

IZA DP No. 4362

Do No-Loan Policies Change the Matriculation Patterns of Low-Income Students?

Glen R. Waddell
Larry D. Singell, Jr.

August 2009

Do No-Loan Policies Change the Matriculation Patterns of Low-Income Students?

Glen R. Waddell

*University of Oregon
and IZA*

Larry D. Singell, Jr.

University of Oregon

Discussion Paper No. 4362
August 2009

IZA

P.O. Box 7240
53072 Bonn
Germany

Phone: +49-228-3894-0
Fax: +49-228-3894-180
E-mail: iza@iza.org

Any opinions expressed here are those of the author(s) and not those of IZA. Research published in this series may include views on policy, but the institute itself takes no institutional policy positions.

The Institute for the Study of Labor (IZA) in Bonn is a local and virtual international research center and a place of communication between science, politics and business. IZA is an independent nonprofit organization supported by Deutsche Post Foundation. The center is associated with the University of Bonn and offers a stimulating research environment through its international network, workshops and conferences, data service, project support, research visits and doctoral program. IZA engages in (i) original and internationally competitive research in all fields of labor economics, (ii) development of policy concepts, and (iii) dissemination of research results and concepts to the interested public.

IZA Discussion Papers often represent preliminary work and are circulated to encourage discussion. Citation of such a paper should account for its provisional character. A revised version may be available directly from the author.

ABSTRACT

Do No-Loan Policies Change the Matriculation Patterns of Low-Income Students?*

We empirically examine whether there is discernable variation in the matriculation patterns of low-income students at public flagship institutions in the United States around changes in institutional financial-aid policies that target resident, low-income students with need-based aid. While enrollment responses cannot be attributed to these programs, we do find that institutions that introduce income-targeted aid subsequently enroll financially needier and geographically more-distant students. These findings imply that "improved" access may actually displace some needy students in favor of others.

JEL Classification: I23, I21, J24

Keywords: low income, financial aid, no loan, Pell

Corresponding author:

Glen R. Waddell
Department of Economics
University of Oregon
Eugene, OR 97403-1285
USA
E-mail: waddell@uoregon.edu

* We thank Tim Bramby and seminar participants at University College Dublin, University of Edinburgh, University of Dundee, University of Aberdeen, and University of Guelph for comments on earlier versions of this manuscript. We also thank Peter Stiffler for helpful research assistance. Any and all errors remain our responsibility.

1 Introduction

Despite significant resources being spent on need-based financial aid in the United States, in the last three decades the gap between low- and high-income students' matriculation rates has not only persisted but has widened (e.g., Ellwood and Kane (2000), Havemen and Wilson (2007)). As such, there continues to be significant interest in understanding the underlying patterns of matriculation and any potential sensitivity to policy aimed at narrowing this gap. Public flagship institutions represent an important gateway into higher education for needy students. Having access to the universe of Pell recipients from 1999-2007, we build on existing literature by examining whether income-targeted institutional-aid programs initiated at these institutions influence the matriculation patterns of low-income students along several dimensions of access. The Pell Grant is the largest single source of need-based aid representing \$13.7 billion of appropriations in 2007/2008 and representing over nearly 25 percent of all college going students.¹

Need-based aid programs have grown, in part, out of a recognition that the rising real cost of college and student debt levels have potentially threatened the access of low-income students to college. For example, while real tuition and fees rose by 63 (43) percent at private (public) universities between 1993 and 2005 (*Trends in College Pricing 2005*, The College Board, Table A1), Mishel, Bernstein, and Allegretto (2007) find that the real wage premium rose by 27 percent. Consistent with tightening financial constraints, Snyder, Tan, and Hoffman (2006) show that students are more likely to require student aid to attend college, that students are covering a smaller portion of their college costs with grants, and, are taking out nearly twice the level of debt in real terms than in the previous decade (e.g., over \$20,000 in 2004). Partly because federal need-based aid programs have not kept pace with rising costs (McPherson and Schapiro (1998)), a growing number of states and institutions have moved toward providing need-based aid, arguably in order to maintain or increase the accessibility of college to needy students.

In our consideration of the efficacy of income-targeted aid initiatives, we exploit data covering all dependent Pell Grant recipients in academic years 1999 through 2007, thinking of efficacy being evidenced in three ways. First, we consider efficacy as evidenced in institutional enrollments. Establishing the nature of any link between these programs and enrollment is, of course, of fundamental interest if not importance. This is particularly true as the extant literature can imply that this population of students is quite inelastic with respect to aid (e.g., Curs, Singell, and Waddell (2007)). Where conditions allow for a valid test – which we will discuss in some detail below – we find no systematic enrollment responses to the policy innovations.

Second, we consider efficacy as evidenced by the incomes of matriculating-students' families. Without family-income data, directly, we will analyze the expected family contributions (*EFC*) of matriculating low-income students. In general, *EFC* is constructed to represent the amount the applicant-student and/or student's family can be expected to contribute toward financing the degree being sought. This expected contribution is estimated from information regarding the family's income, allowances against their income, the number of children, and their assets. To focus on the typical college-age student, our analysis uses only

¹Federal Pell Grant End-of-Year Report 2007-2008, U.S. Department of Education, Washington, D.C..

dependent students who by definition rely on parental contributions to attend college. We establish that the relative need of those who matriculate to at aid-innovating institutions after the introduction of targeted-aid programs supports that these aid innovations have lead to a shift toward lower-income Pell students.

Third, we consider efficacy as evidenced in institutions' geographic basins of attraction. Here we are asking whether these policies change the geographic dispersion of the low-income students enrolling at these institutions. As other associated costs are unchanging over the period of our analysis, we will be inclined to interpret any increase in an institution's basin of attraction as akin to them being better able to draw in low-income students. That is, with an objective of aid policies being to facilitate the matching of needy students with institutions, any increase in an institutions basin of attraction is consistent with prior constraints on the matriculation of low-income students being relaxed. We show that these "basins of attraction" do change, with students tending to travel relatively farther distances to aid-innovating institutions after the policy changes. We interpret this evidence as suggestive that aid policies such as those analyzed can relax geographic barriers (direct, or indirect) to the enrollment of low-income students. While evaluating the match itself remains outside our purview, we expect that such movement is in the direction of improving the potential matching of students with institutions.

In the following section we discuss the related literature and address current shortcomings that, in part, motivate our analysis. In Section 3 we set up our empirical test given the changes to financial-aid policy that fall within the period of data availability. In the process of doing so, we will identify our exact identification strategy and the way in which one should consider "treatment" and "control" groups for the purpose of the analysis. In defining this test, we will also address the inappropriateness of alternative specifications (conducting a similar analysis at private institutions, for example) and set up a falsification exercise that we will perform along side our examined relationships throughout the paper. In Section 4, we describe the data used in our analysis and provide some summary statistics that are new to the literature. We then devote separate sections to each of the three relationships we examine; enrollment in Section 5, financial need in Section 6, and basins of attraction in Section 7. We then share some concluding remarks and discuss several important questions that remain unanswered.

