

This PDF is a selection from a published volume from the National Bureau of Economic Research

Volume Title: College Choices: The Economics of Where to Go, When to Go, and How to Pay For It

Volume Author/Editor: Caroline M. Hoxby, editor

Volume Publisher: University of Chicago Press

Volume ISBN: 0-226-35535-7

Volume URL: <http://www.nber.org/books/hoxb04-1>

Conference Date: August 13-15, 2002

Publication Date: September 2004

Title: How Financial Aid Affects Persistence

Author: Eric Bettinger

URL: <http://www.nber.org/chapters/c10101>

How Financial Aid Affects Persistence

Eric Bettinger

The Pell Grant program is the largest means-tested financial assistance available to postsecondary students across the United States. Students from all types of degree granting postsecondary institutions can apply for Pell Grants. In 2000–2001, the federal government awarded almost \$8 billion in Pell Grants among more than 3.8 million students, roughly one-third of all college students (College Board 2001). President Bush’s 2003 budget allocates over \$10.9 billion dollars for an estimated 4.5 million Pell Grant recipients (U.S. Department of Education 2002), potentially representing a 32.4 percent increase in the number of students receiving Pell Grants since the 1990–1991 school year. Yet despite this continued expansion of the Pell Grant, researchers have only limited evidence on the causal effects of these grants.

Most Pell Grant-related research focuses on the effects of Pell Grants on enrollment decisions, specifically focusing on initial enrollment and choice amongst colleges (see Kane 1999; Ehrenberg and Sherman 1984; Leslie and Brinkman 1987; Seftor and Turner 2002). However, there is surprisingly little research measuring the causal effect of Pell Grants on student outcomes in college (e.g., persistence, graduation). Regardless of whether Pell Grants affect initial enrollment patterns, Pell Grants may independently affect student outcomes.

Moreover, studying the effects of need-based aid on student outcomes may also be important since Pell Grant-eligible students are more likely to

Eric Bettinger is assistant professor of economics at Case Western Reserve University, and a faculty research associate of the National Bureau of Economic Research.

Special thanks go to Rob Sheehan and Andy Lechler for helpful comments and for providing the data. I am also grateful for comments from David Cooper, Jim Rebitzer, Bridget Long, Jon Guryan, and especially Caroline Hoxby. All errors are my own.

be on the margin of “stopping out.”¹ At Ohio four-year colleges in 1999–2000, 18 percent of full-time freshmen who were *not* eligible for Pell Grants withdrew from college by the next year, while 28 percent of students who were eligible for a Pell Grant did not enroll the following year. It is an open question whether Pell Grants and other need-based aid programs affect these margins. This paper attempts to resolve these questions. Using unique student data from Ohio, this paper measures the causal relationship between need-based aid and student retention.

A study of the effects of need-based aid on student retention may be of interest to both policymakers and educational researchers. As Sarah Turner argues in her chapter in this volume, policymakers have often paid more attention to improving student access to college rather than improving student retention once in college. However, as college enrollment rates continue to increase while completion rates do not, both policymakers and researchers have begun focusing more heavily on indicators of student retention (e.g., DesJardins, Ahlburg, and McCall 1999; St. John, Hu, and Tuttle 2000; St. John, Hu, and Weber 2000), and the relationship between financial need and persistence is central to many of these studies (e.g., DesJardins, Ahlburg, and McCall 2002; St. John, Musoba, and Simmons 2003).

There are a number of reasons why more research has not investigated the effects of Pell Grants on student collegiate outcomes. One reason is that researchers have difficulty distinguishing between the effects of family characteristics and the effects of Pell Grants. Pell Grants are a means-tested program. Comparisons between Pell Grant recipients and non-Pell Grant recipients (e.g., Wei and Carroll 2002) may be difficult to interpret since Pell Grant recipients are poorer and may be more likely to drop out, even in the absence of need-based aid. To correct for such bias, researchers must sufficiently control for family characteristics.

Additionally, identifying the effects of Pell Grants is difficult since much of the variation in the size of students' Pell Grants is correlated with students' college enrollment decisions. For example, college choice and the size of a student's Pell Grant are directly connected. Students who attend more expensive (and often higher-quality) schools are eligible for larger Pell Grants than students at other colleges or universities. Pell Grants are also more generous for full-time rather than part-time students. Even in the absence of Pell Grants, students who benefit most from college are more likely than other students to attend more expensive schools and to attend full-time. However, since Pell Grant awards are systematically higher for these same students, it may be difficult to identify the effects of the Pell

1. “Stopping out” refers to students who withdraw from school after their first year. These students are not “dropouts” because many of these students do not leave school permanently, and their undergraduate credit hours do not “expire.” I use these terms interchangeably throughout the paper.

Grant separate from college enrollment effects. To avoid this bias, researchers must exploit variation in Pell Grants that is independent of college choice (e.g., discontinuities in the Pell Grant formula).

A final reason that researchers have been unable to identify the effects of Pell Grants on outcomes is the absence of accurate data, in particular, the absence of accurate persistence and detailed financial data. Some researchers have measured persistence at a particular university; however, in these data, researchers cannot distinguish between a student who transferred to another school and one who withdraws from college. Other survey-based data rely on students' self-reports of their college experience. These data may not be as reliable as administrative data since students may not recall or do not wish to report small periods of time when they withdrew from college.

Financial data are equally as difficult to obtain. Most of the students who receive Pell Grants do not attend elite, expensive institutions, nor do they have substantial family support. Pell Grant recipients (and much of the variation in their awards) typically come from less-expensive colleges and their family contributions are much smaller. Moreover, the variation in Pell Grants is typically small. Exact financial data are necessary both to identify the small variations in Pell Grants and to employ creative identification strategies. Survey data (e.g., High School and Beyond, National Educational Longitudinal Survey) do not offer the level of detail necessary to identify accurately the level of students' Pell Grants.

To examine the effects of the Pell Grant, this paper presents evidence from data gathered by the Ohio Board of Regents (OBR). These data do not have the shortcomings of other data sets and offer a level of detail on both persistence and financial variables that is not available in other data. Since 1998, OBR has collected comprehensive data on college enrollment in Ohio's public two- and four-year colleges. As a result, the OBR data tracks students within and across schools. With the data, researchers can distinguish between students who withdraw from school and students who transfer to other Ohio schools. Moreover, through collaborative agreements, OBR has expanded the data to include students' American College Test (ACT) scores and data from the Free Application for Federal Student Aid (FAFSA). The FAFSA data are the exact data used by institutions to determine the amount of students' Pell Grant eligibility.

The level of detail in the financial data also facilitates the use of statistical tools that are impractical using other data. In particular, the level of detail allows researchers to identify small discontinuities in the Pell Grant formula. These discontinuities may be exploited to identify the causal effects of the voucher. While this paper may not completely resolve biases from college choice and enrollment or family background, the discontinuity analysis may be the best available method for dealing with such biases.

The paper presents evidence on the effects of Pell Grants using both

panel and cross-sectional variation.² The panel specifications suggest that need-based financial aid reduces students' stop-out behavior. In identifying this effect, the paper shows that students who receive Pell Grants after their first year are a unique subset of students who applied for financial aid in their first year. Failure to control for this selection may confound causal estimates of the Pell Grant program.

The paper also presents evidence relying on cross-sectional variation. The paper estimates the effects of Pell Grants close to existing discontinuities in family size. The results based on discontinuity approaches suggest that Pell Grants increase persistence; however, the results are not robust to alternative specifications.

Section 5.1 of this paper presents a simple economic model of student persistence under uncertainty. Section 5.2 of the paper explains the OBR data in greater detail. Section 5.3 of the paper presents the empirical strategies and results. Section 5.4 discusses policy implications of the results and concludes.

5.1 Economic Model

Economists often model educational attainment as investment in human capital. Even basic economics classes teach that students will choose an education level that maximizes the expected present discounted value (PDV) of future wage payments less the expected PDV of educational costs. There have been a number of permutations to this model—factoring in scholarship aid, allowing the returns to education to vary, and showing how predicted education levels vary with expectations (Manski 1993). This paper investigates the relationship between financial aid and outcomes. Rather than use a traditional human capital model, the paper models students' dropout behavior using a multistage investment model.

Multistage investment models are particularly useful in cases where the agent must reevaluate the project after an initial period of time. For example, Myers and Majd (1990) investigate optimal abandonment rules for firms. Dixit and Pindyck (1994) review other examples of multistage investments.

The phenomenon of interest—students' stop-out behavior—is similar to these multistage investments. In the initial period, students must decide whether to attend the first year of college. After completion of the first year, students must then reevaluate whether to complete the next year. About 20 percent of first-time freshmen withdraw from four-year colleges after the first year.

