Research in Economic Education

In this section, the *Journal of Economic Education* publishes original theoretical and empirical studies of economic education dealing with the analysis and evaluation of teaching methods, learning, attitudes and interests, materials, or processes.

PETER KENNEDY, Section Editor

Class Attendance and Exam Performance: A Randomized Experiment

Jennjou Chen and Tsui-Fang Lin

Abstract: The determination of college students' academic performance is an important issue in higher education. Whether students' attendance at lectures affects students' exam performance has received considerable attention. The authors conduct a randomized experiment to study the average attendance effect for students who choose to attend lectures, which is known in program evaluation literature as the average treatment effect on the treated. This effect has long been neglected by researchers when estimating the impact of lecture attendance on students' academic performance. Under the randomized experiment approach, the results suggest that class attendance has a positive and significant impact on college students' exam performance. On average, the effect of attending lectures corresponds to a 9.4 percent to 18.0 percent improvement in exam performance for those who choose to attend classes. In comparison, the improvement is only 5.1 percent, using the empirical method of existing studies, which measures the overall average attendance impact.

Keywords: attendance, experiment, treatment effect, undergraduate
JEL codes: A22, I21

***Jennjou Chen** is an assistant professor of economics at National Chengchi University, Taiwan (e-mail: jennjou@nccu.edu.tw). **Tsui-Fang Lin** is an assistant professor of public finance at National Taipei University, Taiwan.* *The authors thank Peter Kennedy, Kumiko Imai, Thomas Kane, Hung-Lin Tao, three anonymous referees, and seminar participants at the 2006 SOLE and 2007 ASSA meetings for their thoughtful comments. They gratefully acknowledge the financial support provided by Taiwan's National Science Council through grant NSC 95-2415-H-004-012.*

Determinants of college students' academic learning are an important topic in research on higher education. Most educators are concerned about the effect of teaching on a student's learning. Using exam performance as a proxy for learning results, many researchers have studied the determinants of college students' learning. Whether students' attendance at lectures and classroom discussions affects their exam performance has received considerable attention; many researchers have explored the effect of a student's class attendance on exam performance.

Researchers in the fields of education and psychology, by estimating correlations between exam performance and class attendance, have generally found that a student's class attendance has a positive effect on exam performance (Anikeeff 1954; Brocato 1989; Gunn 1993; Jones 1984; Rocca 2003; Van Blerkom 1992). Economists and other social scientists are also interested in class attendance effects. Most economists have used student semester-level data, and they have found results similar to those described in education and psychology literature: During a semester, the more lectures a student attends, the better overall grade he or she obtains (Schmidt 1983; Jones 1984; Park and Kerr 1990; Romer 1993; Durden and Ellis 1995; Devadoss and Foltz 1996; Dolton, Marcenaro, and Navarro 2003).

Some researchers have recently linked exam questions to students' attendance records and constructed a longitudinal data set to investigate class attendance effects (Marburger 2001, 2006; Rodgers 2001; Stanca 2006; Lin and Chen 2006). In such data sets, researchers repeatedly observed the same student's responses to different questions, as well as different students' responses to the same question. Hence, time invariant characteristics of both students and exam questions can be controlled in their statistical models. These rich data sets allow researchers to address some other interesting issues in addition to attendance effects.[1]

Attendance at a lecture can be viewed as a treatment to students and, thus, investigation of the attendance effect is indeed an estimation of treatment effect. In program evaluation literature, two kinds of treatment effects are frequently mentioned: average treatment effect and average treatment effect on the treated. *Average treatment effect* refers to the expected effect of attendance on academic performance of a randomly drawn student. *Average treatment effect on the treated* refers to the average effect of attendance on performance of those who actually participated in the classroom.

In reality, there might be two types of students—type A, who choose to attend lectures regularly, and type B, who are less conscientious about attending lectures. To simplify the story, η_A and η_B refer to the effects of attending lectures on type A and type B students, respectively. There is good reason to believe that η_A and η_B are not equal, and, indeed, values of η_A are higher than η_B. Type A students attend lectures more regularly, and this may imply that they obtain greater benefits from attendance, which makes their learning more effective.

