Borrowing Trouble? Human Capital Investment with Opt-In Costs and Implications for the Effectiveness of Grant Aid*

Benjamin M. Marx† and Lesley J. Turner‡

April 2017

Abstract

We estimate the effect of grant aid on City University of New York (CUNY) students' borrowing and attainment using a regression discontinuity/kink design based on the federal Pell Grant formula. Each dollar of grant aid reduces loans by $1.80 among borrowers. We only find crowd-out of this magnitude in colleges that, like CUNY, "offer" no loan aid and require students to opt into borrowing. We develop and empirically support a model that shows opt-in or other fixed borrowing costs can lead grants to crowd out large amounts of loan aid, lowering some students' attainment by reducing their liquid resources.
Keywords: student loans, opt-in costs, human capital investment. JEL: I22, D14, D91, H52.

*We are grateful to the Office of Policy Research at the City University of New York for providing the data used in this study. We also thank Colin Chellman, Simon McDonnell, and Andrew Wallace for sharing their invaluable knowledge of the CUNY administrative data with us, and Brian Cadena, Pamela Giustinelli, Sara Goldrick-Rab, Bruce Kogut, Wojciech Kopczuk, Mike Lovenheim, Paco Martorell, Amalia Miller, Brendan O'Flaherty, Bernard Salanié, Petra Todd, Miguel Urquiola, and seminar participants at the University of Maryland-Baltimore County, George Washington University, University of Michigan, the Congressional Budget Office, University of Virginia, Washington DC economics of education working group, University of Maryland, Harvard University, Michigan State University, Stanford University, University of Illinois Institute of Government and Public Affairs, Federal Reserve Bank of New York, City University of New York, Federal Reserve Board of Governors, the 2013 ASSA annual meetings, the 2013 IRP summer research workshop, the 2014 NBER Spring Economics of Education meeting, and the 2015 AEFP annual meeting for helpful comments and suggestions. We thank Yuci Chen, Megan Gehret, Luke Godwin-Jones, Stephanie Rennane, and Heath Witzen for excellent research assistance.

†Department of Economics, University of Illinois, 214 David Kinley Hall, 1407 W. Gregory, Urbana, Illinois 61801, MC-707. Email: benmarx@illinois.edu.

‡Department of Economics, University of Maryland, 3114 Tydings Hall College Park, MD 20742, NBER, and CESifo. Email: turner@econ.umd.edu.

1 Introduction

Federal and state governments in the United States provide substantial subsidies to college students with the aim of increasing low-income individuals' educational attainment in an era of rising postsecondary costs.[1] However, students' borrowing and educational investment decisions may be affected by the structure and complexity of financial aid programs, limiting the effectiveness of grants and loans. In the spirit of a growing body of research examining how the cognitive biases of older Americans influence the effectiveness of savings incentives, this paper shows how default settings in student loan procedures can reverse the intended effects of grants and lead to under-investment in human capital.

We estimate the impact of need-based grant aid on college students' borrowing decisions and educational attainment using a combined regression discontinuity/regression kink design based on the formula that determines Pell Grant aid as a function of the federal government's measurement of need. We study City University of New York (CUNY) students who are eligible or nearly eligible for a Pell Grant. We find no evidence that Pell Grant aid increases enrollment or attainment on average. In contrast, Pell Grant aid has large, negative, and statistically significant impacts on borrowing. A $1.00 increase in Pell Grant aid reduces first-year students' borrowing by $0.43, implying a greater than $1.80 reduction in loans among students who would have borrowed if they had not received additional Pell Grant aid. This finding of loan crowd-out in excess of 100 percent is at odds with traditional models of human capital investment under credit constraints.

Our findings can be rationalized with a model that allows for a fixed cost of borrowing. Students do not face a monetary fixed cost to borrow. However, CUNY, like many U.S. colleges, provides all students with a financial aid package that includes a $0 loan "offer."[2] As a result, students must actively request loan aid, and this opt-in requirement may impose discrete cognitive or hassle costs on borrowers. In a simple two-period model of students' joint borrowing and schooling choices, a fixed borrowing cost produces a discontinuity in students' budget sets, resulting in a borrowing constraint at the first dollar of debt. The fixed cost generates a threshold level of debt, below which students are unwilling to borrow. A subset of students who would otherwise borrow amounts near this threshold will be induced to stop borrowing by a marginal increase in grant aid, resulting in a greater than one-for-one reduction in borrowing.