To formalize the model, let person i 's wage at time t (w_{it}) be modeled as a

2. It is important to distinguish between the effects of Pell Grants in general and the effects of Pell Grants conditional on initial enrollment. This paper focuses on estimating the effect of persistence conditional on a student having enrolled. I discuss in the following the relationship between the conditional effects of Pell Grants and the unconditional effects.

function of years of college (s_t) and ability (a_t), which is not perfectly known to the student. Let the cost of education at time t (c_t) be the difference between announced tuition (T_t) and financial aid. Financial aid contains two components: the need-based component is a function of initial wealth (I_0) and the number of children attending college at time t (n_t); the merit-based component is a function of perceived ability at time t . Let $E_t[\]$ denote the expectation operator conditional on information at time t .

$$(1) \quad w_{it} = f(s_{it}, a_i)$$

$$(2) \quad c_t = T_t - g(I_0, n_t) - h(a_t)$$

A student will attend a first year of college if the expected value of increased lifetime earnings for that year exceeds the cost of attending college (including forgone earnings).

$$(3) \quad E_0 \left\{ \sum_{t=2}^T R^{t-1} [f(s_{it} = 1, a_i) - f(s_{it} = 0, a_i)] \right\} > E_0 [f(s_{i1} = 0, a_i) + T_1 - g(I_0, n_1) - h(a_i)]$$

At the start of the first year, a student will indicate an intention to attend a second year as well so long as

$$(4) \quad E_0 \left\{ \sum_{t=3}^T R^{t-2} [f(s_{it} = 2, a_i) - f(s_{it} = 1, a_i)] \right\} > E_0 [f(s_{i2} = 1, a_i) + T_2 - g(I_0, n_2) - h(a_i)].$$

We could solve the decision rules for the maximum tuition level that a student would be willing to pay. For simplicity, let's assume that tuition is fully known one year in advance.

$$(5) \quad E_0(T_1^*) = T_1^* = E_0 \left\{ g(I_0, n_1) + h(a_i) - f(s_{i1} = 0, a_i) + \sum_{t=2}^T R^{t-1} [f(s_{it} = 1, a_i) - f(s_{it} = 0, a_i)] \right\}$$

and

$$(6) \quad E_0(T_2^*) = E_0 \left\{ g(I_0, n_2) + h(a_i) - f(s_{i2} = 1, a_i) + \sum_{t=3}^T R^{t-2} [f(s_{it} = 2, a_i) - f(s_{it} = 1, a_i)] \right\}$$

These tuition levels are likely the formulae that students use to make any decisions about the second year of school that must be made during the first year. For example, a student wanting to transfer to another university must file that application during the first year. Also, students who want financial aid in their second year must file applications during their first

year. With these types of decisions in mind, there are a few insights that come from comparing these two tuition values:

1. For a given level of ability, if the returns to schooling are linear (or even concave) in schooling and scholarship aid does not change, then the maximum tuition that a student will pay falls over time. Hence, many students may rationally choose to get only one year of school.

2. Even if the returns to schooling are convex and scholarship aid does not change, then the maximum tuition a student is willing to pay may still decrease over time, leading to more planned attrition.³

3. Even expected changes in financial aid can alter the maximum that students would be willing to pay, leading to students to plan on withdrawing or transferring.

4. Since students must apply for second-year financial aid during their first year, they will do so only if they perceive that their benefits exceed costs in both periods.⁴

There are also a number of decisions about the second year that can be made after the first year—for example, the decision to withdraw from college altogether. Students make these decisions after gaining another year of information on which to base their decisions. The student will choose to attend another year if the expected value of the increase in lifetime earnings for the second year exceeds the cost of attending college that year (including forgone earnings).

$$(7) \quad E_1 \left\{ \sum_{t=3}^T R^{t-2} [f(s_{it} = 2, a_i) - f(s_{it} = 1, a_i)] \right\} > E_1 [f(s_{i2} = 1, a_i) + T_2 - g(I_0, n_2) - h(a_i)]$$

We could rewrite this decision rule solving for the maximum tuition levels that the student would be willing to pay in order to actually attend a given year of college:

$$(8) \quad T_2^* = E_1 \left\{ g(I_0, n_2) + h(a_i) - f(s_{i2} = 1, a_i) + \sum_{t=3}^T R^{t-2} [f(s_{it} = 2, a_i) - f(s_{it} = 1, a_i)] \right\}$$

3. Holding scholarship aid constant and as $T \rightarrow \infty$, the maximum tuition rises only if the following inequality is satisfied:

$$(1 - r)E_0 [f(s_{it} = 2, a_i) - f(s_{it} = 1, a_i)] > E_0 [f(s_{it} = 1, a_i) - f(s_{it} = 0, a_i)]$$

If in the extreme case, there is a “sheepskin” effect of a degree (i.e., returns only to a two- or four-year degree), then the inequality is always satisfied. Typical models of sequential investment show that the willingness to pay increases over time. The key difference is the usability of capital. Students may be able to drop out of colleges and use a year of college in the labor market.

4. There may be a small group who apply for financial aid even though they expect not to attend the second year. There is an option value to applying for financial aid because ability is not known perfectly (see discussion by Sarah Turner in chapter 1 in this volume).

Notice that the difference between equations (8) and (6) is the information set. Students have a chance to update their expectations with information from their first year of school. As the model stands, the updating comes in terms of ability. Similar to the model in Manski (1993), students discover their ability by attending college. Knowing the ability then changes the willingness to pay.

We could have also changed this model by introducing uncertainty in the financial aid formula. Unexpected changes in financial aid might lower the maximum tuition price that students might be willing to pay. For example, if a student's expected financial aid offer falls, the maximum that a student would be willing to pay declines. The student may wish to transfer to a cheaper school or drop out altogether.

A simple insight of the model is that changes in financial aid matter. Previous work on the effects of financial aid has looked at relationships between student outcomes and both changes and levels of aid. Recent work by Wetzell, O'Toole, and Peterson (1999) look at changes in financial aid for students at Virginia Commonwealth University. They find that increases in need-based financial aid likely improved student retention. Other work by Singell (2001) looks at the effects of the level of financial aid in the first year. He finds that the higher the student's levels of need-based financial aid, the more likely the student is to graduate.

While the model in this paper suggests that decisions about enrollment in the second year rely on financial aid changes rather than levels, there may be reasons that the level of financial aid in the first year matters. If the level of financial aid in the first year creates some inertia or helps to shape expectations about the financial aid offer in the second period, the level may affect the student in the next year. One example of this type of effect is the application for second-year financial aid. The higher a student's Pell Grant in the first year, the more likely the student will apply for a second-year award. If students expect to get a low second-year award, they may never even apply for financial aid.

The level of financial aid may even have deleterious consequences on the student. The model implicitly assumes that financial aid does not change students' behavior in other ways. But for example, if a student receives financial aid, he or she may be more detached from college. The student may not fully engage and take college seriously since his or her money is not on the line. In this way, Pell Grants eliminate the "sunk cost fallacy" for the student.⁵ Since Pell Grant recipients did not make this initial investment (the federal government did instead), they may not have as strong of an incentive to work hard in school as students who made this initial investment themselves. As a result, Pell Grant recipients may perform worse. The model may capture some of this through the updates on students' abilities.

5. The "sunk cost fallacy" suggests that people often devote greater resources and more effort in areas where they have already made an investment.

Controlling for students' performance in school gives a clear indicator of whether they are exerting effort. The empirical results will investigate this hypothesis more fully.

Although they are outside of the scope of this paper, there are other outcomes in which we might be interested that are related to either the level or change in financial aid. For example, we might be interested in how financial aid affects the number of credits that a student successfully attempts. Students without financial aid may be reluctant to take loans and may spend more time working on the side. On the other hand, students without financial aid may want to cram in more credits per semester to try to reduce the number of semesters they have to attend (and as a result the total cost of college). We might also be interested in knowing how grade point averages (GPAs) vary with financial aid. In particular, if the level of financial aid affects hours attempted, completed, or GPAs, it might also affect students' perception of their abilities and, in the context of the previous model, affect their likelihood of completing college.

5.2 Data

The data for this project come from the OBR. Through a collaborative agreement with the OBR, the OBR has allowed me to access anonymous student data from Ohio's public institutions. The data are provided by the respective institutions to the OBR and include information on student demographics, enrollment, credit hours completed, and GPAs.