2 Literature

In the general area of inquiry into student responses to aid, much of the recent work has risen out of natural experiments in federal, state, and institutional-aid programs (e.g., Angrist (1993), Bound and Turner (2002), Long (2004), Cornwell, Mustard, and Sridhar (2006)). Among them, there are several examples where general aid initiatives (i.e., not income-targeted) are used to consider the matriculation of low-income students (e.g., Hansen (1983), Kane (1994), Dynarski (2003)). Kane (2007), for example, analyzes the D.C. Tuition Assistance Grant Program (DCTAG), which was made available to all residents of the District of Columbia regardless of income, and reports that the number of first-time freshman (dependent) Pell applicants and recipients from D.C. increased between 1999 and 2001, around the time of the 1999 introduction of DCTAG. Having the universe of Pell recipients available to

us, we note that 41 of the 50 states experience increases in the number of Pell recipients over this interval of time, though, which some would consider suggestive that the increases in D.C. are not attributable to the policy alone. Kane (2007) does report that the increases in D.C. “were considerably larger than the changes observed in Maryland or Virginia,” although no claim is made that these differences are statistically significant. He also finds that the take-up rate associated with DCTAG is relatively stable across neighborhood-level income deciles, suggesting that low-income students may not respond differently to the program than those of with less financial need.

Using SAT data and a controlled design, Abraham and Clark (2006) also analyze the DCTAG, reporting that applications and enrollment rates at eligible institutions rise with the DCTAG program relative to a control group. This increase is also shown to be primarily at less-selective institutions, which is consistent with Kane (2007), who reports the enrollment impact was largest at nonselective public four-year colleges and, in particular, at predominantly black institutions.

Georgia’s HOPE Scholarship, which was also more general in application than the initiatives we study here, had eligibility requirements that varied by income. As such, Dynarski (2000) considers the potential influence of aid adjustments on low-income populations by exploiting this variation in eligibility. She finds that college enrollment among those from higher-income families (i.e., above \$50,000) responds to the scholarship while that from lower-income families does not. In a second paper, Dynarski (2004a) also finds that “HOPE had a substantially greater effect on white attendance than black and Hispanic attendance,” which she anticipates, given that these groups tend to have lower incomes. Georgia’s HOPE Scholarship, however, had an explicit merit component, thus confounding any differential responsiveness across income and ability measures. Likewise, studying California’s CalGrant financial-aid program, Kane (2003) exploits variation in college price due to merit-based financial aid, and finds large impacts of grant eligibility, although these findings are most appropriately generalized to high-performing students. In particular, findings may not be easily generalized to the low-income population more broadly.

There are several examples that have looked at the Pell Grant specifically, in part an implicit recognition of the prominence of the Pell Grant in the federal financial aid system. For example, using variation in the Pell-eligibility formula in the late 1980s that increased the generosity of the program for financially independent students, Seftor and Turner (2002) report increased access for non-traditional students. Singell, Waddell, and Curs (2006) use institutional level data to show that the HOPE scholarship increased the relative number of Pell students attending institutions in Georgia, particularly at less-selective four-year institutions.

Last, consider two recent studies that relate closely to ours insofar as they examine the efficacy of an aid initiative that were specifically need-based. That being said, they are implicitly targeted toward the academically able as the policy changes occur at elite institutions. First, Avery et al. (2006) uses administrative and Census data to evaluate the first year of Harvard University’s Financial Aid Initiative (HFAI), which increased aid to low-income students. Using estimated family incomes for plausible Harvard applicants, they find that HFAI attracted a larger and slightly poorer pool of applicants who, once admitted, enrolled at rates similar to the prior year. This is consistent with HFAI being effective at recruiting low-income students out of some otherwise-untapped supply of qualified students

who would not have applied to Harvard in the absence of such aid.² Second, Linsenmeier, Rosen, and Rouse (2006) examine the enrollment effects of changes in financial aid policies (i.e., replacing loans with grants) at an anonymous institution (but clearly selective given its attributes). They report that matriculation by low-income students rose insignificantly in response to the introduction of the no-loan program, but that low-income minority students were more responsive, with matriculation rates rising some 8 to 10 percent. Again, however, these findings are most appropriately generalized to high-performing students.

The prospect of encouraging low-income students matriculating to the top U.S. institutions is an important consideration. However, analyzing the matriculation patterns of low-income students to these elite institutions is arguably to consider a very special case from among the population of low-income students engaging with post-secondary education. In fact, our data indicate that of the 384,386 first-year Pell students attending four-year institutions in 2000, roughly 0.3 percent enrolled at Ivy League institutions. In the same year, 15.5 percent are found at public flagships, some 49 times more. Not surprisingly, Pell students who enroll at public flagships also tend to originate from areas with lower income and lower college-going rates, making for more ideal conditions in which to consider the prospect of whether and how low-income students might respond to targeted aid.³

At the institution level, low-income students are in large measure absent from the elite private institutions, yet it is these institutions that have provided the evidence of aid efficacy thus far, where Pell students typically account for just over one percent of incoming undergraduate classes. For example, the U.S. Department of Education reports that for incoming first-year undergraduates in 2000, Harvard University enrolled 108 Pell students (of 6,660) and Princeton University enrolled 50 (of 4,570). At the same time, the average public flagship enrolled roughly 749 first-time Pell students. In fact, the University of Alaska, Anchorage, while enrolling fewer first-time Pell students than any other state flagship in 2000, enrolled two times as many Pell students (i.e., 128) as did Princeton. With such limitations influencing the external validity of previous studies, we see ample room to contribute to understanding what is happening to low-SES students around income-targeted aid initiatives such as those arising at public flagships.

²Rothstein and Rouse (2007) also studies a need-based aid initiative, although focusing on debt accumulation and employment outcomes as a result of the initiative. Using data for a selective university that adopted a no-loan policy under which the loan component of financial aid awards was replaced with grants, they assess the causal effect of student debt on employment outcomes, reporting that debt affects students' academic decisions during college and can be associated with graduates choosing higher-salary jobs and not choosing poorly paid "public-interest" jobs. Suggestive evidence is also provided that credit constraints and debt aversion interfere with students' abilities to optimize over the life cycle. For example, debt is found to reduce students' donations to the institution in the years after they graduate and increases the likelihood that graduates default on pledges made in their senior years.