In previous research, measures of the impact of attending lectures on students' performance have been estimated as weighted average of η_A and η_B. This is the so-called *average treatment effect* in program evaluation literature. It measures the average attendance effect for a randomly selected student, who has received the treatment (i.e., has attended a lecture). Because the randomly selected student could be a type A student or a type B student, the effect is a weighted average of η_A and η_B. This weighted-average attendance effect can also be a measure of potential benefits of enforcing a mandatory attendance policy, because under the policy, both type A and type B students would be required to attend lectures. Thus, most researchers focus on the estimation of the average treatment effect because the issue of whether to make attendance compulsory has received great attention in higher education research.

However, estimation of the average treatment effect on the treated, or the magnitude of η_A, is not only interesting but also as important as the average attendance effect for the following reasons. First, the average attendance effect on the attendees can be viewed as the actual output produced by the course or by the instructor, which deserves special attention. This idea is similar to the example of job-training programs; what researchers and policymakers really want to know is the impact of job-training in terms of the outcomes for program participants and not for an average participant.

Second, the comparison between the average attendance effect (i.e., weighted average of η_A and η_B) and the effect on the treated (i.e., η_A) is interesting and insightful. If a student decides to attend lectures regularly, it implies that the expected benefit of going to classes is greater than the opportunity costs. Therefore, given similar opportunity costs of attending lectures for type A and type B students, we would expect that, on average, the average weighted attendance effect is smaller than the effect on the treated (those who actually attend classes). That is, the weighted average of η_A and η_B is smaller than η_A; and this implies that η_B is smaller than η_A. This hypothesis can be tested by estimating and comparing the two attendance effects. To complement the current attendance effect literature, our focus in this article is to investigate the average treatment effect on the treated.

Our main purpose in this article is to construct a randomized experimental approach to estimate the average attendance effect on the treated. We discuss details of the randomized experiment, followed by the data used for this study, statistical models, and estimation results.

THE RANDOMIZED EXPERIMENT

Our main goal is to construct a randomized experiment to estimate the average attendance effect on the attendees. One difficulty in estimating the average treatment effect on the treated arises from the problem of finding the desired counterfactuals. In this case, we needed to estimate what would have been the grades of students who attended the lectures if they had not attended the classes. One way to circumvent the problem of finding the desired counterfactuals is to run a randomized experiment. The pros and cons of social experiments are detailed in Burtless (1995) and Heckman and Smith (1995).[2]

The key idea for the randomized experiment is that the instructor taught the same course in two sections in the sample semester. Under our experiment design, some course material was randomly covered explicitly in one section but not in the other section and vice versa for some other material. This allowed us to observe if those who chose to attend were affected, by comparing the scores of the students in the one section with the scores of the students in the other section on certain questions. By restricting our sample to students who had chosen to attend and by using a randomization procedure, we could estimate the treatment effect on the treated.

Because the attendees were a mixture of type A and type B students, the treatment effect on the treated that we estimated mixed the impact on type A and type B students. Thus, with the purpose of examining the type A versus type B difference, we further restricted our sample to different types of attendee students by their attendance frequency. We discuss the details of various types of attendees in the section on estimation results.

Note that the treatment in this experiment is receiving the lecture. This attendance effect can also be viewed as the effect of omitted lecture material on learning, if the treatment is defined the other way around. For example, instead of defining *receive the lecture* as the treatment, we might define *forcing students to skip some topics* as the treatment.

In discussing the theoretical basis of our experiment, the following notation is similar to that used in Heckman and Smith (1995):

Y_1: grade outcomes associated with attending the lecture.
Y_0: grade outcomes associated with not attending the lecture.
$d = 1$, attending the lecture; $d = 0$, not attending the lecture.

