who taught any undergraduate: Column 1 shows the raw means at the course-level, and Column 2 weights by the number of students taught. Approximately 44% of students were taught by faculty with a PhD, 21% by adjuncts, and 14% by graduate students. The right side is limited to faculty of courses where the course was the undergraduate's first introduction to a department. Graduate student instructors were much more likely to be women or minority than full-time instructors.

In addition to the basic instructor characteristics summarized in Table 3, the Board of Regents provided additional data on graduate student instructors. The graduate student data include course taking information (similar to the data for undergraduates), longitudinal data on credits earned and GPA, and detailed degree information. Additional data record the details of employment (i.e., employer, earnings) between the last quarter of 1998 and the second quarter of 2006, but only for employment in Ohio.

Table 4 provides mean characteristics for the sample of graduate student instructors in our analysis of graduate student outcomes; students who began their graduate studies at an Ohio public institution in 1998 and 1999. All students in the sample taught undergraduates at least once (within six years of beginning graduate school); 22% taught just once, while 26% taught in five or more terms. Within six years of starting graduate school, 16% had completed a PhD, and 53% were employed in Ohio with 15% working at an Ohio college or university (excluding their own graduate institution). Another 30% were still enrolled at their university during the sixth year. As Column 2 shows, these characteristics are fairly consistent when focusing exclusively on students from academic departments.

## 4. Results

### 4.1. OLSs estimates of the effects of graduate student instructors

We begin with OLS estimates of how early exposure to graduate student teaching affects future undergraduate interest in the subject. In general, the OLS estimates suggest negative effects although, as previously discussed, these may be influenced by non-random sorting of instructors and students even after the inclusion of different fixed effects. Table 5 reports OLS estimates predicting the total number of subsequent credit hours taken in a given subject with varying fixed-effects approaches to address some forms of selection. Columns 1–4 of Table 5 report outcomes following exposure to graduate student teaching in an initial course in a subject during any term. The naïve estimates in Column 1, which exclude any fixed effects would suggest having a graduate student instructor reduces the number of future credits taken in the same subject. When we account for systematic differences from term-to-term and fixed differences across campuses and subject areas, Column 2, the effect is reduced by about 60%. Similarly, when we further control for differences across courses in the same department, Column 3, the effect is again reduced by about half; it approaches zero but is still statistically significantly negative when we cluster at the student level.<sup>10</sup> Finally, if selection patterns differ over time, a course-by-time fixed effect will reduce bias. When we add these last fixed effects in Column 4, the magnitude of the negative effect increases.

The differences in estimates across Columns 1 through 4 of Table 5 reinforce the importance of considering how departments, courses, and sections of the same course differ in terms of how students and instructors are assigned to each other. Accounting for non-random variation in which undergraduates enroll in which sections, and who teaches those sections is critical to identifying the true effect.

Coefficients on ACT score and the subject being the undergraduate's major field are also reported in Table 5. Across specifications, higher ACT scores predict more courses taken in any subject, and particularly in the student's major (the interaction of ACT and major). The estimated marginal effect



of “in intended major” is zero when ACT score is a very low seven or eight points (such an ACT score is more than three standard deviations below the mean for our sample).

On the right hand side of Table 5 we report estimates of credits taken following exposure to graduate student teaching during only the undergraduate's first term (as opposed to any term). We hypothesize that selection may be weakest during the first undergraduate term since at this early stage new matriculants know the least about differences in subjects, specific courses, and instructor types. Additionally, they may have the least flexibility in scheduling given the need to complete baseline requirements. For this sub-sample, Columns 5–7 replicate Columns 1–3 (there is no Column 4 analog since time cannot vary within a cohort when the analysis is restricted to students' initial fall term). The estimated effect of exposure to graduate student instruction in the first term is, as before, negative.

#### 4.2. Selection on observables

The fixed effects strategies employed in the specifications in Table 5 address selection across departments, courses, and the timing of enrollment in a course. They do not, however, address selection across sections within a course. In Table 6 we explore the extent of this sorting by predicting the likelihood of having a graduate student instructor using specifications similar to Table 5 except, of course, excluding the graduate student regressor which has become the dependent variable. Students with higher ACT scores are less likely to have a graduate student instructor, though this relationship is practically meaningful only when the subject is the student's intended major. Although the coefficient on “in intended major” is positive and large, undergraduates with ACT scores near or above average do not have a higher probability of seeing a graduate student in their intended major. Additionally, undergraduates attending selective institutions are more likely to have graduate student instructors. These patterns generally hold across fixed effects specifications, and whether considering all courses or focusing just on courses in students' first term.