Furthermore, the fixed cost can produce heterogeneous impacts of grant aid on attainment. On one hand, grant aid increases the attainment of students constrained by the fixed cost, even though these students do not face a traditional credit constraint. Conversely, additional grant aid may *reduce* educational attainment

[1] During the 2015-16 academic year, the U.S. Department of Education provided $28 billion in Pell Grants and $44 billion in Direct Loans to undergraduate students (see Department of Education's Title IV Program Volume Reports).

[2] Eligibility for federal loans does not depend on the school's "offer". As described in Section 2, a student's dependency status and class standing determines her overall federal loan eligibility, while the composition of loans depends on her unmet need.

of students whose optimal debt is shifted to a level at which the fixed cost binds. For such students, a small increase in grant aid reduces the consumption-smoothing benefit of borrowing to the point at which it is no longer worth paying the fixed cost. Instead, these students reduce schooling to shift consumption to the present. Thus, our model generates ambiguous predictions for the average impact of grant aid on educational attainment in the presence of a fixed borrowing cost. Empirically, we find small, insignificant impacts of Pell Grant aid on the attainment of CUNY students near the Pell Grant eligibility threshold, with confidence intervals excluding average impacts larger than one additional credit earned per $1,000.

We provide reduced form evidence supporting the existence of a fixed cost of borrowing. Borrowing responses to Pell Grant aid occur primarily along the extensive margin. Quantile treatment effects suggest smaller impacts of Pell Grant aid on borrowing at higher quantiles. Within-student regressions of educational attainment on Pell Grant aid and loan take-up provide evidence of the heterogeneous effects predicted by our model. First, among students who do not borrow - a group that includes those constrained by the fixed cost - Pell Grant aid significantly increases credits attempted. Second, after accounting for Pell Grant aid, students who switch from borrowing to not borrowing attempt and earn significantly fewer credits, as the model predicts for students whose borrowing decisions are affected by a binding fixed cost.

We argue that the fixed cost we identify is generated, in large part, by CUNY's default loan offer choice. The majority of students enrolled in community colleges that participate in federal loan programs, as well as all CUNY students, face a default loan offer of zero and must opt into borrowing. In a nationally representative sample of public college students, crowd-out only significantly exceeds 100 percent in schools that employ a zero default loan offer. Among colleges that make nonzero loan offers, crowd-out is significantly less than 100 percent.

Our findings contribute to a growing literature on the importance of defaults and framing in agents' financial decision making.[3] Pallais (2015) shows that an increase in the number of free ACT score reports (equivalent to a cost reduction of $6) substantially increased the quality of college attended by low-income students. Law school students randomly assigned to receive a tuition waiver were significantly more likely to choose public sector careers than students assigned to debt forgiveness with the same present discounted value (Field 2009). Default options could have particular sway in students' borrowing decisions, as Cadena and Keys (2013) provide evidence that even in settings without opt-in costs, undergraduates avoid student loans due to knowledge of their own lack of self-control. Our paper documents the effects of loan offer defaults on students' borrowing decisions and shows how this choice can counteract the ability of grant aid to increase human capital.

[3] Madrian and Shea (2001), Choi et al. (2006), and Chetty et al. (2014) show that default options affect decisions related to investment, saving, and 401(k) participation.

Our paper also contributes to a large literature examining the impact of higher education prices on financially needy students' attainment and the role of complexity in financial aid systems. Grant aid programs that are easily understood and accessed have been shown to increase college-going (Deming and Dynarski 2010). The Pell Grant Program targets students who are especially needy; the average award received by first-year, Pell Grant-eligible students in our sample represented 10 percent of family adjusted gross income (AGI) and 34 percent of the total cost of attendance.[4] However, the complexity of the federal student aid application process may limit the impacts of Pell Grant aid (Dynarski and Scott-Clayton 2006; Dynarski and Scott-Clayton 2008; Bettinger et al. 2012; Scott-Clayton 2013). Pell Grant aid has not been found to increase most low-income students' college enrollment or college quality.[5] Likewise, we find no evidence of enrollment effects, with an additional $1,000 in Pell Grant aid resulting in an insignificant 0.3 percentage point (2 percent) increase in the probability that an applicant enrolls in a given CUNY institution.