We were interested in the mean impact of attending lectures on exam performance for students who chose to attend classes. The average attendance effect on the attendees is shown as

$$E(Y_1|d = 1) - E(Y_0|d = 1). \tag{1}$$

To estimate the effect, we need to know what the grades would be if the students did not attend the class. This implies that we needed an estimate for $E(Y_0|d = 1)$ because it is unobserved by researchers. In general, $E(Y_0|d = 0)$ cannot be used as a proxy for $E(Y_0|d = 1)$ because students who chose not to attend lectures might be different in many ways from those who chose to attend classes, such as unobserved individual intelligence and motivation. As a result, the process of selecting to attend or not to attend classes might become an issue; and it would bias our results if we use $E(Y_0|d = 0)$ to replace $E(Y_0|d = 1)$.

Our main focus was to generate an experimental group of students who would have participated but were randomly denied access to the treatment. By doing so, we could use this randomly selected group as our control group and obtain their responses as the desired counterfactuals, $E(Y_0|d = 1)$. Ideally, the instructor could randomly select some students and ask them to leave the classroom at the beginning of each lecture. However, this approach comes with at least two major problems. First, the instructor would have some difficulty in convincing the university officials to allow the instructor to do such an experiment because asking students to skip lectures is something a university usually does not want to do. Second, perhaps a more problematic issue would be that students' decisions to attend (or not to attend) lectures might be altered once students learned that there was a possibility of their being denied access to classes.

Because of these two potential problems, we used the following different approach to estimate the counterfactuals. The instructor taught the same course in two sections in the sample semester. At each class meeting, the same PowerPoint presentation was used in both sections, and the lecture slides were posted on the course Web site after each class meeting. During the sample semester, the instructor randomly selected the dates, sections, and some materials and topics that would be covered in only one section but not in the other section. The lecture slides that were randomly skipped in one of the sections had to be taught in the other section. Consequently, we could observe and compare students' performance from receiving and not receiving the treatment.

In addition, students were told to be responsible for materials and topics shown in the slides, including the ones skipped by the instructor. This implied that materials and topics not covered by the instructor might appear in the exams and that students would need to prepare and study those materials by themselves to be able to answer the corresponding exam questions. In this study, about 8 percent of the exam questions were not covered by the instructor, and yet they appeared in the exams.

Let $d^* = 1$ denote the students who attended lectures and $d^* = 0$ denote everyone else. Also, let $r = 1$ denote the group of students who were randomly assigned to the treatment group for particular exam questions (i.e., materials and topics corresponding to specific exam questions were covered by the instructor) and $r = 0$ denote the group of students who were denied access to treatment for the same questions (i.e., materials and topics corresponding to exam questions were randomly skipped by the instructor).

By introducing variables d^* and r, we can rewrite equation (1) as

$$\begin{aligned} E(Y_1|d = 1) - E(Y_0|d = 1) = E(Y_1|d^* = 1, r = 1) \\ - E(Y_0|d^* = 1, r = 1), \end{aligned} \tag{1'}$$

where $d = 1$ is replaced by $d^* = 1$ and $r = 1$.

We could reasonably expect that

$$E(Y_0|d^* = 1, r = 1) = E(Y_0|d^* = 1, r = 0). \tag{2}$$

$E(Y_0|d^* = 1, r = 0)$ is the expected grade for students who chose to attend lectures but did not actually receive certain treatments because materials and topics corresponding to certain questions were randomly skipped. The original problem was that we could not observe $E(Y_0|d^* = 1, r = 1)$ in equation (1′). $E(Y_0|d^* = 1, r = 1)$ is the average grade that would be obtained if the students had not attended the lecture. This partial observation issue is a common problem in estimating the average treatment effect on the treated. By running the randomized experiment, we could observe $E(Y_0|d^* = 1, r = 0)$ and use it as a replacement for $E(Y_0|d^* = 1, r = 1)$. Hence, by use of equation (2), the average attendance effect on the attendees could be shown as

$$\begin{aligned} E(Y_1 - Y_0|d = 1) &= E(Y_1 - Y_0|d^* = 1, r = 1) \\ &= E(Y_1|d^* = 1, r = 1) - E(Y_0|d^* = 1, r = 1) \\ &= E(Y_1|d^* = 1, r = 1) - E(Y_0|d^* = 1, r = 0). \end{aligned} \tag{3}$$

Thus, randomization (i.e., $r = 1$ or $r = 0$) served as an instrumental variable by creating variations among students who chose to attend lectures, because some of them received the treatment (i.e., $r = 1$), whereas some of them did not (i.e., $r = 0$). In so doing, we could estimate the counterfactuals for the attendees and obtain the average attendance effect on them accurately.