³The typical Pell student at a public flagship originates from a ZIP code with per-capita income of \$20,830 and with 16 percent of adult residents having college degrees, while those at Ivy League schools are coming from ZIP codes with \$24,815 and with 18 percent of adult residents having college degrees, on average. (Source: 2000 U.S. Census.)

3 Identification

While there are thirty assistance programs of this type introduced over the 1999 through 2007 period, our identification strategy has us exploit only a subset: those programs initiated at public institutions that limited potential beneficiaries to their own state residents. Namely, we exploit such policies at Minnesota (introduced in 2005), Tennessee (2005), Michigan (2006), Florida (2006), Indiana (2007), Washington (2007), and Illinois (2007). That this would define our strategy is revealed in considering the actual experiment underlying any relevant efficacy test.⁴

In particular, note that treatment is not at the institution level but, rather, at the student level. By extension, where an institution does not limit the set of potential beneficiaries by some exogenous mechanism (e.g., state of residence) there is no well-defined group of comparable students against which one can measure the response to treatment. That is, where no low-income students are excluded from treatment the “control group” is empty.⁵ As none of the private institutions have limited the scope of their initiatives by any exogenous measure (e.g., a student’s state of residence), they, along with four public flagship institutions, will not contribute to our test.⁶ In testing the efficacy of targeted-aid programs, then, our strategy exploits the existence of state borders and the typical practice of public universities to offer income-targeted aid only to residents within their respective states. We then measure the response of treated residents with respect to their own flagship against non-treated non-residents with respect to their flagship institutions.

To illustrate our strategy more specifically, consider a flagship public institution in state i that newly introduces a program of financial assistance to low-income residents in state i who matriculate to the state’s flagship. Low-income residents of this state may or may not respond to the policy (which is what we aim to determine). However, as residents of state $j \neq i$ (i.e., nonresidents from i ’s perspective) are excluded from the institution’s offer, we expect no direct response from j ’s residents.⁷ In fact, residents of state j face the same relative margins in their matriculation decisions before and after the introduction of

⁴Among income-targeted programs, there are four typical allocation mechanisms. Given our sample, and with the exception of the University of Minnesota, we are considering “no-loan” policies that eliminate loans for low-income students. In 2005, Minnesota adopted a “Pell Grant match” policy that match the student’s Federal Pell Grant. The other types of mechanism include, “loan cap” policies that institute a low cap on student loans for low-income students (e.g., Brown), and “no parental contribution” policies that eliminate the parental contribution but retain the student contribution or the standard self-help level (e.g., Yale, Stanford).

⁵As state flagships “treat” all resident students in the state, there is likewise no within-state experiment to be run. In general, one could imagine a regression-discontinuity approach to the research question, using students around the margin of receiving financial support, or, around the margin of admission based on non-financial consideration. See van der Klaauw (2002) for an example of a regression-discontinuity approach to examining the relationship between financial aid and enrollment.

⁶There are four public flagship institutions that also introduce targeted aid but do not limit the potential recipients exogenously to in-state students (i.e., North Carolina adopted in 2003, Virginia in 2004, Michigan State in 2006, and Maryland in 2007). We drop these institutions and their associated students.

⁷Indirectly, they may appear to have responded if, for example, in-state enrollments rise sufficiently to crowd out nonresident students. However, to the extent that nonresidents (i.e., students who do not reside in i) are crowded out by any increase in resident matriculants to i , the bias will be in favor of finding no effect. For example, as students in i respond to i -specific policy and crowd out students from j , we would expect j ’s flagship to gain resident students *relative to* i ’s flagship institution, which attenuates our findings.

targeted aid in state i . Moreover, in the absence of the flagship institution in state j initiating a similar program, we would expect no behavioral response whatsoever from j 's residents, toward their own flagship institution in particular. It is in this way that we adopt low-income resident of all states $j \neq i$ as a control group against which to measure the response of i 's low-income students to the policy change at i 's flagship. We will therefore draw claims of efficacy only to the extent that i -students' patterns of matriculation change with the policy change at i 's flagship institution *relative to* the observed patterns of j -students' matriculation to j 's flagship. The experiments we perform are therefore comparisons of the matriculation patterns of low-income students when their own public flagship institution initiates a change in policy to the matriculation patterns of low-income students in states where no such initiative is adopted.

Before we continue, note that the practice of distinguishing residents from nonresidents in the allocation of financial assistance provides an opportune falsification exercise. Quite simply, if there is a significant response from residents of j (i.e., who are not treated) to the introduction of income-targeted aid at the flagship institution in state i we would be forced to conclude either that there were unobservable time-varying attributes driving matriculation patterns more generally at institution i (and not just of the type for which we expect to see response only from in-state students) or that the experiment was altogether invalid. We therefore perform falsification tests through a comparison of the matriculation patterns of *nonresident* low-income students at flagships that initiate relevant policies to the matriculation patterns of *nonresident* low-income students at other flagships that do not. As nonresident students are not being treated by the policy, there should be no systematically occurring relative difference in the realized matriculation patterns at treated flagship institutions from this group of students.⁸

4 Data

Our primary data source is the Pell "recipient file" held by the offices of the U.S. Department of Education, obtained through our request under the Freedom of Information Act (FOIA). While the dataset includes all Pell recipients over the academic years 1999 through 2007 (at roughly six million observations per year), we use only those students who are recorded as first-time, dependent recipients in their first year at one of the 70 official state flagship institutions. Independent students are dropped from the analysis as they are recognized by the literature as having different attributes and patterns of behavior (e.g., Seftor and Turner (2002)). We also drop student-level observations where we are unable to match them with an institution within the IPEDS dataset using the mapping of Pell-provided institution code

⁸In terms of potential bias in the falsification exercise, note that to the extent residents of i (i.e., part of j 's nonresident population) might forgo an out-of-state alternative when they respond to i 's policy (e.g., they forgo the flagship in state $j \neq i$), the bias will be in favor of finding relative differences in matriculation patterns between treatment and control. For example, as students in i respond to i -specific policy, j 's flagship may well lose nonresident students *relative to* treated flagship institutions, which would yield a positive coefficient on the treatment variable in the nonresident falsification model. We will therefore interpret finding no significant relative difference as secondary evidence of the efficacy of the targeted-aid programs. The use of out-of-state, low-income students could also be problematic if the school itself chooses to substitute in-state, low-income students for out-of-state students.

to IPEDS institution code. We are confident, however, that this problem is primarily one that exists outside of the well-established institutions we analyze here. Moreover, as this type of data issue would keep entire institutions from appearing in our final sample, that we have the entire set of state flagships is consistent with having no missing students at these institutions.⁹

To receive federal aid, a student must first complete a Free Application for Federal Student Aid (FAFSA) form, which provides aid administrators with the information needed to determine the size of an applicants Pell Grant. Related research has relied on indirect measures for the number of low-income students, such as minority enrollments or other student background measures that are correlated with income (e.g., Kane (1994), Dynarski (2004b)). To the contrary, our analysis exploits unique student-level Pell data to directly examine the effects of changes in aid policy on low-income students. Furthermore, the program size ensures that to a large extent our data constitute the most comprehensive representation of poor students attending U.S. higher-educational institutions.