Since we can rule out all but small impacts of Pell Grant aid on enrollment, we are able to estimate the impact of need-based grant aid on the attainment of students who have already made the decision to attend college. Despite increases in low-income youths' college entry over the past decades, college completion rates in this population remain low (Bound et al. 2010), underscoring the importance of determining when need-based aid can increase *enrolled* students' attainment. Only a handful of studies address this question. Using data on public college students in Ohio, Bettinger (2004) estimates a positive correlation between a student's Pell Grant aid and subsequent reenrollment, with an additional $1,000 correlated with a 3 to 4 percentage point increase in year-to-year persistence.[6] Goldrick-Rab et al. (2012) show that Pell Grant recipients randomly assigned to receive an additional $1,000 in grant aid experienced a 2 to 4 percentage point increase in persistence, but only when additional aid did not displace federal loans. In contrast, Denning (2016) finds positive effects of grant aid when combined with higher federal borrowing limits. Our results suggest that default settings in student loan programs have important implications for the effectiveness of public higher education expenditures and the ability of need-based grant aid to increase attainment among low-income students who choose to go to college.

2 The CUNY System and Need-Based Student Aid

Our main analyses focus on students enrolled in the CUNY System. CUNY is the largest urban public university system in the country, encompassing 17 two- and four-year colleges that serve over 250,000 under-

[4]Nationwide, the average Pell Grant represented 17 percent of average family income in 2012 (authors' calculations using Tables 1 and 2-A in U.S. Department of Education 2013).

[5]See, for example, Kane (1995), Rubin (2011), Turner (2014), and Carruthers and Welch (2015). In contrast, Seftor and Turner (2002) find that the Pell Grant Program's introduction increased enrollment of some non-traditional, older students.

[6]These estimates are not robust to controlling for institution fixed-effects and interactions between Pell Grant aid and borrowing are not examined.

graduate students annually. CUNY institutions have low tuition, limited institutional grant aid, and operate in a state with generous need-based grant aid. Similar to those attending other urban public institutions, CUNY students have low retention and graduation rates. Among first-time fall 2006 freshmen, only 10 percent of associate degree seeking students graduated in three years and only 41 percent of bachelor's degree seeking students graduated within six years. Appendix A provides additional details on the CUNY system, including tuition, fees, and estimated living expenses.

A centralized application system determines eligibility for federal student aid. Current and prospective students must submit a Free Application for Federal Student Aid (FAFSA) to the U.S. Department of Education every year. FAFSA inputs include a detailed set of financial and demographic variables, such as income, untaxed benefits, assets, family size and structure, and number of siblings in college. The federal government calculates each student's expected family contribution (EFC) via a complicated, non-linear function of these inputs. Eligibility for Pell Grant aid is determined by a student's EFC and cost of attendance (COA), which includes tuition, fees, and estimated living expenses. Specifically, students with EFC below a set annual threshold are eligible for the minimum Pell Grant. Every $1 decrease in EFC leads to a $1 increase in (statutory) Pell Grant aid, up to the maximum Pell Grant. Only students with zero EFC receive the maximum award.[7]

After submitting a FAFSA, New York State residents are prompted to complete the online Tuition Assistance Program (TAP) application. TAP grants are determined by New York State taxable income. TAP grant aid is not awarded based on EFC and TAP provides grants to students much higher in the income distribution - up to $80,000 in New York State taxable income for dependent students. In the years we examine, the maximum TAP award was $5,000.

According to the Integrated Postsecondary Education Data System (IPEDS), 68 percent of first-time, full-time, fall 2012 entering undergraduate CUNY students received a Pell Grant, while less than 3 percent received institutional grant aid. Thus, CUNY institutions' ability to adjust institutional aid in response to Pell Grants, as illustrated by Turner (2014), will be limited. CUNY institutions also have limited federal work-study funds: fewer than 9,000 students participated in the federal work-study in 2012, implying that less than 4 percent of CUNY undergraduates receive work-study aid.