DATA

We conducted a survey of 114 students who took the public finance course at a private university in Taiwan in the spring semester of 2005. All students who major in industrial economics are required to take this course in their third year of study. Students were in two separate sections: 67 students in the first section, and 47 students in the second. Students freely chose to register in either section. Both sections were taught by the same instructor but had different time schedules. One class met at 3 p.m., and the other met at 5 p.m. Also, the same PowerPoint presentation was used in both sections, and lecture slides were posted on the course Web site after each class meeting. There were 12 two-hour class meetings in addition to two exams and one project presentation during the sample period. The same exam questions were taken by all students in both sections at the same time. Attendance was recorded at each class meeting during the sample semester.

In this article, the dependent variable was a binary variable indicating students' exam performance. Fifty multiple-choice questions were asked in the midterm exam, and 57 multiple-choice questions were asked in the final exam. There were 12,028 observations, which came from 114 students and their responses to the 107 exam questions.[3] We assigned 1 to the binary variable if students answered the exam question correctly; otherwise the binary variable was 0.

There were two main independent variables: *actual attendance* (i.e., d^* in equation [3]) and *experimental attendance* (i.e., r in equation [3]). *Actual attendance* was used to obtain the average attendance effect, whereas *experimental attendance* was used to estimate the average attendance effect on attendees. The binary variable, *actual attendance*, was coded as 1 if a student had attended the lecture that covered the material relevant to the corresponding exam question (i.e., $d^* = 1$), as discussed in the random experiment section. *Actual attendance* was coded as 0 if a student missed the class that day (i.e., $d^* = 0$).

Among the students who attended lectures, we created a binary variable, *experimental attendance*. *Experimental attendance* was coded as 1 if a student had attended the lecture ($d = 1$), and the instructor had taught the material in that lecture ($r = 1$). *Experimental attendance* was coded as 0 if a student had attended the lecture ($d = 1$), but the instructor had randomly chosen not to cover materials corresponding to specific exam questions in that lecture ($r = 0$).

The average actual attendance rate was 91 percent, which was higher than that in some previous studies (Romer 1993; Marburger 2001). Note that the sample course, public finance, is a required course for students in their junior year. Students are more likely to attend lectures when they are in their junior and senior years, as pointed out by Rocca (2003). Therefore, a 91 percent class attendance rate seems reasonable. If we further restricted our sample to students who chose to attend lectures, we found that the average experimental attendance was about 92 percent, which also implied that 8 percent of the course materials were randomly skipped.

Table 1 shows the percentage of correct answers to exam questions by students' attendance records and types of exam questions. The percentage of correct answers of questions was computed for two groups: attendees and nonattendees. The first column presents the percentage of correct answers among attendees, and the second column presents the percentage of correct answers by nonattendees. In addition, exam questions were divided into two types: type X and type Y. Type X exam questions correspond to materials covered in the lectures, and type Y questions correspond to unlectured materials for some attendees.

The percentage of correct answers of type X questions for attendees was 64.6 percent, and that for nonattendees was 63.2 percent. The difference was not significant; the average scores on type X questions were very similar between attendees and nonattendees. As for type Y questions, if they were covered in the lectures, the percentage of correct answers of these questions for attendees was 62.7 percent. However, if type Y questions were randomly skipped for comparison purpose, the percentage of correct answers was 51.1 percent. That is, when the attendees were randomly assigned to the control group and did not receive the lecture treatment, their average score was much lower. Thus, without controlling students' individual effects and exam question effects, the attendance effect on the attendees was about 11.6 percent. Note that nonattendees achieved low scores in type Y questions, and the percentage of correct answers in their case was only 44.9 percent.