#### 4.3. Instrumental variables estimates

Given the sorting concerns observed in Table 6, we remain concerned about issues of selection across course sections as well as other unobserved differences, which our previous OLS specifications fail to adequately address. In response to these lingering sources of potential bias we employ an instrumental variables strategy that uses variation over time in the composition of each department's faculty as an instrument for exposure to graduate student teaching. Ultimately this approach results in estimates quite different from ordinary least squares.

Table 7 reports coefficients and F-tests on excluded instruments from first stages predicting undergraduates' exposure to graduate student instruction as a function of variation in faculty composition within departments over time (Eq. (2)). As before, the left side of Table 7 uses the sample of all courses while the right side focuses on courses taken during the first term. The pattern across specifications and samples is fairly consistent: the probability of taking a course from a graduate student instructor is significantly related to departmental deviations from the steady-state proportions of most full-time faculty of different ranks, particularly assistant professors. The results are strongest when focusing on courses during the first term. And, given the coefficient signs, graduate students appear to be a substitute for assistant professor teaching in particular.<sup>11</sup> However, the more conservative standard errors clustered at the section level make the pattern of statistical significance somewhat less consistent. In general, F-tests suggest that the combined excluded instruments do explain substantial variation in undergraduate exposure to graduate student teaching. However, when clustering at the section level, the F-statistics are much lower and sometimes below conventional thresholds.

Table 8 reports two-stage least squares estimates of the effect of graduate student teaching on undergraduates' subsequent interest in the subject. Each cell represents the coefficient of interest from a separate regression. All regressions in Table 8 include campus-by-department-by-course and term fixed effects, as well as the standard set of student-level covariates employed across all analyses. The underlying first stages are reported in Columns 3 and 6 from Table 7 for the left and right sides of Table 8 respectively.

In contrast to the OLS results in Table 5, the 2SLS results suggest that overall—pooling all subjects (Row 1)—taking a first class in from a graduate student increases the probability of majoring in that subject. Across their college career, taking an initial course in a subject from a graduate student increased undergraduates' likelihood of choosing that subject as a major by 9.6 percentage points. This is almost a tripling of the baseline probability of 3.5% reported in Column 3 of Table 1. Exposure to graduate student teaching also increased the number of subsequent credits taken in the subject by about 2.3 or about one quarter of a standard deviation. These results are statistically different from zero when we allow for clustering at the student level, but allowing for clustering at the course-by-section level suggests the estimates may be much noisier. By contrast, it does not appear that undergraduates' major choices or subsequent course taking are influenced when their exposure to graduate instructors occurs during their first college term.

The additional rows of Table 8 report effect estimates for three subsamples: academic subjects (e.g., humanities, sciences, mathematics), and professional subjects (e.g., education, business, engineering). The academic and professional samples are mutually exclusive and exhaustive; the STEM sample draws from both.

The contrast between academic and professional subjects is notable. Undergraduates taught by graduate students are very likely to subsequently major in the subject: an estimated effect of 81 percentage points on major choice. The estimates for course taking behavior are complementary: 83 points more likely to take any additional courses and an average increase of 30 additional credits in the subject. These differences are robustly significant to both assumptions about error clustering.

Although these results are consistent with the large effects for professional students exposed to adjunct instructors in prior research (Bettinger & Long, 2010), we caution against making strong inferences regarding what appear to be large differences for professional subjects. There is reason to believe that both the timing of undergraduates' first course in a professional subject, as well as the use of graduate student instructors in these departments, differs in substantively important ways compared to academic departments. Turning to focus just on courses in academic disciplines, undergraduates taught by graduates students are about 2.4 percentage points more likely to major in the subject subsequently although estimates for course taking are fairly imprecise.

The results in Table 8 suggest that there are positive impacts on student engagement in subsequent courses. Other research has suggested that high quality instruction is directly related to subsequent engagement (e.g. Coates, 2005). There still remains the possibility that other graduate student characteristics or practices (e.g. ease of grading) could explain the differences in student outcomes noted in Table 8.12

#### 4.4. Effects of teaching on graduate student outcomes

While the instrumental variables estimates address issues of selection regarding the impact of graduate student instructors on undergraduates, the estimates in Table 8 are partly a function of the relative quality of graduate student instructors and full-time faculty in our Ohio public institution sample. Of particular concern is that, conditional on hiring graduate student instructors, departments likely do not randomly choose individuals to fill those teaching slots. Faculty may prefer that their most promising students take opportunities to teach as preparation for future careers or to preserve the quality of the department's instruction. Alternatively, faculty may prefer that their most promising students focus on research not teaching. To explore the selection of graduate students into teaching roles, we next estimate the relationship between the frequency with which a graduate student taught courses and the graduate student's own probability of future academic success.