Finally, CUNY students are eligible to borrow through the federal Direct Loan Program. The terms of federal loan aid depend on a student's course load, tenure, and unmet need. Specifically, a student's unmet need, equal to COA minus EFC and grant aid, determines her eligibility for subsidized loans. First-year students are eligible for subsidized loans equal to the lesser of unmet need and $3,500. Dependent first-year students can borrow an additional $2,000 in unsubsidized loans while independent students can borrow an

[7] Appendix A contains annual minimum and maximum Pell Grant awards.

additional \$6,000. All students are eligible for unsubsidized loans and even students that do not qualify for subsidized loans can still borrow unsubsidized loans up to the overall maximum (\$5,500 for first-year dependent students and \$9,500 for first-year independent students). Subsidized loans do not accrue interest until six months after a student leaves school. Appendix A provides additional details on federal student loan programs, including interest rates. Most CUNY students are eligible for the maximum subsidized loan.[8]

Prospective students can list up to 10 schools that they are considering attending on their FAFSA. Because the FAFSA requires information on prior-year taxable income, prospective students generally wait to complete the FAFSA until after their family has filed their tax return (at best, early February). Most prospective CUNY students apply to college in advance of completing a FAFSA via the CUNY System's centralized application. Even though CUNY applicants are admitted on a rolling basis, applications submitted after February 1st are not guaranteed consideration. Students are notified of their EFC by the Department of Education shortly after submitting a FAFSA but do not learn of their federal aid eligibility until after they have been admitted to a college, generally six to eight weeks after submitting the CUNY application. Upon admission, each college provides a student with her financial aid package, which specifies grant aid (federal, state, and institutional) and nonbinding loan "offers." During the months leading up to the fall semester, the student decides whether to accept the admissions offer and how much (if any) federal loan debt to incur.

Colleges participating in federal loan programs must allow students to borrow up to their full federal loan entitlement, but they have discretion over the presentation of loans offers (Scott-Clayton 2013). Almost all four-year colleges list the maximum amount of federal loan aid students can borrow in their award packages, while CUNY and a substantial number of community colleges offer a loan amount of zero to all students. Students in these institutions must submit a separate application in which they specify their desired amount of federal loan aid.[9] This practice may impose fixed costs in the form of both psychic costs of deviating from the zero loan default and transaction costs associated with applying for loan aid.

Most community colleges have a zero default loan offer. We gathered information on loan procedures for all community colleges that participate in federal loan programs.[10] Of the 792 non-CUNY institutions (92 percent) for which we were able to ascertain packaging practices, 449 follow CUNY's practice of a zero loan default option (Appendix Table B.1). These schools contained 51 percent of undergraduates attending non-CUNY community colleges that participated in federal loan programs. An additional 19 institutions (containing 4 percent of students) only offer subsidized loans. Community colleges that do not offer loans

[8] Less than 1 percent of CUNY undergraduate students received parent PLUS loans and no CUNY students used private loans during the period we examine.

[9] Appendix Figure B.1 displays a sample CUNY financial aid award letter and Appendix Figure B.2 displays a sample loan application from Hunter College.

[10] Data on enrollment and federal student loan program participation was drawn from Cochrane and Szabo-Kubitz (2014) and http://projectonstudentdebt.org/files/pub/CC_participation_status_2013-14.pdf.

have substantially lower borrowing rates (16 percent) than schools with a nonzero default (29 percent).

3 Data and Sample

Our primary sample includes first-time, degree seeking freshmen who entered the CUNY system in the fall of the 2004-05 through 2010-11 (hereafter 2005 through 2011) academic years. We observe students in their first three years after entry and distinguish between first-year and returning (second- and third-year) students. We only observe students' FAFSA information in the 2007 through 2011 academic years, and therefore the two earliest entry cohorts only contribute to our sample of returning students. We restrict our sample to US citizens or permanent residents who completed a FAFSA and received an EFC within $4,000 of the threshold for Pell Grant eligibility.[11]