5 Enrollment response

In order to accommodate the policy implementation across multiple time periods, we set up the following difference-in-difference model with a full set of time-period indicators and a policy indicator defined to be unity for institutions and time periods that are subject to the policy.

As a general framework, then, we are interested in the estimate of β in the following time-series/cross-section model:

$$\ln(LIenroll_{it}) = \mathbf{x}_{it}\delta + T_{it}\beta + \alpha_i + \gamma_t + \epsilon_{it}, \quad (1)$$

where $LIenroll_{it}$ captures the low-income enrollment at institution i in academic year t (which we measure as the number of Pell recipients), and T_{it} is the treatment variable, defined to be unity for institutions and time periods (i, t) that are subject to the treatment. The model includes a full set of institution effects, α_i , and a full set of time effects, γ_t . As institution fixed effects will not account for other time-varying factors that influence $LIenroll_{it}$, \mathbf{x}_{it} includes the log-cost of attendance (e.g., tuition, fees) and the state-level log-population of first-year, low-income students enrolled in four-year institutions within each state. Institution-specific errors are captured in ϵ_{it} . An identifying assumption of the above model is that low-income students in treatment and control groups are on similar trends in low-income enrollment before the introduction of targeted aid. We find no evidence of differential trending.¹⁰

⁹We note that the student counts provided in the “recipient files” of the Department of Education are a perfect match with the administrative records of our own institution, the University of Oregon, supporting a certain confidence in the data.

¹⁰The existence of differential pre-treatment trends, where such to exist, would introduce bias into the estimated treatment coefficient. For example, if resident, low-income enrollment was rising faster at treated flagship institutions then this methodology would falsely attribute any continuation of this trend to the targeted-aid programs themselves. Modeling $\ln(LIenroll_{it})$ as cubic in t , we find no significant differences by treatment and control group in pre-treatment trends.

In Column (1) of Table 1 we report the estimated coefficients of Equation (1), which operates as an empirical test of the efficacy of targeted-aid programs as measured by the total enrollment of resident low-income students. With respect to the treatment variable, the point estimate suggests that targeting low-income students has not increased the low-income enrollment levels among the treated. Given the degree of noise in the estimated coefficient, however, we conclude that while targeted-aid policies cannot be credited with attracting needy in-state students to state flagship institutions, there is insufficient evidence to claim that the policies have been detrimental. This is consistent with Linsenmeier et al. (2006), who find no significant overall enrollment effect associated with similar policy changes at an elite institution. We also consider that to the extent that information transmission plays a role in any enrollment response, one might see enrollment effects with a lag. (Although this is inconsistent with the responsiveness of *EFC* and student proximity we document below.)

We report the estimates from the falsification exercise in Column (2) of Table 1, where we expect no significant predictive power in the treatment variable. While this specification reveals significant patterns with respect to costs of attendance and the number of Pell students in the population, as expected, the treatment variable remains insignificant in predicting the low-income enrollment of nonresident students. This is expected given that no such student is treated by the initiated policies.¹¹

– Table 1 approximately here –

While not a falsification test, per se, analyzing the enrollment of low-income students at alternative institutions may well suggest a possible path of substitution. For example, we considered all non-flagship publics in each state and measure the low-income enrollment in a similar pre-post design, defining treatment at the state level as the state’s residents having been treated by the implementation of targeted aid programs by the state’s flagship institution. In so doing, we find that the non-flagship, public institution do lose low-income students when the state’s flagship introduces targeted aid to low-income families. However, while the point estimates (not reported) are larger negatives they are measured imprecisely and are not separable from zero, as in Table 1. Likewise, if we adopt all similarly sized non-flagship publics as the comparator group we find no significant enrollment response.¹²

Table 1 also highlights the importance of residency status in considering the elasticity of enrollment to costs of attendance. In fact, we find no enrollment response with (time-series) variation in costs among resident, low-income students, where non-resident enrollment increases significantly as costs increase, with a 10-percent increase in tuition yielding a 5-student increase in low-income enrollment at the mean enrollment level in the sample (relative to a mean nonresident enrollment of 132 in 1999). The positive elasticity suggests that it may be difficult to separate price effects from the correlation of price and quality (either real or

¹¹Cast differently, one might be inclined to model the (untreated) nonresident enrollment as a kind of difference in difference in difference specification. If the two samples are pooled, with the resident low-income enrollment response to treatment also measured against the “response” on nonresident low-income enrollments, there is the suggestion that enrollment increases with treatment.

¹²Note that this is an indirect test of the efficacy of the initiatives in a way that analyzing changes to expected family income or basins of attraction would not be. For example, while we might expect that any enrollment response at a treated flagship would be to draw low-income students away from other non-flagship publics in the state, there is no reason to believe that such would draw those of different wealth or proximity.

perceived, since time-invariant institutional heterogeneity is absorbed in the error structure of the model). We will return to this point when considering the distance models below where we learn more about the underlying patterns by considering the price effect across the student-proximity distribution. An alternative explanation that is possible where institutions enjoy excess demand from out-of-state students and are in search of higher revenues in that they may admit more nonresident students while raising nonresident tuition. The state-level population of low-income students attending four-year institutions has the expected sign, and magnitudes that seem quite plausible given the selectivity of the institutions within our sample (i.e., elasticities of roughly 1).

6 Family income

As in the enrollment model above, our empirical objective is to capture any changes in expected family contributions (*EFC*) that systematically relate to the timing of “no-loan” policy implementation, using untreated students as “control” observations. With *EFC*, however, the implementation of such analysis is somewhat nontrivial. For example, *EFC* is determined formulaically by the FAFSA and is a function of student attributes (e.g., increasing in family income, number of dependent children in the student’s family). In practice, however, no family is expected to contribute negatively and we only observe the student-specific *EFC* after the application of a non-negativity constraint.