STATISTICAL MODELS

This study used a micro-level data set to explore the average attendance effect for students who chose to attend lectures. We used the following probit model to describe the relation between a student's exam performance and various learning input variables:

$$y_{ij}^* = \eta r_{ij} + \alpha_i + \gamma_j + \varepsilon_{ij}, \qquad \text{and}$$
$$y_{ij} = \begin{cases} 1 \text{ } if \text{ } y_{ij}^* \geq 0 \\ 0 \text{ } if \text{ } y_{ij}^* < 0 \end{cases}, \qquad i = 1, 2, \ldots, I, \ j = 1, 2, 3, \ldots, J. \tag{4}$$

I is the total number of students, and J is the total number of exam questions. The variable y_{ij} corresponds to student i's observed exam performance on question j. The variable y_{ij}^* is the unobserved propensity for exam performance. The variable r_{ij} is the *experimental attendance*, which equals 1 if student i attended the lecture when question j was covered; r_{ij} equals 0 if student i attended the lecture when question j was not covered. The variable η is the attendance effect. The variable α_i represents student i's time-invariant individual effect, γ_j represents question j's effect, and ε_{ij} is a random disturbance term.

We restricted our sample to attendees, and the parameter of interest in this study was η, the average attendance effect on the attendees. We employed both probit and probit-with-individual-dummies models. We call probit-with-individual-dummies

TABLE 1. Sample Means of Exam Performance (by Attendance Record and Type of Exam Questions)

		Attendees ($d^* = 1$)			Nonattendees ($d^* = 0$)		
Exam question		Mean	SD	Sample size	Mean	SD	Sample size
Type X	($r = 1$)	0.6459	0.0050	9224	0.6321	0.0163	878
Type Y	For attendees for whom this material was covered in lecture ($r = 1$)	0.6270	0.0162	893	0.4492	0.0347	207
	For attendees for whom this material was not covered in lecture ($r = 0$)	0.5108	0.0174	826			

Note. Type X questions are the ones based on material covered in lecture for all attendees. Type Y questions are the rest of the questions (randomly skipped in either one of the sections).

models *probit fixed-effects models* hereafter. By the definition of a randomized experiment, the treatment was randomly assigned within the estimation sample and was not correlated with x_{ij}, α_i, γ_j, and ε_{ij}. This implied that probit estimation of the attendance effect would yield consistent results even though individual effects were not controlled in the probit model. Thus, we expected that both probit and probit fixed-effects models were consistent and produced similar estimates.

ESTIMATION RESULTS

Table 2 presents the estimation results for the average attendance effect (i.e., weighted average of η_A and η_B), and it replicates previous observational studies such as Marburger (2001, 2006) and Stanca (2006). To be consistent with prior research, both attendees and nonattendees were included in the analysis sample. In addition, we removed observations corresponding to unlectured topics from the sample. The number of observations was 11,097 in this case.

The first column shows estimation results of the probit model, and the second column shows estimation results of the probit fixed-effects model. In both models, the dependent variable was a binary variable, indicating whether or not the students answered the exam questions correctly. Independent variables in the probit model included *actual attendance* and exam question dummies. In addition to these independent variables, individual time-invariant dummies were also used in the probit fixed-effects model. We reported both the coefficients and marginal effects for the *actual attendance* variable. Marginal effects were evaluated at the sample means of the independent variables. In this article, we discuss the marginal effects results because it is more intuitive to interpret them.

From the data in Table 2, we found that attendance produced a significant and positive impact on students' exam performances. The marginal effect of *actual attendance* in the probit model was 8.6 percent, and it declined to 5.1 percent in the probit fixed-effects model. Thus, after accounting for individual heterogeneity, we obtained a smaller attendance effect in the probit fixed-effects model than in the probit model. This result is similar to that in previous research. For instance, Stanca (2006) found that least squares overestimate the impact of attendance on exam performance. In Stanca (2006), *average attendance* effects ranged from 7.3 percent to 9 percent, and the size of the attendance effect declined to 4 percent in the fixed-effects model. Also, Marburger (2001, 2006) found that absenteeism increased the probability of answering the exam question incorrectly. Absenteeism effects ranged from 7.5 percent to 14.6 percent in Marburger (2001) and from 9 percent to 14 percent in Marburger (2006).