The OBR has collaborative arrangements with other agencies that allow them to expand the data. For example, the OBR links the student records to ACT and Scholastic Aptitude Test (SAT) records. Most Ohio students take the ACT exam, and the ACT records include the highest test score of the student and the most recent responses to the ACT survey (which includes student-reported data on high school performance). The OBR also links students to their respective FAFSA. The FAFSA data include detailed information about the finances of both students and their families. From the FAFSA, the variable of most interest is the "Estimated Family Contribution" that colleges use to award grants based on financial need.

One important limitation of the data is that they only include information about need-based financial aid. From FAFSA data, we know students' eligibility for federal grants and loans. We also know students' eligibility for Ohio's Instructional Grant program, a state-run need-based financial aid award. The data do not include information about merit-based financial aid. Ohio institutions are reluctant to divulge merit-based awards since these rewards are central to their recruitment strategies. While I do not observe merit aid, I observe students' GPAs once in college, their ACT scores, and their (self-reported) high school GPAs. If these variables adequately control for student ability and if colleges determine need- and merit-based

awards separately, then not knowing students' merit-based awards should not affect the estimated results.

Another limitation of the data is that they only include students attending Ohio public universities. Students from Ohio that attend universities in other states, including the nation's elite schools, and students that attend private schools in Ohio are excluded from the sample.⁶ These exclusions are both a weakness and a strength of these data. Excluding elite students may make the results not generalizable to all college students; however, excluding elite students gives us the opportunity to describe how financial aid affects students at nonelite schools. These nonelite schools educate the majority of college students and may be places where financial constraints are more binding.

Another concern related to the inclusion of Ohio public institutions alone is the measurement of dropout behavior. Students who transfer from Ohio public institutions to institutions located in other states are indistinguishable in the data from students who withdraw from Ohio public universities. This potential bias, however, should be very small because the percentage of students who probably transferred makes up a small fraction of the total number of observed dropouts.⁷

I focus entirely on the incoming freshman class in the 1999–2000 school year. These are the first students for whom FAFSA data are available through the OBR. I include students who enrolled in any college, including community colleges, for the first time in 1999, and I track these students through the 2000–2001 school year.

Table 5.1 provides summary statistics for the sample. At four-year institutions, about 10 percent of incoming full-time freshmen are from other states, and students are much less likely to be commuter students than at two-year colleges. At two-year colleges, which include local and state-run community colleges and technical colleges, about 2 percent of all students live on campus. Similar to other national surveys, the average age of first-time freshmen at two-year colleges is considerably higher than at four-year colleges, and students complete fewer semester hours in their first year (thirteen at four-year colleges as compared to eleven at two-year colleges).

Seventy-five percent of incoming freshmen at Ohio's four-year colleges

6. Ohio State University and Miami University are the top ranked public universities in Ohio. In the 2004 version of *U.S. News & World Report's* college rankings, they rank 60th and 64th, respectively, among national universities with doctoral programs. Other high ranking institutions in Ohio (e.g., Oberlin) are private colleges.

7. The Integrated Postsecondary Education Data System (IPEDS) tracks the number of transfers at each institution but does not record the state of residence of transfer students although it does track the states of residence for incoming freshmen. Assuming that transfer students are geographically representative of the incoming freshman class, then one would expect around 650 Ohio students to transfer to the non-Ohio schools with substantial Ohio enrollments. If we further assume that *all* 650 transfer students just finished their first year of school, then about 4.3 percent of observed dropouts are actually transfer students.

Table 5.1 Average Student Characteristics

	4-Year College	2-Year College
Out-of-state student	.103	.032
Lives on campus	.557	.024
Age	18.8 (2.5)	21.0 (6.2)
Nonwhite	.134	.173
Hours completed by fall 1999	13.4 (4.7)	11.2 (6.3)
Left institution after 1 year	.278	.491
Left higher education after 1 year	.201	.431
Took ACT exam	.750	.446
ACT composite score (36 = max)	21.8 (4.3)	18.9
Filed FAFSA for fall 1999	.653	.628
Uncovered financial need (\$)	423.0 (716.1)	24.1 (172.4)
Uncovered financial need conditional on being >0	1,081.8 (773.7)	82.0 (310.5)
Filed FAFSA for fall 2000	.490	.399
Uncovered financial need (\$)	1,261.5 (951.4)	21.3 (93.1)
Uncovered financial need conditional on being >0	1,400.2 (900.2)	41.3 (126.5)
Change in Pell Grant (conditional on Pell eligibility in 1999 or 2000)	1,691 (994)	881 (762)

Source: Author's calculations from unpublished data from the Ohio Board of Regents.

Notes: Standard deviations appear in parentheses for nonbinary variables. Data are for full-time students who first entered Ohio public colleges and/or universities in fall 1999. Uncovered financial need equals tuition less the estimated family contribution from the FAFSA less any Pell Grant for which the student was eligible.

took the ACT exam while only 45 percent of students at two-year colleges took the exam. The four-year college students performed better than the two-year college students. Throughout the paper, I will at times restrict the sample to students who took the ACT exam. Not only do I know these students' test scores, but I also have additional (self-reported) data on these students' high school experiences.

Throughout the paper, I will also restrict the sample at times to those students who filed a FAFSA in both fall 1999 and fall 2000. About 65 percent of four-year students and 63 percent of two-year students submitted FAFSAs in 1999; however, only 49 percent and 40 percent, respectively, filed FAFSAs in 2000. As explained below, not observing FAFSA data for many applicants leads to substantial biases in the results using panel identification. The average, uncovered financial need is small across all students, but conditional on it being positive, the uncovered financial need is slightly greater than \$1,400 for students at four-year schools.

Table 5.2 shows some basic least squares regressions of student stop-out

Table 5.2 Association between Financial Aid and Stop-Out Behavior (dependent variable: student stopped out)

	All Students				Students Taking ACT Exam	
	(1)	(2)	(3)	(4)	(5)	(6)
Level of financial aid (000s)	.033 (.002)	.006 (.003)	.0002 (.003)	-.005 (.003)	.005 (.003)	.002 (.003)
Log of parents' income		-.036 (.005)	-.030 (.005)	-.042 (.005)	-.030 (.006)	-.024 (.005)
Out-of-state student		.072 (.077)	.055 (.073)	.041 (.073)	.365 (.152)	.165 (.172)
Age		.022 (.004)	.018 (.004)	.020 (.004)	.018 (.005)	.016 (.005)
Male		.029 (.005)	.003 (.004)	.003 (.004)	.017 (.005)	-.002 (.005)
Lives on campus		-.082 (.007)	-.057 (.007)	-.112 (.004)	-.079 (.007)	-.062 (.007)
Took the ACT		-.113 (.007)	-.063 (.007)	-.088 (.007)		
ACT score					-.0003 (.0007)	.004 (.001)
Freshman grade point average			-.138 (.003)	-.136 (.003)		-.132 (.003)
Includes high school GPA controls	No	No	No	No	Yes	Yes
Includes race fixed effects	No	Yes	Yes	No	Yes	Yes
Includes campus fixed effects	Yes	Yes	Yes	No	Yes	Yes
R ²	.108	.112	.204	.192	.091	.169
No. of observations	37,028	30,851	29,778	29,778	24,627	24,012

Source: Author's calculations from unpublished data from the Ohio Board of Regents

Notes: White standard errors are in parentheses. Data are for full-time students who first entered Ohio public colleges and/or universities in fall 1999. The sample varies across columns due to missing data on student characteristics or first-year GPA data.

behavior on the level of students' financial aid. These regressions are not meant to show the causal effect of Pell Grants but rather to demonstrate associations between the stop-out behavior, financial aid awards, and other covariates. These regressions are also useful in understanding the types of biases present in the data. Comparing the various specifications will help identify important biases.

Column (1) shows a regression of whether a student drops out or not regressed on the student's financial aid award. The regression includes fixed effects that control for the school that the student attends. The estimated coefficient is positive and significant, suggesting that larger awards are positively associated with dropout behavior. As mentioned before, however, these coefficients are significantly biased for a number of reasons. For example, when we include controls for an individual's socioeconomic back-

ground and personal characteristics in column (2), the estimated relationship drops significantly.⁸ The estimated relationship drops dramatically (from 0.033 to 0.006) and is still significant. Column (5) is similar to column (2) except that I focus only on students who took the ACT exam, including controls for a student's high school performance and entrance exam scores. The estimated coefficient is similar to column (2) but is no longer significant. The other rows in columns (2) and (5) suggest that wealthier students are less likely to stop out; out-of-state students are more likely to drop out; older students and men are more likely to withdraw; and students living on campus and students who took the ACT (and performed well on it) are less likely to stop out.