In Column 1 of Table 9 we predict the probability of completing a doctoral degree within six years of starting graduate school. Each additional term in which a graduate student teaches increases the probability of degree completion by 2.1 percentage points in the full sample.<sup>13</sup> The point estimate is the same, 2.1 percentage points, when we limit the sample to graduate students in academic departments, and appears strongest in the humanities. These coefficients suggest that teaching one additional term in the humanities or math is associated with a 13% increase in the probability of degree completion within six years.

This predicted increase in degree completion rates could be an effect of the “treatment” of teaching undergraduates, but could just as likely be an artifact of how graduate student instructors are selected. The positive coefficients suggest positive selection: departments may prefer to selectively hire their relatively successful doctoral students to fill teaching slots. Institutions and departments may have heterogeneous preferences and our estimation includes campus-by-department fixed effects in part as a response; however, the results are not much changed without these fixed effects.

Degree completion is a broadly held goal for all graduate students. In addition, many graduate students are also working toward securing employment at a college or university after graduation. Again, opportunities to teach could have a causal effect on graduate students’ ability to obtain a post-secondary position. Or alternatively, a positive relationship could suggest that departments select graduate student instructors with preference for those who are on a trajectory toward college or university faculty.

In Column 2 of Table 9 we predict the probability of being employed by an Ohio college or university, other than one's own graduate institution, within six years of starting graduate school. Of note, unlike much of the analysis in this paper, “Ohio” here includes both public and private colleges and universities because ODJFS tracks all employment.<sup>14</sup> Across all departments, we estimate that the probability of subsequently working in higher education increases by 1.4 percentage points for each additional term a graduate student teaches while a candidate; a 9% increase over the baseline probability of about 0.15 (see Table 4 Column 1). The estimates for subject-specific sub-samples are mostly not statistically significant. As with the degree completion results, if the mechanism here is selection the selection appears positive, and potentially directed toward students on a future faculty trajectory.

We extend the analysis of future employment by comparing employment at a research university to employment at other colleges and universities.<sup>15</sup> Graduate student teaching is positively correlated with future employment at both research (Column 3) and non-research institutions (Column 4). In the full sample each additional term teaching raises the probability of employment 0.4 percentage points at a research university and 1.1 percentage points at other colleges and universities. If this difference is driven by selection it suggests that departments are more likely to choose graduate students headed toward jobs in colleges and universities where teaching skills are relatively more valuable. If the difference is causal, then the experience of teaching while a graduate student leads more individuals to jobs in college-level teaching.

The mechanisms behind the results in Table 9 could include some causal role for teaching opportunities. For example, graduate students who teach more frequently could develop skills preferred by non-research-focused institutions, or could miss out on opportunities to develop research skills needed for securing employment at research universities. These effects, if they exist at all, are relative since the results in columns 3 and 4 of Table 9 suggest that teaching predicts a greater chance of employment at both types of institutions.

Which interpretation—selection or causation—is correct? For our purposes, it may not matter. Both interpretations give the result that graduate student teaching generates positive results. If the results represent causal effects, then graduate student teaching is a positive experience and helps graduate students graduate in a timely way and gain employment in higher education, at least in their initial careers. Graduate student teaching leads to positive effects on graduate students and mostly positive or no effects on the undergraduate students they teach. If, on the other hand, the results represent selection of university-bound graduate students, then departments are successfully identifying the graduate students who will have positive effects on students, or at least no differential effects compared to full-time faculty.

## Acknowledgments

The authors thank the Ohio Board of Regents for their support during this research project. Funding was provided by the Bill & Melinda Gates Foundation OPP52525. The research reported here was supported by the Institute of Education Sciences, U.S. Department of Education, through Grant R305C110011 to Teachers College, Columbia University. Taylor was also supported in part by

the Institute of Education Sciences, U.S. Department of Education, through grant R305B090016 to Stanford University. All opinions and mistakes are our own.