Table 3 presents the estimation results of the *average attendance* effect on the attendees under the randomized experiment setting, which is the main focus of this study. We restricted our samples to observations with *actual attendance* equal to 1. There were 114 students, and the sample size was 10,919. In this case, we obtained the *average attendance* effect on the attendees. In addition, to estimate the *attendance effects* for type A students, who chose to attend lectures regularly

TABLE 2. Estimation Results for the Average Attendance Effect

Variable	Probit (with exam question dummies only)		Probit fixed effects (with both sets of exam and individual dummies)	
	Coefficients	Marginal effects	Coefficients	Marginal effects
Dependent variable	Answer the question correctly (yes = 1, no = 0)			
Independent variable				
Actual attendance	0.2386* (0.0478)	0.0862* (0.0173)	0.1422* (0.0543)	0.0509* (0.0194)
Sample size	11,097		11,097	

*Significant at .05 Type I error level. White (1980) robust standard errors are in parentheses.

TABLE 3. Estimation Results for the Attendance Effects on Attendees

Variable	Probit (with exam question dummies)				Probit fixed effects (with both sets of exam question and individual dummies)			
	Coefficients	Marginal effects	Sample size	Number of students	Coefficients	Marginal effects	Sample size	Number of students
Dependent variable	Answer the question correctly (yes = 1, no = 0)							
Independent variables								
Experimental attendance								
Sets of analysis samples								
Attendees who never missed a lecture	0.4689*	0.1800*	5,026	48	0.4717*	0.1802*	5,026	48
	(0.0951)	(0.0376)			(0.0972)	(0.0385)		
Attendees who missed fewer than two lectures	0.3832*	0.1466*	7,641	76	0.3825*	0.1455*	7,641	76
	(0.0774)	(0.0305)			(0.0789)	(0.0311)		
Attendees who missed fewer than three lectures	0.3065*	0.1165*	9,455	99	0.3067*	0.1169*	9,455	99
	(0.0699)	(0.0274)			(0.0712)	(0.0278)		
All attendees	0.2474*	0.0935*	10,919	114	0.2509*	0.0943*	10,919	114
	(0.0656)	(0.0255)			(0.0669)	(0.0259)		

*Significant at .05 Type I error level. White (1980) robust standard errors are in parentheses.

(i.e., η_A), we used another three sets of analysis samples. The definitions of these samples are

1. attendees who never missed a lecture,
2. attendees who never missed a lecture or missed only one lecture, and
3. attendees who never missed a lecture or missed only one lecture or missed only two lectures.

Among the 114 students, there were 48 students who never missed a lecture; 76 missed fewer than two lectures. In addition, there were 99 students who missed fewer than three lectures. In these models, the dependent variable was a binary variable, indicating whether or not the students answered the exam questions correctly. Independent variables in the probit model included *experimental attendance* and exam question dummies. In addition to these independent variables, individual time-invariant dummies were also used in the probit fixed-effects model.

There are two important findings in Table 3. First, the marginal effects of *experimental attendance* in both probit and probit fixed-effects models were nearly identical in all sets of analysis samples. For instance, marginal attendance effects for students who never missed any lecture was 18.0 percent in both models. The finding held for the other three sets of analysis samples. As emphasized in the random experiment and the statistical model sections, randomization served as an exogenous instrumental variable. Thus, whether time-invariant individual characteristics were controlled in the probit model or not, both probit and probit fixed-effects models should yield consistent estimates.

Second, we found that the more frequently a student attended lectures, the greater the benefits obtained from attending. Among students who attended lectures regularly, attending lectures yielded a positive, significant, and larger impact on performance of those who attended more often. For example, the average attendance effect for students who never missed any lecture was 18.0 percent; it was 14.7 percent for attendees who missed fewer than two lectures. The attendance effect was even lower at 11.7 percent for attendees who missed fewer than three lectures. Finally, for all attendees, the attendance effect declined to 9.4 percent. This interesting finding was intuitive and consistent with our prediction because students who had decided to attend lectures regularly may have had a higher return from attending classes than those who were less likely to attend.