In column (3), I add controls for students' grades during their freshman years of school. As previously mentioned, a Pell Grant may have had a negative effect since students with a Pell Grant may have had less of a financial commitment to schooling and may not have worked as hard. Including GPA should control for these students and may further weaken the estimated relationship between financial aid and stop-out behavior. As shown in table 5.2, the estimated relationship is smaller than in columns (1), (2), and (5). The result is also not statistically different from zero. In column (6), I estimate a similar regression for students who took the ACT exam. Again, the estimated relationship is indistinguishable from zero.

In column (4), I estimate the relationship, controlling for personal and family characteristics and grades during a student's first year. I exclude the fixed effects for students' campuses of attendance. These fixed effects also control for differences in quality, price, and other unobservable campus characteristics (e.g., the strength of a campus's freshman intervention programs). These fixed effects control for the fact that students attending lower-quality schools (who also receive smaller financial aid awards because tuition is smaller) are more likely to withdraw than students attending better schools (who receive higher financial aid awards for similar reasons). Without these fixed effects, we would expect the estimated relationship between financial aid and student stop-out behavior to be even smaller and maybe even negative. This is exactly what column (4) of table 5.2 shows. The estimated relationship suggests that higher financial aid awards are negatively associated with student stop-out behavior.

5.3 Estimating the Causal Effects

There are three sources of variation that economists can use to identify the effects of Pell Grants: time series, panel, and cross-sectional. In this sec-

8. The sample sizes fall across columns in table 5.2. The first column includes all students in the sample. The second column includes only those students for whom demographic data are available. Columns (3) and (4) further restrict the sample to students who had a freshman GPA reported to the OBR. Columns (5) and (6) are based on the subsample of students who took the ACT exam.

tion, I discuss the feasibility of each of these identification strategies using the OBR data. I also present the basic empirical results for each identification strategy.

Each strategy aims at identifying the effects of Pell Grants *conditional* on initial enrollment. However, Pell Grants may affect enrollment, and, as a consequence, some students who would not be enrolled without Pell Grants would be included in the analysis, making the parameter estimated in this paper more difficult to interpret. If a Pell Grant affects enrollments, then the parameter estimated in this paper would be a combination of two different effects: (1) the effect of Pell Grants on the persistence behavior of students who would have attended college in the absence of the Pell Grant; and (2) the effect of Pell Grants on the persistence behavior of students who would not have attended school without the Pell Grant.

Under some circumstances, the latter effect does not matter. For example, any potential bias depends on the degree to which Pell Grants affect enrollment. If the Pell Grant has no effect, then the parameter estimated in this paper represents only the effect of the Pell Grant on all students eligible for Pell Grants. If the Pell Grant has a small effect, then any bias in the estimated parameter is likely to be small. Most of the research on Pell Grants and enrollment suggest that the Pell Grant has had either no effect or a very small positive effect on enrollment (see Kane 1999; Leslie and Brinkman 1987).

Additionally, if the unobservable factors that can potentially bias any estimate of the effect of Pell Grants on persistence are similar for those people for whom Pell Grants influenced enrollment decisions and those people for whom they did not, then the effect of Pell Grants estimated in this paper should be the same for both groups. For example, suppose that unobservable family characteristics affect both the size of students' Pell Grants and their likelihood of dropping out. We typically think that these unobserved characteristics likely bias our estimate of the effects of Pell Grants downward. If the bias is the same for students for whom Pell Grants did and did not affect initial enrollments, then the bias should be symmetric across both groups, and, while the overall parameter may be biased by these unobservables, the overall parameter is not made up of two components. If we can control for these unobservables in our empirical design, we can estimate a single parameter that is easily interpreted. However, if the determinants of persistence are not constant across groups, then the parameter will remain a combination of the two aforementioned effects.

5.3.1 Time Series Identification

One way to identify the effects of Pell Grants is to compare changes in students' outcomes after systematic changes in Pell Grants occur. For example, the Pell Grant program began in 1973. Previous work by Kane (1995) compares low-income student enrollment rates before and after the Pell Grant program was established. Kane finds that college rates grew

about 2.6 percentage points slower for low-income students than other groups, suggesting that the Pell Grant had little effect. Other systematic changes in Pell Grant formulae are described in Mortenson (1988). For the study at hand, I am presenting evidence for a single cohort, so time series variation will not be useful in identifying the effects of Pell Grants.

5.3.2 Panel Identification

Another way to identify the effects of Pell Grants is to look at changes in students' Pell Grants over time. While this seems like a promising strategy since the OBR data contain two years of data for a single cohort, there are a number of reasons why this strategy might be limited.

To see the limitations and possibilities of this identification strategy, we need to understand how variations over time are generated for a single individual. There are really three basic reasons that a student's Pell Grant would change from one year to the next. First, the generosity of the Pell Grant may change. This could be the result of systematic changes in the Pell Grant formula or by a change in college tuition. Such changes are likely to be exogenous, and if they generate enough variation, they may help researchers to accurately identify the effects of Pell Grants. Unfortunately, there is little variation over time in the period of time that the OBR data are available. From the 1999–2000 school year to the 2000–2001 school year, the maximum Pell Grant increased from \$3,125 to \$3,300, a 5.6 percent increase. Over the same time, tuition at Ohio schools increased by 5 percent across the board (OBR 2001).

Another source of variation comes from changes in students' college choices. Students may transfer to another school after the first year. The corresponding change in tuition will generate variation in students' Pell Grants. Unfortunately, this source of variation does not help identify the effects of Pell Grants. Students who transfer may have different abilities than those students who do not transfer. For example, a student with high ability may transfer from a two- to a four-year college to gain access to more opportunities. This student's Pell Grant would automatically increase. However, this increase is correlated with the student's ability and may confound causal estimates of the Pell Grant. Another reason why using variation from transfer behavior may be misleading is that the size of students' Pell Grants may affect transfer behavior, making it even more difficult to interpret and identify the effects of Pell Grants using variation caused by student transfer behavior. Thus, changes in Pell Grants resulting from transfer decisions will not generate variation in Pell Grants that can be legitimately used to identify the effect of Pell Grants alone.

A final reason that students' Pell Grants may change is due to changes in students' circumstances. Some changes may be legitimate sources of variation. For example, a family of four with one child in college may have a second child come of college age, causing the existing college student's Pell

Grant to increase. Similarly, a change in family size (e.g., birth of another child or separation) may increase a student's Pell Grant. Even the natural aging of parents should increase students' Pell Grants, although only slightly. However, there are other changes in family circumstances that may not be legitimate sources of variation. For example, changes in income due to unemployment or health shocks may reduce family income and consequently increase students' Pell Grants from year to year. These sources of variation may also affect the likelihood that a student persists in college.

For panel identification strategies to be successful, variation in the Pell Grants must come from sources that are exogenous from changes in students' stop-out behavior. As mentioned previously, the most legitimate changes come from changes in the Pell Grant formula, changes in tuition, and changes in family size or sibling attendance. I can use changes from these legitimate sources as instruments for actual changes in financial aid.

Constructing the instrument from changes in the Pell Grant formula and tuition is straightforward. I simply impute what students' Pell Grants would have been during the 2000–2001 school year, assuming that their financial and family information is unchanged from the 1999–2000 school year. The imputed 2000–2001 Pell Grant does not include variation from changes in students' (or their families') circumstances. It only includes variation arising from changes in the Pell Grant formula and tuition.

Imputing student data not only allows me to isolate exogenous variation, but it also allows me to estimate data for many students for whom financial data are missing. In the 1999–2000 school year, 35,233 students filed FAFSAs. However, 12,143 of these students did not file FAFSAs in the 2000–2001 school year (hereafter referred to as the “nonfilers”). These nonfilers are not a random subset of all students. These nonfilers include 2/3 of all students who withdrew from college after the 1999–2000 school year. For these individuals, I am missing financial data and information about changes in their siblings' college attendance for the 2000–2001 school year. Because the imputation assumes that students' financial information is unchanged from their first to second year of college, I can estimate data for these nonfilers.

If I had data for all nonfilers, I could estimate out the causal effect of Pell Grants by using the imputed grant as an instrument for the actual grant. The instrumental variable estimate would be an unbiased estimate of the effect. Unfortunately, because the actual data are not available for nonfilers, I can only estimate reduced-form regressions of stop-out behavior on the imputed Pell Grant. Because there is a significant, positive relationship between students' imputed Pell Grants and their actual Pell Grants, the reduced-form estimates of the relationship between imputed Pell Grants and stop-out behavior should give us a sense of the sign and significance of the effect of actual Pell Grants, but the reduced-form estimates will not give a precise estimate of the magnitude of such an effect.