## References

1. Bair, C. R., & Haworth, J. G. (2004). Doctoral student attrition and persistence: A meta-synthesis of research. In J. C. Smart (Ed.). *Higher education: Handbook of theory and research: XIX*. New York: Agathon Press.
2. Bettinger, E. P. (2004). Is the finish line in sight? financial aid's impact on retention and graduation. In Caroline M. Hoxby (Ed.), *College choices: The economics of where to go, when to go, and how to pay for it*. Chicago: University of Chicago Press and NBER.
3. Bettinger, E. P., & Long, B. T. (2010). Does cheaper mean better? the impact of using adjunct instructors on student outcomes. *Review of Economics and Statistics*, 92(3), 598–613.
4. Bettinger, E. P., & Long, B. T. (2004). Do college instructors matter: The effects of adjuncts and graduate assistants on students' interests and success. National Bureau of Economic Research working paper 10370.
5. Bettinger, Eric P., & Long, B. T. (2006). The increasing use of adjunct instructors at public institutions: Are we hurting students? In Ronald G. Ehrenberg (Ed.), *What's happening top public higher education* Westport, CT: American Council on Education / Praeger.
6. Borjas, G. J. (2000). Foreign-born teaching assistants and the academic performance of undergraduates. *American Economic Review*, 90(2), 355–359.
7. Bowen, W. G., & Rudenstine, N. L. (1992). *In pursuit of the PhD*. Princeton, N.J.: Princeton University Press.
8. Carrell, S. E., & West, J. E. (2010). Does professor quality matter? Evidence from random assignment of students to professors. *Journal of Political Economy*, 118(3), 409–432.
9. Coates, H. (2005). The value of student engagement for higher education quality assurance. *Quality in Higher Education*, 11(1), 25–36.
10. Ehrenberg, R. G., & Mavros, P. G. (1995). Do doctoral students' financial support patterns affect their times-to-degree and completion probabilities? *Journal of Human Resources*, 30(3), 581–609.
11. Fleisher, B., Hashimoto, M., & Weinberg, Bruce A. (2002). Foreign GTAs can be effective teachers of economics. *Journal of Economic Education*, 33(4), 299–325.
12. Hanushek, E. A., & Rivkin, S. G. (2010). Generalizations about using valueadded measures of teacher quality. *American Economic Review*, 100(2), 267–271.
13. Hoffmann, F., & Oreopoulos, P. (2009). A professor like me: The influence of instructor gender on college achievement. *Journal of Human Resources*, 44(2), 479–494.
14. Jacobs, L. C., & Friedman, C. B. (1988). Student achievement under foreign teaching associates compared with native teaching associates. *The Journal of Higher Education*, 59(5), 551–563.
15. Knotts, H. G., & Main, E. C. (1999). Teaching Ph.D. students to teach political science: The emory TATTO program. *PS: Political Science and Politics*, 32(3), 607–610.
16. Koehnecke, D. S. (1991). Boundaries of Graduate Assistants: Bouncing off Boundaries. In *Proceedings of the Annual Meeting of the Conference on College Composition and Communication*.
17. Mattson, K. (2000). The academic labor movement: Understanding its origin and current challenges. *Social Policy*, 30(4), 4–10.
18. McLeod, S. H., & Schwarzbach, F. S. (1993). What about the TAs? Making the Wyoming resolution a reality for graduate students. *WPA: Writing Program Administration*, 17(1–2), 83–87.
19. Meyers, S. A., & Loreto, R. P. (2000). Training in the teaching of psychology: What is done and examining the differences. *Teaching of Psychology*, 27(4), 258–261.
20. Mortenson, T. (2002). *Chance for college by age 19 by state in 2000*. Oskaloosa, IA: Postsecondary Education Opportunity.
21. Norris, T. (1991). Nonnative English-speaking teaching assistants and student performance. *Research in Higher Education*, 32(4), 433–448.
22. Loreto, R. P., & Altmaier, E. M. (1994). The relationship of prior training and previous teaching experience to self-efficacy among graduate teaching assistants. *Research in Higher Education*, 35(4), 481–497.
23. Reeve, J., & Jang, H. (2006). What teachers say and do to support students' autonomy during a learning activity. *Journal of educational psychology*, 98(1), 209–218.
24. Seagram, B. C., Gould, J., & Pyke, S. W. (1998). An investigation of gender and other variables on time to completion of doctoral degrees. *Research in Higher Education*, 39(3), 319–335.
25. Sharnoff, E. (1993). Neither fish nor fowl: Graduate students, unionization, and the academy. In *Proceedings of the Annual Meeting of the Modern Language Association*.
26. Skinner, E. A., & Belmont, M. J. (1993). Motivation in the classroom: Reciprocal effects of teacher behavior and student engagement across the school year. *Journal of educational psychology*, 85(4), 571–581.
27. Slevin, J. F. (1992). The next generation: Preparing graduate students for the professional responsibilities of college teachers. Washington, D.C.: Association of American Colleges.
28. Smith, K. S. (2001). Pivotal events in graduate teacher preparation for a faculty career. *Journal of Graduate Teaching Assistant Development*, 8(3), 97–105.
29. Vaughn, W. (1998). Apprentice or employee? Graduate students and their unions. *Academe*, 84(6), 43–49.
30. Watts, M., & Lynch, G. J. (1989). The principles course revisited. *American Economic Review*, 79(2), 236–241.
31. Wise, M. (1967). Who teaches the teachers? *Improving College Teaching* Washington, D.C.: American Council on Education.