In Table 2, we provide several alternative specifications as a transparent attempt to document that our results are not peculiar to any particular exercise and to reveal to the extent possible the variation that exists in *EFC* around the introduction of these aid initiatives. In Column (1), we first report a linear probability model of the form

$$Prob(EFC_{sit} > 0) = \mathbf{x}_{sit}\delta + T_{sit}\beta + \alpha_i + \gamma_t + \epsilon_{sit}, \quad (2)$$

where EFC_{sit} captures the expected family contribution of resident low-income student s at institution i in academic year t . In (2), T_{sit} is the treatment variable, defined to be unity for students and time periods that are subject to the treatment. Again, as in the enrollment model, we include a full set of institution indicators in α_i , a full set of time effects in γ_t , the log-cost of attendance (e.g., tuition, fees) and the log-population of first-year low-income students enrolled in four-year institutions in \mathbf{x}_{sit} . Student-level errors are captured in ϵ_{sit} . As was the case in Eq. (1), the identifying assumption of Eq. (2) requires that low-income students in treatment and control states trend similarly prior to the introduction of treatment. Finding no evidence of differential trending in the data, we can confirm the validity of this assumption.¹³

The dependent variable is set to unity for all positive *EFC* and is otherwise zero, yielding coefficients on the treatment variable that are intended to capture the marginal change in the propensity for a low-income student at a given institution to have at least some financial assistance expected from the student’s family. The results of the baseline model in Column (1) indicate that there is a significant decrease in the probability that students

¹³Modeling $Prob(EFC_{sit} > 0)$ as cubic in t , we find no significant differences by treatment and control group in pre-treatment trends.

will matriculate with the expectation of financial assistance from family associated with the introduction of targeted aid programs of the type being considered. That is to say, conditional on enrollment, treated low-income students are roughly 2-percent less likely to have positive *EFC*. As specified, however, there is nothing in the linear probability model of Column (1) bounding the projected probabilities to the unit interval. Yet, importantly, institution fixed effects can be estimated without yielding bias.

– Table 2 approximately here –

In Column (2), we report the estimated coefficients of what amounts to the comparable conditional logit model, to bound the projected probabilities to the unit interval. With a rather large dataset, where each group has many observations and there are multiple positive outcomes per group, conditional logit can be a very difficult estimation to perform. We therefore exploit the equivalence of the conditional logistic regression likelihood function and the Cox proportional hazard likelihood function in in Column (2).¹⁴ Doing so again reveals that associated with the introduction of targeted aid programs, there is a significant decrease in the likelihood that students will matriculate with the expectation of family assistance. Conditional on enrollment, this specification suggests that treated low-income students are roughly 2.9-percent less likely to have positive *EFC*.

In Column (3) we report on a probit specification with institution indicator variables. While not yielding the unbiased estimation of parameters as in the Chamberlain procedure of Column (2), this alternative specification yields an estimate of the key parameter of interest that is suggestive of the same empirical regularity in the data. As the treatment coefficient estimate does not substantively change across columns (1) through (3), we are inclined to suggest that neither the unboundedness of the OLS specification nor the potential bias from the incidental parameter problem in the probit specification are cause for concern. Moreover, neither the sign nor economic significance across these alternatives – treated students being roughly 2-to-3 percent likely to have levels of family wealth that yield positive *EFC* – raise concern that our results are sensitive in any meaningful way.

Given the nature of the Pell Program’s formulaic determination of need, *EFC* has a continuous component (which is increasing in family wealth, *ceteris paribus*) that has information of interest in considering the question of efficacy. However, the true expected family contribution is some latent variable, EFC^* , such that $EFC_{sit} = \max(0, EFC_{sit}^*)$. In Column (4), we report the estimated coefficients of a Tobit-equivalent to the specification in Eq. (2). Here again we see evidence that targeted aid has drawn students of relatively lower means, with the estimated coefficient on the treatment variable implying a 25-percent decrease in the underlying family contribution (i.e., EFC^*). We interpret this large increase as partly reflecting the mass at zero *EFC*.

¹⁴For example, the `-clogit-` command in Stata uses an exact calculation to handle multiple positive outcomes. Thus, if there are 5,000 observations per group and 1,000 positive outcomes within a group, `-clogit-` must calculate how many ways 1,000 subjects can be drawn from a pool of 5,000. $\prod_{k=4,001}^{5,000} k$ exceeds that which can be recorded, even in double precision. Our models fail to converge using Stata on a Linux box with 16GB of RAM. In order to achieve convergence, then, we note that `-stcox-` can be used to estimate `-clogit-` models and vice versa. While there is no reason to do so in general, `-stcox-` offers alternative methods by which one can handle the multiple positive outcomes within group. The default method it exploits is the Breslow approximation to the exact calculation, by which we are able to handle multiple positive outcomes in our data. A nice demonstration of equivalency is available at stayconsistent.wordpress.com.

– Table 3 approximately here –

The associated falsification exercises are reported in columns (1) through (4) of Table 3, which shows no such pattern in the data. That is, in each case a comparable specification on nonresident low-income students shows no systematic relationship around the introduction of these aid programs. Holistically, these four specifications and their accompanying falsification tests confirm that targeted aid programs of the nature adopted by flagship public universities lead to more-needy student attending these institutions. Given that the number of needy students does not increase at flagship institutions, which may be the result of explicit or implicit capacity constraints, these findings may be seen as somewhat provocative as they can imply that “improved” access may actually displace some (relatively wealthy) needy students in favor of others.

7 Basins of attraction

Avery et al. (2006), Linsenmeier et al. (2006) and Rothstein and Rouse (2007) suggest that the types of programs under our analysis should permit a better matching of students with institutions. Here, we consider the geographic basins of attraction around each university in our sample of flagships – the area from which the institutions draw low-income students. In particular, we anticipate that to the extent credit constraints on low-income students are systematically relaxed with the introductions of additional aid opportunities, improved matching may well be exhibited in the distances from which students travel to attend the state’s flagship institution.

7.1 Average student proximity by targeted-aid regime

As before, for our analysis of basins of attraction, we set up the following difference-in-difference model with a full set of time-period indicators and a policy indicator defined to be unity for institutions and time periods that are subject to the policy.

As a general framework, we are interested in the estimate of β in the following model:

$$\ln(\text{Distance}_{sit}) = \mathbf{x}_{sit}\delta + T_{sit}\beta + \alpha_i + \gamma_t + \epsilon_{sit}, \quad (3)$$

where Distance_{sit} captures the distance (km) between (resident, low-income) student s ’s home addresses and institution i in academic year t , and T_{it} is the treatment variable, defined to be unity for students and time periods that are subject to the treatment by institution i . As in the enrollment model, the distance models include a full set of institution effects, α_i , a full set of time effects, γ_t , and in \mathbf{x}_{sit} the log-cost of attendance (e.g., tuition, fees) and the state-level log-population of first-year, low-income students enrolled in four-year institutions. Student-level errors are captured in ϵ_{sit} .¹⁵

– Table 4 approximately here –

¹⁵As in earlier specifications, the identifying assumption here is again that low-income students in treatment and control states trend similarly prior to the introduction of treatment. Finding no evidence of differential trending in the data, we can confirm the validity of this assumption.