Table 5.3 Ordinary Least Squares Regressions of Stop-Out Behavior on Changes in Pell Grants: Results with Panel Data (dependent variable: student stopped out)

	Full Sample: Students Filing FAFSA in 1999–2000			4-Year Students (4)	Students Filing FAFSA Both Years (5)
	(1)	(2)	(3)		
Imputed increase in financial aid (000s)	–.086 (.002)	–.092 (.002)	–.092 (.002)	–.064 (.003)	.018 (.008)
Level of financial aid in 1999–2000 (000s)		.025 (.002)	.026 (.002)	.020 (.002)	.006 (.002)
Includes covariates	No	Yes	Yes	Yes	Yes
Includes campus fixed effects	Yes	No	Yes	Yes	Yes
No. of observations	35,233	35,233	35,233	21,506	24,116

Source: Author's calculations from unpublished data from the Ohio Board of Regents.

Notes: Standard errors are reported in parentheses. Samples in columns (1) to (3) include all students who filed FAFSAs in 1999–2000. Column (4) focuses only on students who attended a four-year college in 1999–2000. Column (5) includes only those students who applied for financial aid in both years. Covariates include the following: an indicator for whether the student was from out of state; age; gender; whether the student lives on campus; whether the student took the ACT exam; student's freshman GPA; and controls for race.

Table 5.3 contains the reduced-form regressions regressing student stop-out behavior against the imputed increase in a student's Pell Grant from one year to the next measured in thousands of dollars. For students who withdrew or transfer, I impute the Pell Grant that students would have received if they remained at the same institution as their initial enrollment. Column (1) shows the results with fixed effects for school of attendance but no covariates. Column (2) shows the results without fixed effects for college of enrollment but with covariates for gender, age, campus living conditions, whether the student took the ACT exam, and GPA in the student's first years. Column (3) includes both fixed effects and covariates.

As column (1) shows, students whose Pell Grants increase are less likely to drop out. Without covariates, the coefficient suggests that a \$1,000 increase in a student's Pell Grant leads to an 8.6 percentage point decrease in the likelihood that the student withdraws. With covariates, the estimated coefficient implies that a \$1,000 increase in a student's Pell Grant corresponds to a 9.2 percentage point decrease in the likelihood that the student withdraws. These estimates suggest strongly and consistently that increases in financial aid decrease the likelihood that students withdraw from school.

Column (4) repeats the analysis, focusing only on the students whose initial college enrollment was at a four-year campus. As before, the estimated coefficient is negative and significant. A \$1,000 increase in students' imputed Pell Grants corresponds to a 6.4 percent reduction in the likelihood that students withdraw from college.

Column (5) reports estimates when nonfilers are excluded from the analysis. As mentioned before, the students who filed FAFSAs in both years are a nonrandom subset of all students. Now the estimate is positive and significant. A \$1,000 increase in a student's financial aid corresponds to a 2 percentage point increase in the likelihood that a student withdraws. I include the estimate of column (5) to provide some hint of what the bias may be from excluding the nonfilers in the previous columns. When we include the nonfilers, we get significant, negative relationships between increases in students' Pell Grants and the likelihood that students drop out; however, when these students are omitted, the estimates are positive and significant.

One might be able to further refine the estimates of students' Pell Grants in the cases where data are missing by using information about students' siblings. If ages or graduation dates were known or could be approximated, I could include this information in the estimation of what students' Pell Grants would have been in the 2000–2001 school year. Unfortunately, little information is available about students' siblings for the nonfilers.⁹

What conclusion should be drawn from the panel identification specifications? First, panel identification has only limited power to actually identify the effects of Pell Grants. Much of the variation created over time in a student's Pell Grant comes from sources that may also affect the probability that the student withdraws from school. It would be inappropriate to use this type of variation to identify the effects of Pell Grants. Second, the fact that many students, especially those who plan to withdraw from school, do not file FAFSAs in both years makes it difficult to estimate the effect of financial aid. When we impute data for these people, we find estimates suggesting that increases in financial aid reduce the likelihood that students withdraw from college.

5.3.3 Cross-Sectional Identification

One might also identify the effects of need-based financial aid by comparing the need-based awards of different students at a single moment in time. There are a number of reasons why students may have different need-based awards. Students may differ from each other in terms of personal income and assets, family income and assets, family size, parental age, college of attendance, and enrollment status (full- versus part-time). All of these differences will lead to differences in students' need-based financial aid. Much of this variation will not be helpful in identifying the effects of

9. Some information about the family (parental age, family size in 1999–2000, number of children in college in 1999–2000, parental marital status) may help predict changes in the number of children attending college; however, their predictive power is limited. When I model changes in sibling attendance on these variables, I get a very low R^2 . After rounding the predicted values to the nearest integer, the specification predicts that 0.05 percent of students who filed FAFSAs in both periods would have had a change in the number of siblings attending college. In reality, 20.5 percent of students had a change in the number of siblings attending college.

Pell Grants. These sources of variation will also likely affect students' dropout behavior, independent of need-based awards.

However, there is some variation across individuals that might be useful. In particular, differences in family size and the number of children in college may facilitate identification in a cross section. The Pell Grant formula contains a number of discontinuities, the largest of which is based on family size and the number of students attending college. Even these sources of variation may not be exogenous. For example, if a family can only afford to send one child to school, they may choose the student who has the most potential to benefit from college. This student's Pell Grant would likely be smaller than it would be for a comparable family that sent multiple children to school. However, in this example, comparing this solitary student to other families with multiple children attending would lead to a bias because the family with fewer children in college sent a child to college with a greater chance of succeeding. As a result of this potential bias, I will primarily focus on results that take advantage of discontinuities in family size. I will also briefly show estimates based on discontinuities in both family size and the number of children in college.

Table 5.4 shows the changes in Pell Grants that accompany changes in

Table 5.4 Pell Grant by Family Size and the Number of Children in College

Number in Family	Number of Children in College			
	1	2	3	4
<i>Income = \$40,000</i>				
2	\$2,175	\$2,475	n.a.	n.a.
3	\$2,875	\$2,775	\$2,775	n.a.
4	\$3,125	\$3,125	\$2,975	\$2,875
5	\$3,125	\$3,125	\$3,125	\$3,075
6	\$3,125	\$3,125	\$3,125	\$3,125
<i>Income = \$50,000</i>				
2	\$975	\$1,775	n.a.	n.a.
3	\$1,575	\$2,175	\$2,275	n.a.
4	\$2,325	\$2,425	\$2,575	\$2,575
5	\$3,125	\$2,825	\$2,775	\$2,775
6	\$3,125	\$3,125	\$3,075	\$2,975
<i>Income = \$60,000</i>				
2	\$0	\$975	n.a.	n.a.
3	\$400	\$1,475	\$1,775	n.a.
4	\$1,075	\$1,825	\$2,175	\$2,275
5	\$1,875	\$2,275	\$2,375	\$2,475
6	\$2,675	\$2,675	\$2,675	\$2,675

Source: Author's calculations.

Notes: Calculations assume that the families have zero assets and no student contribution in the computation of the estimated family contribution. Calculations also assume that students attend high-cost institutions. n.a. = not applicable.




Fig. 5.1 Pell Grants by family size

Source: Author’s calculations.

Notes: The estimated Pell Grants are formulated assuming no assets or student contribution. The Pell Grants further assume that only one person from the family is attending a high-cost institution.

family size. The table shows three different schedules linking family size and the number of children attending college.¹⁰ Each schedule corresponds to a different income level (\$40,000; \$50,000; or \$60,000). For example, among the families with \$50,000 in income, a family of two with one in college would receive a Pell Grant of \$975. If the family was actually a three-person family with one in college, then the Pell Grant would be \$1,575 per person. Figure 5.1 shows similar comparisons for different family sizes. There are three lines corresponding to family sizes of three, four, and five. The Pell Grants shown in the figure assume that only one child is attending college. As seen, the differences in family size can lead to systematic differences in students’ Pell Grants. These systematic differences create discontinuities that can be exploited to estimate the effect of Pell Grants on students.

Assuming that the differences between family size are unrelated to a stu-

10. The computations assume that the family has no assets and that the students do not contribute to the family’s estimated family contribution.

dent's success in college, comparisons can be made between families of different sizes who have the same number of children in college. However, as table 5.4 shows, there is heterogeneity in income (and thus Pell Grant) within a given family size. For discontinuity analysis to work, the families on either side of the discontinuity should be similar except for the discontinuity. As a result, when making comparisons across family sizes, we need to stratify the groups so that comparisons are made across relatively homogeneous groups (e.g., people with similar income and assets).