Comparing the estimation results in Table 2 and Table 3, we also found that the weighted average attendance effect, 5.0 percent, was much lower than the average attendance effect on the attendees, which ranged from 9.4 percent to 18.0 percent. The results suggested that the mean attendance effect for students who chose to attend classes regularly (i.e., η_A) was greater than the mean attendance effect when students were randomly selected to attend lectures (i.e., weighted average of η_A and η_B). This finding was also consistent with our intuition because students who decided to attend lectures might get a higher return from attending classes than those who were randomly selected to attend.

Some might be concerned about issues regarding random assignment of treatment in this study. For instance, our estimation results might suffer from some biases because students might expect that materials not covered by the instructor were less likely to appear in the exams, even though they were told to be responsible for the skipped materials. To examine whether students' perception might have been an issue and biased our results, we further divided our sample into two sets: the midterm exam sample and the final exam sample.

If students felt that materials not covered by the instructor were less likely to appear in the exam, we would expect to find different attendance impacts in midterm and final samples. For example, in the midterm, if students assumed that materials not covered by the instructor would not appear in the exam and then they realized, after taking the midterm examination, that they were wrong, they would pay the same attention to the skipped materials as the remaining topics, when preparing for the final exam. If this was the case, we would expect to see a smaller attendance effect in the final exam samples. Otherwise, we would expect to find similar results in midterm and final exam samples.

Table 4 presents estimation results for the average attendance effect on the attendees in midterm and final exams' samples. In these two models, the same dependent and independent variables were used. It is important to note that we found fairly similar results in both cases for all attendees. For the midterm exam, the average attendance effect on the attendees was 8.6 percent; for the final exam, the average attendance effect on the attendees was 9.4 percent. Because these two estimated statistics were not significantly different from each other, we did not need to worry about the perception issue here. Moreover, this also assured us of the robustness of our estimation results.

CONCLUSION

This study contributes to literature on class attendance effects by using a randomized experiment to estimate the average attendance effect on the attendees. We conducted a classroom experiment to control for students' endogenous class attendance choices and explored the impact of class attendance on exam performance. Our data set provided us with a great opportunity not only to replicate previous observational studies in estimation of attendance effect but also to clearly identify the causal link between attendance and exam performance in an experimental setting. Under our randomized experiment, the mean outcomes of the experimental treatment and the control groups provided estimates of the average attendance effect on the attendees.

Our estimation results show that under the randomized experiment, simply estimating the probit model, without controlling students' heterogeneity, still yields consistent estimates. In addition, both probit and probit fixed-effects models produce similar estimates of attendance effects. On average, the effect of attending lectures corresponds to a 9.4 percent to 18.0 percent improvement in exam performance for students who choose to attend lectures. Moreover, the more frequently a student attends lectures, the greater the benefits he or she obtains from attending. Last, the average attendance effect on the attendees is much larger than the average attendance effect. We found that the improvement is only 5.1 percent using the empirical methodology of existing studies, which measure the overall average attendance impact.

TABLE 4. Estimation Results for the Attendance Effect on Attendees by Exam

	Midterm exam				Final exam			
	Probit (with exam question dummies only)		Probit fixed effects (with both sets of exam and individual dummies)		Probit (with exam question dummies only)		Probit fixed effects (with both sets of exam and individual dummies)	
Variable	Coefficients	Marginal effects	Coefficients	Marginal effects	Coefficients	Marginal effects	Coefficients	Marginal effects
Dependent variable	Answer the question correctly (yes = 1, no = 0)							
Independent variable								
Experimental attendance	0.2279* (0.1157)	0.0846* (0.0444)	0.2401* (0.1173)	0.0884* (0.0448)	0.2564* (0.0796)	0.0982* (0.0312)	0.2389* (0.0830)	0.0906* (0.0323)
Sample size	5,247				5,672			

*Significant at .05 Type I error level. White (1980) robust standard errors are in parentheses.