In Column (1) of Table 4 we report the estimated coefficients of Equation (3) which operates as an empirical test of the efficacy of targeted-aid programs measured in terms of the average proximity of resident low-income students (i.e., the distance between their home addresses and the institution). In short, the data reveal a significant increase in the relative ability of policy-innovating institutions to attract low-income students, with the point estimate suggesting the average resident low-income student travels 4.6-percent farther when treated. Recalling that there are no significant enrollment responses to treatment, these findings again seem to suggest that “improved” access may be displacing some needy students in favor others. In Column (2), we report the estimates from the falsification exercise. As was the case in the earlier falsification attempts, we find no significant predictive power in the treatment variable. This is expected, of course, given that no such student is treated by the initiated policies.

7.2 Distance quantiles by targeted-aid regime

Figure 1 produces the kernel density estimates of low-income student proximity to their associated institutions in 2000. As can be seen in the figure, the distribution of campus proximity is skewed right. This pattern is expected and occurs in all sample years. For example, in 2000, the median proximity was 66 km while the mean proximity was 118 km. In some sense, there is a smallness to the low-income-student market at state flagships, with 25 percent of (first-year, dependent) Pell students at flagship institutions in our sample reporting their official residence (prior to matriculation) within 13.7 km of campus. In this way, it seems a valuable exercise to establish the robustness of the above proximity result to different points along the distribution, for these programs, by their nature, may well experience differential effects on students across the distance range. Thus, as a complement to the analysis of Equation (3), we set up a quantile-regression model with a particular interest in the robustness of the treatment effect to considering different ranges of the proximity distribution.

In Table 5 we largely establish that the proximity effect is realized over a large range of the distribution of proximity. In Figure 2 we report the coefficient estimate and confidence intervals corresponding to the treatment variable derived from equivalent specifications over the entire range of proximity quantiles. Overall, the clear implication is that there are significant and general increases in the basins of attraction where institutions have introduced income-targeted aid and not where aid regime was unchanged. The associated falsification exercises are reported separately in Table 6, which again reveal no predictive power in the treatment variable but for the interval around the 75th percentile.¹⁶

– Table 5 approximately here –

– Table 6 approximately here –

While these relative increases are not present in the matriculation patterns of (untreated) nonresident students, Table 6 does reveal an interesting pattern across the proximity distribution in the underlying behavior of low-income students around changes in *COA*. Note that

¹⁶While proportional measures of treatment response suggest a monotonically decreasing effect throughout the range of distances, we note that the absolute distance implied by the estimates is inverted-U shaped.

the price effect (on distance quartiles) reverses sign over higher distance quantiles. Given that the dependent variable is not enrollment, we are somewhat cautious but nonetheless inclined to interpret this in light of the earlier nonresident enrollment result – *COA* being positively related to low-income nonresident enrollment levels after absorbing time-invariant institution effects into the error structure. Clearly, *COA* positively influences the 10th-quantile distance yet positively influences distance at or above the median. As the confounding of price and quality effects are often suggested as a reason for the positive correlation of price and enrollment, this can be viewed as a challenge to the often-implied prior that students closer to the institution are better informed about the true quality. One interpretation of this quantile-dependent *COA* pattern is that the more geographically distant is a student, the greater are the student’s individual costs to attending and the greater their incentive to determine quality independent of price. (Price then has the effect of decreasing market size.) We thus allow that institution controls may be insufficient to capture time-varying student attributes.

8 Conclusion

We use unique data on all Pell students between 1999 and 2007, that constitutes the most comprehensive representation of poor students entering higher education over this period, to examine the access effects of the introduction of targeted, need-based aid programs available to in-state students at several flagship public institutions. In particular, we use a series of difference-in-difference regressions to examine whether institutions that adopted targeted need-based aid programs for in-state students experienced any systematic increase in the number of (first-time, first-year, dependent) Pell students, the income of matriculating Pell students or their proximity to campus.

Our analysis strongly suggests that targeted-aid programs influence the behavior of needy students, but in a subtle way. Specifically, enrollment models generate negative but insignificant point estimates of enrollment gains from adopting a targeted need-based aid program. At the same time, models of *EFC* and proximity (as measured by the distance between the institution and their official residence prior to matriculation) indicate that institutions that introduce these programs subsequently enroll more-financially-needy and more-geographically-distant students. Quantile regressions also indicate that these aid programs generate a non-linear (i.e., an inverted-U) effect on proximity with maximum effect well below the mean distance. Overall, the data support that targeted need-based aid programs do not change the overall number of needy students, but instead change the profile and composition of students, and toward students who historically may have had less access to these institutions.

As a falsification exercise, we replicate these analyses using nonresident Pell students who are ineligible for the targeted need-based aid. We find no systematic responses from nonresident, low-income students, which supports the assumption that there are no unobservable time-varying attributes driving matriculation patterns at treated institutions more generally.

These findings are important as they suggest that institutions can influence broad access objectives by using need-based aid and that Pell students are responsive to financial incen-

tives. These findings are also somewhat provocative insofar as they suggest that measurable influences are more compositional, in nature. That is, as in the higher-ability population considered in Linsenmeier et al. (2006), we find no significant enrollment effect associated with similar policy changes at flagship public institutions, where low-income students are much more prevalent and more representative of the wider population of needy students. One might consider that active capacity constraints may well imply that “improved” access in a matching sense may actually displace some needy students in favor of others. Collectively, our findings are consistent with Rothstein and Rouse (2007) insofar as increased basins of attraction are suggestive that need-based aid programs reduce credit constraints. Our findings are also consistent with the Avery et al. (2006) hypothesis that these programs permit students to better match within the hierarchy of institutions.

While it is possible that improved access and consequent displacement could work towards a need-blind process where needy students are appropriately matched within the hierarchy of institutions, further work must be done to understand the full general-equilibrium enrollment and access effects of such programs. For example, Linsenmeier et al. (2006), analyzing student responses to one of the earliest adopters of targeted need-based aid in the country, finds statistically significant enrollment gains in low-income minority populations. However, it is unclear whether such programs would have similar efficacy in light of the adoption of similar programs at comparable institutions. Given data such as ours, we are unable to test for differential effects across race. However, future work should consider samples of institutions with competing need-based aid programs in order to tease out the general-equilibrium effects of need-based aid and their implications for low-income and minority populations across the full distribution of ability. Overall, the patterns we identify in *EFC* and proximity around such initiatives suggest that the impact of these programs may well be subtle.

Table 1: **Institution-level enrollment of low-income students around the adoption of income-targeted aid**

The dependent variable is equal to the log of low-income enrollment (i.e., which we measure as the number of Pell recipients) at institution i in academic year t . All specifications include a full set of institution effects and a full set of time effects over annual institution-level observations.