Intuitively, the easiest way to estimate the effect of the Pell Grant while taking advantage of this discontinuity is to use a Wald estimator (Wald 1940). To find the Wald estimator, one must first isolate two groups that are fairly homogeneous. Across the groups, the Wald estimator is found by taking the ratio of the differences across groups of the dependent variable (stop-out behavior) and the independent variable (size of the Pell Grant). For example, suppose we could identify all people who have low income and few assets and have one child in college. Some of these families are two-person families, and some are three-person families. Assuming that family size is uncorrelated with an individual's success in college, we could estimate a Wald estimator across these groups. Let y_i be the average withdrawal rate for group i . Group i takes on a value of 1 for the group of students in two-person families with one in college and 2 for the group of students in three-person families with one in college. Let x_i be the average Pell Grant for group i . The Wald estimator between these groups would then be

$$\beta_{\text{wald}} = \frac{y_2 - y_1}{x_2 - x_1}.$$

The denominator should be the expected change in the Pell Grant as a result of this discontinuity within this income-asset group. The numerator would be the difference in stop-out rates between these groups.

After computing the first Wald statistic, we could then create a Wald estimator between each income-asset grouping within the sets of two- and three-person families. If we had ten income-asset groupings, we would have ten Wald estimators. These Wald estimators can be combined by taking a weighted average of the estimators (weighted by the number of observations in each group 2). We could similarly create Wald estimators across adjacent groupings of family size. For example, we could compare three- and four-person families. Of course, in the estimation of each Wald statistic we would actually have multiple Wald statistics comparing income-asset groupings across each discontinuity.

While this approach seems straightforward, other discontinuities in the Pell Grant formula complicate the estimation of Wald statistics. For example, there are some income ranges where students would receive the maximum Pell Grant regardless of their family size or the number of children attending college. The Wald statistic would not be defined (or would

be greatly inflated) over these ranges. Similarly, the Wald statistic will not be defined for families that would have received no Pell Grant regardless of their family size or the number of children in college. Because these groups will likely create additional noise in the estimation, we may want to exclude these groups at times.

Before estimating the Wald statistics, we need to create the income-asset groupings needed to create comparisons between homogeneous groups. To create groupings, I reestimate each student's Pell Grant, assuming that he or she belonged to a two-person family with only one person attending college. I then divide this group into six subgroups, based on the revised Pell Grant:

1. People whose Pell Grant in a two-person family with one in college would have been at the maximum of \$3,125
2. People with revised Pell Grants between \$3,124 and \$2,001
3. People with revised Pell Grants between \$2,000 and \$1,001
4. People with revised Pell Grants between \$1,000 and \$401
5. People with revised Pell Grants at the Pell Grant minimum of \$400
6. People with revised Pell Grants equal to zero

Having uniform groupings across cells makes it much easier to estimate the Wald statistics and their standard errors. Creating groupings around Pell Grant values also avoids the problem that wealthier families are more likely to apply if they have more children. These families are identified in subgroup 6. Also, I separate people who would have had the Pell Grant minimum (\$400) because in the Pell Grant formula there is a discontinuity that allows families across a wider range of income to have this value of Pell Grant. I use these revised Pell Grants only for the purpose of creating homogeneous groups (i.e., identifying families with similar assets and incomes). When actually computing the Wald statistics, I use the actual Pell Grants.

Figure 5.2 shows an example of how the discontinuity works. The sample of students is from subgroup 3. These students would have had similar Pell Grants had their family size not been different. The left axis of figure 5.2 plots these students' actual Pell Grants across family size. The Pell Grant increases with the number of children. The right-side axis of figure 5.2 plots the stop-out rates for these students. The stop-out rate declines for students with larger families. Assuming that family size affects stop-out rates only through its effect on Pell Grant size, then stop-out rates are negatively related to Pell Grant size for these students. We could produce similar figures for each homogeneous group of students.

Table 5.5 shows the regression-based Wald estimates for the whole sample. Following Angrist (1991), the efficient combination of Wald estimators is just the instrumental variables estimate of y (stop-out behavior) on x (size of Pell Grant), where dummy variables for each homogeneous




Fig. 5.2 Average Pell Grants and stop-out rates across family size discontinuity

Source: Author's calculations from unpublished data from the OBR.

Notes: Sample is restricted to families whose assets and income are such that their college-age children would have had a Pell Grant between one and two thousand dollars if they had only had one child.

Table 5.5 Wald/Instrumental Variable (IV) Estimates of Effect of Financial Aid on Stop-Out Behavior

	Wald (1)	Wald ACT Sample (2)	Wald Dependent Variable = ACT Score (3)	IV (4)	IV with Campus Fixed Effects (5)	Wald Family and Child Discount (6)
Level of Pell Grant in 1999–2000 (in 000s)	–.037 (.009)	–.012 (.010)	.100 (.099)	–.029 (.004)	.0004 (.004)	–.036 (.009)
No. of observations	6,114	4,470	4,470	29,936	29,936	6,398

Source: Author's calculations from unpublished data from the Ohio Board of Regents.

Notes: White standard errors are reported in parentheses. Columns (1)–(3) and (6) include only students with positive Pell Grants less than the maximum. In columns (4) and (5) the instrument for “Level of Pell Grant” is the residual from a regression of Pell Grant on a quartic of the key variables determining Pell Grants (family income, assets, family size, number of children in college). The IV columns exclude families with income greater than \$115,000 or assets greater than \$150,000. Campus fixed effects are included in the first stage of the specification in column (5) to control for the fact that different school costs will lead to different size Pell Grants.

group and family size combination are used as instruments for x . For example, if all families were two, three, or four people in size, then I would include eighteen dummy variables as instruments (i.e., for each of the six homogeneous income/asset subgroupings defined previously, I would include three dummy variables for the possible family sizes).

Column (1) of table 5.5 shows the instrumental variable estimate when I estimate the regression excluding those for whom there is no variation in Pell Grants.¹¹ The estimated effect is negative and significant. A \$1,000 increase in Pell Grants stemming from differences in family size corresponds to a 4 percentage point decrease in likelihood that a student drops out. The result suggests that systematic differences in Pell Grants lead to differences in stop-out rates for students. Larger Pell Grants reduce students' probabilities of withdrawing.

In column (2) of table 5.5, I include only the students who took the ACT exam. The estimated effect suggests that a \$1,000 increase in a student's Pell Grant leads to a 1.2 percentage point reduction in the probability that a student withdraws although the result is not statistically significant. In column (3), I provide a specification test. Rather than use stop-out behavior as the dependent variable, I use the ACT score. If the specification is correctly identified, there should be no significant differences between ACT scores between groups. Indeed, the estimated relationship is indistinguishable from zero.

Another way to estimate the effects of a Pell Grant is to use an instrumental variable approach where the "delinearized" Pell Grant is used as an instrument for the actual Pell Grant. To do this, I run a regression of the actual Pell Grant on a quartic in the key variables that determine the Pell Grant (family income, family assets, family size, and number of children in the family attending college). The residuals from this regression should be made up primarily of discontinuities in the Pell Grant formula along these dimensions. I then use the residual as an instrument for the actual Pell Grant in a simple regression of stop-out behavior on students' Pell Grants.¹² I restrict the sample to students whose families have less than \$150,000 in assets or less than \$115,000 in annual income. The results appear in columns (4) and (5) of table 5.5. Similar to column (3), I find that a \$1,000 increase in a Pell Grant is associated with a 3 percentage point reduction in the likelihood that a student drops out.

In column (5), I use a similar procedure except that I put campus fixed effects in the regression that predicts students' Pell Grants. This is impor-

11. The maximum income in this group is \$25,000.

12. This is identical to running a regression of stop-out behavior on the Pell Grant value and a quartic in the variables that determine the Pell Grant. These strategies estimate the same estimator with only a negligible change in standard errors. I report the standard errors from the "instrument" approach because it may be intuitively easier to identify what source of variation is behind the estimated effect.

tant because different campus costs can lead to differences in the size of students' Pell Grants. Again, I use the residuals from this first stage as an instrument in the next. The results drop in magnitude and are statistically indistinguishable from zero.

The final column of table 5.5 estimates the effects of Pell Grants using both the discontinuities from family size and the number of children in college. I use the same methodology except now I compare differences in Pell Grants resulting from differences in the number of children going to college as well. The results are similar to those in column (1). I find that a \$1,000 increase in Pell Grants corresponds to a 4 percent reduction in the likelihood that students withdraw.