NOTES

1. For example, Marburger (2006) studied the mandatory-attendance policy effect and found that an enforced mandatory-attendance policy significantly reduced absenteeism and improved exam performance. Lin and Chen (2006) incorporated the spillover effects from absenteeism in estimation of college students' academic performance and found a significant and positive effect of past cumulative attendance on exam performance.
2. The foremost advantage of controlled experimentation is that the random assignment provides clear causal links between treatment and outcome. In nonexperimental data, it is usually not easy to extract the causality between treatment and outcome. Random assignment also obliterates systematic correlation between treatment status and participants' observed or unobserved characteristics. In addition, controlled experimentation is simple to understand for social scientists and policymakers. Some disadvantages of controlled experimentation include high costs, ethical issues of experimentation with human beings, limited duration, attrition and interview nonresponse, partial equilibrium results, and program entry effects.
3. Two students missed the final exam (57×2), and some questions were not answered by some students (56). So $114 \times (50 + 57) - (57 \times 2) - 56 = 12{,}028$.

REFERENCES

Anikeeff, M. 1954. The relationship between class absences and college grades. *Journal of Educational Psychology* 45 (4): 244–49.

Brocato, J. 1989. How much does coming to class matter? Some evidence of class attendance and grade performance. *Educational Research Quarterly* 13 (3): 2–6.

Burtless, G. 1995. The case for randomized field trials in economic and policy research. *Journal of Economic Perspectives* 9 (2): 63–84.

Devadoss, S., and J. Foltz. 1996. Evaluation of factors influencing student class attendance and performance. *American Journal of Agriculture Economics* 78 (3): 499–507.

Dolton, P., D. Marcenaro, and L. Navarro. 2003. The effective use of student time: A stochastic frontier production function case study. *Economics of Education Review* 22 (6): 547–60.

Durden, C., and V. Ellis. 1995. The effects of attendance on student learning in principles of economics. *American Economic Review* 85 (2): 343–46.

Gunn, P. 1993. A correlation between attendance and grades in a first-year psychology course. *Canadian Psychology* 34 (2): 201–02.

Heckman, J., and J. Smith. 1995. Assessing the case for social experiments. *Journal of Economic Perspectives* 9 (2): 85–110.

Jones, H. 1984. Interaction of absences and grades in a college course. *Journal of Psychology* 116 (1): 133–36.

Lin, T. F., and J. Chen. 2006. Cumulative class attendance and exam performance. *Applied Economics Letters* 13 (14): 937–42.

Marburger, R. 2001. Absenteeism and undergraduate exam performance. *Journal of Economic Education* 32 (2): 99–110.

———. 2006. Does mandatory attendance improve student performance? *Journal of Economic Education* 37 (2): 148–55.

Park, H., and P. Kerr. 1990. Determinants of academic performance: A multinomial logit approach. *Journal of Economic Education* 21 (2): 101–11.

Rocca, K. 2003. Student attendance: A comprehensive literature review. *Journal on Excellence in College Teaching* 14 (1): 85–107.

Rodgers, R. 2001. A panel-data study of the effect of student attendance on university performance. *Australian Journal of Education* 45 (3): 284–95.

Romer, D. 1993. Do students go to class? Should they? *Journal of Economic Perspectives* 7 (3): 167–74.

Schmidt, R. 1983. Who maximizes what? A study in student time allocation. *American Economic Review Papers and Proceedings* 73 (2): 23–28.

Stanca, L. 2006. The effects of attendance on academic performance: Panel data evidence for introductory microeconomics. *Journal of Economic Education* 37 (4): 251–66.

Van Blerkom, L. 1992. Class attendance in an undergraduate course. *Journal of Psychology* 126 (5): 487–94.

White, H. 1980. A heteroskedasticity-consistent covariance matrix estimator and a direct test for heteroskedasticity. *Econometrica* 48 (4): 817–38.