	Resident enrollment	Nonresident enrollment (falsification)
	(1)	(2)
Treatment period	-0.0607 [0.0464]	-0.145 [0.0925]
ln(Cost of attendance)	-0.0578 [0.0722]	0.500*** [0.144]
ln(Resident Pell population)	1.074*** [0.0777]	0.243 [0.155]
ln(Nonresident Pell population)	3.225 [2.903]	-1.463 [5.788]
Constant	-43.73 [37.60]	16.74 [74.97]
Institution controls	Yes	Yes
Year controls	Yes	Yes
Observations / Institutions	630/70	630/70
R^2	0.38	0.06

*** $p < 0.01$, ** $p < 0.05$, * $p < 0.1$. Standard errors in parentheses.

Table 2: **Resident EFC around the adoption of income-targeted aid at flagship public institutions**

In columns (1) through (3), the dependent variable is equal to one if the student's expected family contribution (*EFC*) is positive (i.e., a family contribution is expected to be made). Estimated coefficients are reported in Column (1), estimated logit coefficients are reported in Column (2) and estimated probit marginal effects are reported in Column (3). In Column (4), the dependent variable is equal to student *s*'s *EFC* at institution *i* in academic year *t*, continuous for positive values and left-censored at zero. All specifications include a full set of institution and time indicators.

	Linear Probability (1)	Conditional Logit^a (2)	Probit (3)	Tobit (4)
Treatment period	-0.020*** (0.005)	-0.029** (0.013)	-0.017*** (0.005)	-0.182*** (0.052)
ln(Cost of attendance)	0.017*** (0.003)	0.025*** (0.009)	0.015*** (0.003)	0.183*** (0.036)
ln(Resident Pell population)	0.002 (0.012)	-0.005 (0.031)	0.014 (0.012)	-0.127 (0.123)
ln(Nonresident Pell population)	1.733*** (0.313)	2.638*** (0.872)	1.776*** (0.310)	17.807*** (3.319)
Constant	-21.984*** (4.123)			-225.538*** (43.326)
Institution fixed effects	Yes	Yes	Yes	Yes
Year fixed effects	Yes	Yes	Yes	Yes
Observations / Institutions	420,057/70 $F(12, 419975) =$ 148.5	420,057/70 $\chi^2(12) =$ 510.5	420,057/70 $\chi^2(81) =$ 13397.98	420,057/70 $\chi^2(81) =$ 13869.63

Standard errors in parentheses. *** p<0.01, ** p<0.05, * p<0.1. ^a As discussed in the text, we exploit the equivalency of the Cox proportional hazard likelihood and the conditional logit likelihood in achieving convergence.

Table 3: **Nonresident EFC around the adoption of income-targeted aid at flagship public institutions**

In columns (1) through (3), the dependent variable is equal to one if the student's expected family contribution (*EFC*) is positive (i.e., a family contribution is expected to be made). Estimated coefficients are reported in Column (1), estimated logit coefficients are reported in Column (2) and estimated probit marginal effects are reported in Column (3). In Column (4), the dependent variable is equal to student *s*'s *EFC* at institution *i* in academic year *t*, continuous for positive values and left-censored at zero. All specifications include a full set of institution and time indicators.

	Linear Probability (1)	Conditional Logit^a (2)	Probit (3)	Tobit (4)
Treatment period	-0.003 (0.011)	-0.003 (0.029)	-0.004 (0.011)	-0.057 (0.112)
ln(Cost of attendance)	0.051*** (0.006)	0.073*** (0.018)	0.049*** (0.006)	0.571*** (0.066)
ln(Resident Pell population)	-0.005 (0.023)	-0.015 (0.062)	0.003 (0.024)	-0.265 (0.237)
ln(Nonresident Pell population)	-0.233 (1.091)	-0.448 (2.952)	-0.132 (1.093)	-5.324 (11.146)
Constant	3.227 (14.269)			69.719 (144.432)
Institution fixed effects	Yes	Yes	Yes	Yes
Year fixed effects	Yes	Yes	Yes	Yes
Observations / Institutions	88,224/70 $F(12, 88142) =$ 23.60	88,224/70 $\chi^2(12) =$ 76.92	88,224/70 $\chi^2(81) =$ 2279.41	88,224/70 $\chi^2(81) =$ 2480.57

Standard errors in parentheses. *** $p < 0.01$, ** $p < 0.05$, * $p < 0.1$. ^a As discussed in the text, we exploit the equivalency of the Cox proportional hazard likelihood and the conditional logit likelihood in achieving convergence.

Table 4: **Proximity responses to the adoption of income-targeted aid at flagship public institutions**

The dependent variable is equal to the log of the distance (km) between the student's home address and their institution. All specifications include a full set of institution effects and a full set of time effects, over annual institution-level observations.

	Resident proximity	Nonresident proximity (falsification)
	(1)	(2)
Treatment period	0.060*** (0.015)	0.026 (0.028)
ln(Cost of attendance)	0.075*** (0.010)	-0.122*** (0.016)
ln(Resident Pell population)	0.091*** (0.035)	-0.069 (0.059)
ln(Nonresident Pell population)	-0.908 (0.945)	-3.506 (2.741)
Constant	13.993 (12.331)	52.929 (35.506)
Institution controls	Yes	Yes
Year controls	Yes	Yes
Observations / Institutions	420,056/70	88,224/70
Overall R^2	.06	.00
Mean distance (km)	116.1	967.6

*** p<0.01, ** p<0.05, * p<0.1. Standard errors in parentheses.

Table 5: Resident proximity-response to the adoption of income-targeted aid at flagship public institutions, by proximity quantiles

The dependent variable is equal to the log of the X^{th} distance quantile between home addresses and institutions', where X is given by the column heading. All specifications include a full set of institution effects and a full set of time effects, over annual institution-level observations.					
	10th	25th	Median	75th	90th
	(1)	(2)	(3)	(4)	(5)
Treatment period	0.237*** (0.023)	0.131*** (0.018)	0.064*** (0.010)	0.007* (0.004)	-0.004 (0.003)
ln(Cost of attendance)	-0.175*** (0.019)	-0.017 (0.013)	0.065*** (0.007)	0.015*** (0.003)	0.002 (0.002)
ln(Resident Pell population)	0.118** (0.055)	-0.104** (0.042)	0.049** (0.024)	0.070*** (0.009)	0.027*** (0.007)
ln(Nonresident Pell population)	2.222 (1.455)	-1.693 (1.128)	-0.440 (0.651)	0.521** (0.234)	0.342* (0.194)
Constant	-26.671 (18.995)	25.666* (14.735)	8.895 (8.503)	-2.805 (3.052)	0.492 (2.528)
Institution controls	Yes	Yes	Yes	Yes	Yes
Year controls	Yes	Yes	Yes	Yes	Yes
Observations	420,056	420,056	420,056	420,056	420,056
Pseudo R^2	0.11	0.17	0.20	0.23	0.21
Distance (km)	3.6	14.9	64.5	167.5	292.5

*** p<0.01, ** p<0.05, * p<0.1. Standard errors in parentheses.

Table 6: **Nonresident proximity-response to the adoption of income-targeted aid at flagship public institutions, by proximity quantiles**

The dependent variable is equal to the log of the X^{th} distance quantile between home addresses and institutions', where X is given by the column heading. All specifications include a full set of institution effects and a full set of time effects, over annual institution-level observations.