What conclusions should be drawn? First, while the estimates reinforce a negative relationship between the size of one's Pell Grant and stop-out behavior, they are not completely robust to specification. When I focus on the sample taking the ACT (column [2]) and when I include additional campus level controls (column [5]), the results are indistinguishable from zero. Additionally, there is substantial heterogeneity around the discontinuity, and efforts to create comparisons among homogeneous groups may not fully account for the heterogeneity. Overall the results seem supportive of those in the panel identification although the results are not robust to different subsamples.

5.4 Conclusion and Policy Implications

This paper set out to estimate the effects of Pell Grants on student retention. Using panel and cross-sectional variation as sources of identification, this paper attempts to estimate the relationship. The panel identification results suggest strongly that a Pell Grant reduces dropout rates. The regression-discontinuity results show similar results although they are more fragile. The regression-discontinuity results, however, are less likely to be biased by other factors than the specifications that use cross-sectional variation.

The finding that Pell Grants affect student persistence has several policy implications. Most importantly, it implies that federal and state need-based policies and aid matter and that they influence the likelihood that students continue from year to year in college. Even if these aid programs have no effect on enrollment (as Kane [1999] and Leslie and Brinkman [1987] show), need-based programs may impact educational attainment. The finding that aid increases persistence may suggest that front-loaded financial aid programs may improve student retention in the first years of college.¹³

13. DesJardins, Ahlburg, and McCall (2002) find that frontloaded aid programs have a "modest impact" on student retention.

The results in this paper may also support conclusions in Dynarski (2003) and Bound and Turner (2002). These papers focused on beneficiaries of the Social Security Student Benefit Program and the GI Bill, respectively. The papers found that need-based aid affected both enrollment and completion.

While many policymakers may view the finding that need-based aid improves retention as being “good” for society, it is not clear that this is so. As Jonathan Guryan points out in the comment that follows this chapter, more education may not be optimal for all students in society. There may be a number of students who do not have the skills and for whom a college degree may not improve earnings. If these students are the marginal students for whom the Pell Grant program influences college persistence, then the measured effect in this paper may reflect an inefficient use of societal resources.

In order to fully resolve the question of whether increased persistence is “good” for society, we would need to know how an additional year of education affects the earnings for these marginal students. In estimates of the returns of education for the entire population, Jaeger and Page (1996) argue that the returns to a second year of college have significant effects on earnings, particularly if students finish an associate’s degree. These estimates suggest that persistence into the second year of college is a positive outcome; however, the estimated return in Jaeger and Page (1996) is for the whole population. It is not clear that this estimated return would be the same for the marginal students affected by Pell Grants. Evidence from Tobias (2003) suggests that the returns to education are concentrated at the highest-ability students and remain small for lower-ability students. Moreover, as Sarah Turner points out in chapter 1 of this volume, the marginal students are likely less prepared for college and more likely to hold a General Education Development certification (GED) than a high school diploma.¹⁴ There is little empirical evidence on the returns to an additional year of college for these students.

Finally, there are two empirical points that should be considered in interpreting the results in this paper. First, the paper focuses solely on student persistence between the first and second year of college in adjacent years. Researchers have documented the growing trend of students to take breaks at various times during college (e.g., chap. 1 in this volume). DesJardins, Ahlburg, and McCall (2002) examine enrollment probabilities over a longer stretch of time. They find that need-based aid has no long-run effect on enrollment probabilities. They are using a sample in which students repeatedly enter and exit higher education. While the sample they use

14. Many of these students enroll in college remediation. Ongoing research by Bettinger and Long (2003) examines the effect of college remediation on college outcomes and student earnings.

focuses on students who initially enrolled at a four-year campus, their finding suggests that examining the effect of financial aid on retention over a longer stretch of time may be important.

Second, as states continue to gather more complete and expansive data on their students, economists and other researchers will be able to employ methods that exploit variation in Pell Grant formulae. However, even if better data are available, researchers should take care in how they deal with changes in Pell Grant eligibility and missing data. Students who leave school after their first year are less likely to file additional FAFSAs, making it difficult for researchers to measure their financial status. Research that fails to control for these missing data may be biased.

References

- Angrist, Joshua. 1991. Grouped-data estimation and testing in simple labor-supply models. *Journal of Econometrics* 47 (2–3): 243–266.
- Bettinger, Eric, and Bridget Long. 2003. The effect of remediation on student outcomes: The plight of underprepared students in higher education. Case Western University, Department of Economics. Mimeograph.
- Bound, John, and Sarah Turner. 2002. Going to war and going to college: Did World War II and the G.I. Bill increase educational attainment for returning veterans? *Journal of Labor Economics* 20 (4): 784–815.
- College Board. 2001. Trends in student aid 2001. Available at [<http://www.collegeboard.com/press/cost01/html/TrendsSA01.pdf>]. Retrieved 27 August 2003.
- DesJardins, S. L., D. A. Ahlburg, and B. P. McCall. 1999. An event history model of student departure. *Economics of Education Review* 18:375–390.
- . 2002. Simulating the longitudinal effects of changes in financial aid on student departure from college. *Journal of Human Resources* 37 (3): 653–679.
- Dixit, Avinash, and Robert Pindyck. 1994. *Investment under uncertainty*. Princeton, N.J.: Princeton University Press.
- Dynarski, Susan. 2003. Does aid matter? Measuring the effect of student aid on college attendance and completion. *American Economic Review* 93 (1): 279–288.
- Ehrenberg, Ron, and D. Sherman. 1984. Optimal financial aid policies for a selective university. *Journal of Human Resources* 19 (2): 202–230.
- Jaeger, David A., and Marianne Page. 1996. Degrees matter: New evidence on sheepskin effects in the returns to education. *Review of Economics and Statistics* 78 (4): 733–741.
- Kane, Thomas J. 1995. Rising public college tuition and college entry: How well do public subsidies promote access to college? NBER Working Paper no. 5164. Cambridge, Mass.: National Bureau of Economic Research, July.
- Kane, Thomas J. 1999. *The price of admission: Rethinking how Americans pay for college*. Washington, D.C.: Brookings Institution.
- Leslie, Larry L., and Paul T. Brinkman. 1988. *The economic value of higher education*. New York: MacMillan.
- Manski, Charles F. 1993. Adolescent econometricians: How do youth infer the returns to schooling? In *Studies of supply and demand in higher education*, ed.

- Charles Clotfelter and Michael Rothschild, 43–57. Cambridge, Mass.: National Bureau of Economic Research.
- Mortenson, Thomas. 1988. Pell Grant program changes and their effects on applicant eligibility 1973–74 to 1988–89. American College Testing Program Research Series Report. Iowa City, Iowa: American College Testing Program.
- Myers, S. C., and S. Majd. 1990. Abandonment value and project life. *Advances in Futures and Options Research* 4:1–21.
- Ohio Board of Regents. 2001. State-supported Ohio college and university performance report: Student outcomes and experiences. Available at [<http://www.regents.state.oh.us/perfrpt/2001index.html>]. Retrieved 29 August 2003.
- Seftor, Neil, and Sarah Turner. 2002. Back to school: Federal student aid policy and adult college age enrollment. *Journal of Human Resources* 37 (2): 336–352.
- Singell, Larry. 2001. Come and stay a while: Does financial aid affect enrollment and retention at a large public university? University of Oregon, Department of Economics. Working Paper.
- St. John, Edward P., Shouping Hu, and Tina Tuttle. 2000. Persistence by undergraduates in an urban public university: Understanding the effects of financial aid. *Journal of Student Financial Aid* 30 (2): 23–37.
- St. John, Edward P., Shouping Hu, and Jeff Weber. 2000. Keeping public colleges affordable: A study of persistence in Indiana's public colleges and universities. *Journal of Student Financial Aid* 30 (1): 21–32.
- St. John, Edward P., Glenda D. Musoba, and Ada B. Simmons. 2003. Keeping the promise: The impact of Indiana's 21st-Century Scholars program. Indiana Education Policy Center research report. Available at [<http://www.indiana.edu/~iepc/hepolicy.html>]. Retrieved 4 September 2003.
- Tobias, Justin. 2003. Are returns to schooling concentrated among the most able? A semiparametric analysis of the ability-earnings relationships. *Oxford Bulletin of Economics and Statistics* 65 (1): 1–29.
- U.S. Department of Education. 22 May 2002. Paige hails House for not leaving Pell Grant recipients behind. Available at [<http://www.ed.gov/PressReleases/05-2002/15222002.html>]. Retrieved 27 August 2003.
- Wald, Abraham. 1940. The fitting of straight lines if both variables are subject to error. *The Annals of Mathematical Statistics* 11 (3): 284–300.
- Wei, Christina, and C. Dennis Carroll. 2002. Persistence and attainment of beginning students with Pell Grants. NCES Publication 2002169. Washington, D.C.: National Center for Education Statistics.
- Wetzel, James, Dennis O'Toole, and Steven Peterson. 1999. Factors affecting student retention probabilities: A case study. *Economics of Education Review* 23 (1): 45–55.

Comment Jonathan Guryan

Eric Bettinger presents a careful analysis of an overlooked but important question. Economic researchers have noted the large sums of money the

Jonathan Guryan is assistant professor of economics at the University of Chicago Graduate School of Business, and a faculty research fellow of the National Bureau of Economic Research.

United States spends on subsidies for college costs annually and have responded with a number of studies on the effect of subsidies on college matriculation decisions. But few have investigated the impact of these subsidies on what Bettinger calls “stop-out” behavior, the choice to leave college after the first year. One might ask why stop-out behavior is an interesting outcome to examine. The answer is simple: A lot of college students do it. Twenty-eight percent of first-time full-time freshmen in Ohio do not return to school the next year. Depending on whether one considers stopping out an unambiguously negative outcome—a point I will discuss shortly—this number may be a signal of a large problem. What are the likely causes of stop-out decisions? Significant suspects include failing out, realization that the decision to attend college was wrong for the student, changes at home, cost concerns, or an inability or unwillingness to pay tuition. The latter factor is one that is easily manipulated by policy, so Bettinger asks whether subsidies affect stop-out decisions.

Bettinger has compiled a unique and detailed data set for this project. The data from the Ohio Board of Regents track every college student in the state of Ohio for the 1999–2000 and 2000–2001 school years and are merged to detailed financial aid data. The data allow Bettinger to measure stop-out behavior better than has been possible in the past. Previously, researchers could not distinguish students who left college to return home to work from students who transferred to a different school. Bettinger is able to track any student who transfers to another college in Ohio.

The linked detailed financial aid data allow Bettinger to employ two estimation strategies. In the first strategy, Bettinger estimates whether changes in Pell Grants for individual students from the first year of college to the second affect stop-out decisions. He isolates variation in Pell Grant changes that is otherwise unrelated to stop-out behavior. In the second strategy, Bettinger compares the stop-out behavior of students with slightly different characteristics that cause them to have significantly different Pell Grants. Both empirical strategies suggest that larger Pell Grants lead to a smaller propensity to leave college after one year.

I would like to make two points about this research agenda. First, it is important to understand the effect the Pell Grant has on college matriculation if we want to interpret the estimates of the effect of the Pell Grant on stop-out behavior. Because we do not have conclusive evidence on the former question, we must be careful when interpreting the evidence brought to bear on the latter. A simple empirical model will help to clarify this point. Bettinger estimates the effect of Pell Grants on stop-out behavior in two ways. Each method is intended to isolate the causal effect of a dollar of Pell Grant on the likelihood that an enrolled freshman will not return to school the following fall. Consider the following model of stop-out behavior for students currently enrolled as freshmen:

$$(1) \quad \text{stopout}_{it+1} = \beta_0 \text{Pell}_{it+1} + \beta_1 \text{Pell}_{it+1} \cdot I_{it} + \beta_2 I_{it} + \mathbf{X}_{it+1} \delta + \varepsilon_{it+1},$$

where stopout_{it+1} indicates that the student does not return to school in year $t + 1$, Pell_{it+1} is the size of the Pell Grant for which the student is eligible in year $t + 1$, \mathbf{X}_{it+1} is a vector of individual characteristics, ε_{it+1} is a random error term, and I_{it} is a variable that indicates whether the student was induced to attend college by the Pell Grant in year t .¹ Without a good estimate of which students were induced to begin college by the Pell Grant, we are forced to estimate the model excluding $\text{Pell}_{it+1} \cdot I_{it}$ and I_{it} .

For most policy questions, we are interested in knowing the magnitudes of β_0 and β_1 . These parameters tell us how sensitive particular students are to changes in the price of college. However, because we are forced to estimate the model without information indicating which students matriculated because of the Pell Grant, we can only estimate the correlation between the Pell Grant and stop-out behavior among freshmen college students. For ease of exposition, allow me to partial out the variation in demographic characteristics (\mathbf{X}_{it}) and denote with asterisks the orthogonal variation in the remaining variables. The regression estimate of the effect of the Pell Grant on stop-out behavior is

$$(2) \quad \beta = \beta_0 + \beta_1 \cdot \frac{\text{cov}(\text{Pell}_{it+1}^*, \text{Pell}_{it+1}^* \cdot I_{it}^*)}{\text{var}(\text{Pell}_{it+1}^*)} + \beta_2 \cdot \frac{\text{cov}(\text{Pell}_{it+1}^*, I_{it}^*)}{\text{var}(\text{Pell}_{it+1}^*)}.$$

The first two terms confirm Bettinger’s contention that the estimated effect combines the price sensitivity of both groups of students. There is reason to believe that students induced to matriculate by the Pell Grant are more price sensitive on the stop-out margin as well ($\beta_1 < 0$). If we are interested in how many students continue past their first year of college because of the Pell Grant program, then we need to know the magnitude of β_0 . To the extent that $\beta_1 < 0$, Bettinger’s estimate is an overestimate of the magnitude of β_0 .

The final term highlights that there is an additional source of bias. It is likely that $\text{cov}(\text{Pell}_{it+1}^*, I_{it}^*) > 0$. All students induced to begin college by a Pell Grant received a Pell Grant, while most of the rest of the population received no Pell Grant. Additionally, a marginal increase in the Pell Grant may increase the likelihood that $I_{it} = 1$ but should not decrease that likelihood. It is also likely that students induced to begin college by the Pell Grant have an inherently larger propensity to leave college after the first year ($\beta_2 > 0$). If true, this combination of factors would induce a positive bias in the estimate of the causal effect of Pell Grants on stop-out behavior.

1. Note that the cross-sectional regressions presented in section 5.3.3 estimate β_1 , while the panel estimates presented in section 5.3.2 estimate a similar model in which Pell_{it} is replaced with ΔPell_{it} .

The second point I would like to discuss is the idea that reducing stop-out behavior is not necessarily a good outcome. Like all goods, education comes at a cost. The economy as a whole bears the cost of allocating resources to provide education, and the individual investing in education bears a related cost, including both explicit tuition payments and forgone earnings. As economists, we usually assume that an individual weighs the potential benefits of attending college against the costs she must bear to do so. The Pell Grant, like any subsidy of higher education, lowers the net price of college. Viewed this way, it would be quite surprising if there were no resulting change in eligible students' propensity to stay in college. Anyone close enough to the margin between staying in college and dropping out should be significantly influenced by the change in price. The question at hand is how many students are close enough to that margin. Our evaluation of the Pell Grant program should also consider whether the subsidy helps the economy utilize resources more efficiently, by directing education resources to those who would benefit most or by directing the economy's resources to those who value them most.

Indeed, the increased education for the Pell Grant recipient comes at a cost. Resources in the form of teachers' time, library and computing facilities, and classroom space could have been used to educate someone else. The government funds could have been spent on welfare, on the Center for Disease Control, to buttress the Social Security trust fund, or to lower taxes. So we must ask two important questions. First, can students with the help of their parents effectively weigh the benefits of a college education against the costs they will have to bear? And second, are the costs that students must bear to attend college different from the total cost to the economy to provide the resources necessary to educate the student? These questions are addressed in multiple chapters of this book. It seems reasonable that many high school students, when making this decision, respond to incentives they probably should ignore and poorly estimate the costs and benefits of college. Risk-averse students who heavily discount the future may fear the uncertainty associated with the benefits of college that will not come for many years. All of these factors can cause students to drop out of school before they reach the level that would maximize their lifetime wealth.

If Pell Grants correct systematic mistakes made by students, as discussed in other chapters of this book, or if they compensate for higher borrowing costs that are due to discrimination or asymmetric information, then we should conclude such education price subsidies improve how education resources are allocated. If Pell Grants reduce stop-out behavior by students who are highly risk-averse or greatly discount the future, then supporters of college subsidies must justify the paternal nature of the federal government telling college-aged students what will be best for them when they "grow up." If Pell Grants increase the educational attainment of students

solely because the price of college is lower, then it would seem easy to argue that there are better uses of government funds.

Bettinger has taken an important first step in recognizing the prevalence of stop-out behavior and in analyzing the effect of the Pell Grant on students' decisions to continue in college past the first year. He has comprised a unique data set, which is sure to produce additional important research going forward. The results indeed suggest that changes in individual Pell Grants affect year-to-year college attendance. The next important question in this research agenda is why.