	10th	25th	Median	75th	90th
	(1)	(2)	(3)	(4)	(5)
Treatment period	0.007 (0.018)	0.023 (0.016)	0.012 (0.017)	0.050*** (0.017)	0.018 (0.020)
ln(Cost of attendance)	0.022** (0.010)	-0.004 (0.009)	-0.032*** (0.010)	-0.033*** (0.010)	-0.071*** (0.011)
ln(Resident Pell population)	-0.002 (0.037)	-0.044 (0.033)	-0.064* (0.035)	-0.049 (0.035)	0.065 (0.040)
ln(Nonresident Pell population)	0.821 (1.713)	-1.219 (1.568)	-2.593 (1.657)	-2.499 (1.629)	2.431 (1.814)
Constant	-6.067 (22.204)	21.648 (20.322)	40.094* (21.474)	39.576* (21.110)	-23.930 (23.511)
Institution controls	Yes	Yes	Yes	Yes	Yes
Year controls	Yes	Yes	Yes	Yes	Yes
Observations	88,224	88,224	88,224	88,224	88,224
Pseudo R^2	0.29	0.26	0.22	0.21	0.18
Distance (km)	66.5	184.8	453.5	1,171.9	2,973.1

*** p<0.01, ** p<0.05, * p<0.1. Standard errors in parentheses.

Figure 1: **The proximity of resident, low-income students**

This figure reproduces kernel density estimates of the (km) distance between each state's flagship institution(s) and the reported residences of first-year Pell students attending their own-state flagship institutions in 2000 as residents.





Figure 2: Estimated treatment effect at distance quantiles, by residency



References

- Abraham, K. G. and M. W. Clark (2006). Financial aid and students' college decisions: Evidence from the District of Columbia Tuition Assistance Grant Program. *Journal of Human Resources* 41(3), 578–610.
- Angrist, J. D. (1993). The effect of veterans benefits on education and earnings. *Industrial and Labor Relations Review* 46(4), 637–652.
- Avery, C., C. Hoxby, C. Jackson, K. Burek, G. Pope, and M. Raman (2006). Cost should be no barrier: An evaluation of the first year of Harvard's Financial Aid Initiative. NBER Working Papers 12029, National Bureau of Economic Research.
- Bound, J. and S. Turner (2002). Going to war and going to college: Did World War II and the GI Bill increase educational attainment for returning veterans? *Journal of Labor Economics* 20, 784–815.
- Cornwell, C., D. B. Mustard, and D. J. Sridhar (2006). The enrollment effects of merit-based financial aid: Evidence from Georgia's HOPE scholarship. *Journal of Labor Economics* 24, 761–786.
- Curs, B. R., L. D. Singell, and G. R. Waddell (2007). The Pell program at thirty years. In J. Smart (Ed.), *Higher Education Handbook of Theory and Research*, Volume XXII. New York: Springer.
- Dynarski, S. M. (2000). Hope for whom? Financial aid for the middle class and its impact on college attendance. *National Tax Journal* 53(3), 629–661.
- Dynarski, S. M. (2003). Does aid matter? Measuring the effect of student aid on college attendance and completion. *American Economic Review* 93(1), 279–288.
- Dynarski, S. M. (2004a). The new merit aid. In C. Hoxby (Ed.), *College Choices: The Economics of Where to Go, When to Go and How to Pay for It*, pp. 63–100. Chicago: University of Chicago Press.
- Dynarski, S. M. (2004b). Who benefits from the education saving incentives? Income, educational expectations and the value of the 529 and Coverdell. *National Tax Journal* 57(2), 359–383.
- Ellwood, D. and T. Kane (2000). Who is getting a college education: Family background and the growing gaps in enrollment. In S. Danziger and J. Waldfogel (Eds.), *Securing the future*. New York: Russell Sage Foundation.
- Hansen, L. (1983). Impact of student financial aid on access. In J. Froomkin (Ed.), *Crisis in Higher Education*, pp. 84–96. New York: Academy of Political Science.
- Havemen, R. and K. Wilson (2007). Access, matriculation, and graduation. In S. Dickert-Conlin and R. Rubenstein (Eds.), *Economic Inequality and Higher Education: Access, Persistence, and Success*, pp. 17–43. Russell Sage Foundation.

- Kane, T. J. (1994). College entry by blacks since 1970: The role of college costs, family background, and the returns to education. *The Journal of Political Economy* 102, 878–911.
- Kane, T. J. (2003). A quasi-experimental estimate of the impact of financial aid on college-going. NBER Working Papers 9703, National Bureau of Economic Research.
- Kane, T. J. (2007). Evaluating the impact of the D.C. Tuition Assistance Grant Program. *Journal of Human Resources* 42(3), 555–582.
- Linsenmeier, D. M., H. S. Rosen, and C. E. Rouse (2006). Financial aid packages and college enrollment decisions: An econometric case study. *Review of Economics and Statistics* 88(1), 126–145.
- Long, B. T. (2004). How do financial aid policies affect colleges? The institutional impact of the Georgia HOPE Scholarship. *Journal of Human Resources* 39, 1045–1066.
- McPherson, M. S. and M. O. Schapiro (1998). *The Student Aid Game*. Princeton University Press.
- Mishel, L., J. Bernstein, and S. Allegretto (2007). *The State of Working America 2006/2007*. Cornell University Press.
- Rothstein, J. and C. Rouse (2007). Constrained after college: Student loans and early career occupational choices. NBER Working Papers 13117, National Bureau of Economic Research.
- Seftor, N. and S. E. Turner (2002). Back to school: Federal student aid policy and adult college enrollment. *Journal of Human Resources* 37(2), 336–352.
- Singell, L. D., G. R. Waddell, and B. R. Curs (2006). Hope for the Pell: The impact of merit based scholarships on needy students. *Southern Economic Journal* 73(1), 79–99.
- Snyder, T. D., A. G. Tan, and C. M. Hoffman (2006). *Digest of Education Statistics, 2005*. Washington, D.C.: National Center for Education Statistics.
- van der Klaauw, W. (2002). Estimating the effect of financial aid offers on college enrollment: A regression-discontinuity approach. *International Economic Review* 43, 1249–1287.