Table 1 displays the characteristics of first-year students by Pell Grant eligibility. Our focus on students with an EFC near the Pell Grant eligibility threshold yields a sample with average family AGI of approximately $48,000. Pell Grant eligible students receive more TAP and other grant aid (a category that includes aid from smaller state and federal grant programs, as well as institutional aid) than ineligible students, while ineligible students take on greater debt. CUNY students borrow at low rates; only 12 percent of the sample takes on any debt in their first year, despite having substantial need and subsidized loan eligibility. Only 2 percent of first-year borrowers fully exhaust their federal loan eligibility. Pell-ineligible students are more likely to borrow, with 24 percent taking on some debt, compared to 7 percent of ineligible students. Finally, Pell Grant eligible students have different demographic characteristics than ineligible students - they are more likely to be nonwhite, have lower SAT scores, and are less likely to have a college educated parent. These differences in observable characteristics between Pell Grant eligible and ineligible students motivate our use of RD and RK designs to identify the causal impact of grant aid on student outcomes.

3.1 Are CUNY students representative of the national population?

To determine whether our sample of CUNY students resembles Pell-eligible and near-eligible students nationwide, we compare CUNY students who first entered college in fall 2010 to the nationally representative sample of first-year, degree seeking public school students in the 2012 National Postsecondary Student Aid Study (NPSAS), who first entered college in fall 2011. We limit both samples to students with expected family contributions within a $4,000 window around the Pell Grant eligibility threshold. We separate stu-

[11] Noncitizens that are not permanent residents are ineligible for most federal and state grant aid and make up only 6 percent of students in the cohorts we examine. Approximately 12 percent citizens/permanent residents do not complete a FAFSA. Among FAFSA filers, 73 percent have EFCs falling outside the $4,000 window and are excluded. The majority of students dropped by this restriction have an EFC of zero.

dents by initial degree program (associate versus bachelor's) and institutional loan packaging procedures (i.e., whether a student attends a school with default loan offer of zero). Appendix Table B.2 displays the demographic characteristics, cost of attendance, and financial aid received by students in each group.

After taking into account EFC and federal, state, and institutional grant aid, CUNY students have unmet need that is comparable to their peers nationwide. Among associate degree seeking students, mean unmet need at CUNY ($6,000) is just above national averages ($4,700 for nonpackaging NPSAS institutions and $5,600 for schools that package loans). Among bachelor's degree seeking CUNY students, mean unmet need ($6,200) falls between that of students enrolled in schools that package loans ($11,000) and those that do not package loans ($4,900).

Despite relatively similar levels of unmet need, CUNY students borrow at lower rates than their NPSAS counterparts, with 14 percent of associate degree seeking students and 18 percent of bachelor's degree seeking students taking up loans. Outside of the CUNY system, schools' packaging practices are predictive of loan take-up. Borrowing rates for associate degree seeking students are around 26 percent in schools with zero loan defaults compared 50 percent in schools with nonzero defaults. The difference is even larger for bachelor's degree seeking students (28 versus 76 percent, respectively). These patterns suggest that the low borrowing rates in the CUNY system are at least partially driven by loan opt-in costs. We formally test whether the impact of Pell Grant aid on borrowing significantly varies with schools' loan packaging practices using micro data from the NPSAS in Section 5.5.

In terms of their demographic characteristics, CUNY associate degree seeking students are younger and more likely to be classified as dependent students. Both associate and bachelor's degree seeking CUNY students are more likely to be Black or Hispanic, less likely to be White, and less likely to have parents who attended college, although these differences are smaller when NPSAS students attending nonpackaging schools are used as the comparison group. Finally, CUNY students are more likely to be first- or second-generation immigrants, reflecting the fact that the majority of CUNY students matriculate from the New York City public school system. Because of these differences in demographic characteristics and the fact that groups that are overrepresented in the CUNY system may be less willing to take on loans, we show that our main estimates do not vary significantly with race, dependency status, parental education, and immigrant status in Section 5.5.

4 Conceptual Framework

This section describes how we generate causal estimates of the impact of Pell Grant aid on students' borrowing and attainment outcomes using Pell Grant eligibility as an instrument for grant aid and provide evidence that CUNY students face a fixed borrowing cost. We start by presenting a standard lifecycle model, which generates the prediction that crowd-out of loans in response to increases in grant aid will be bounded from above by 100 percent. In Section 4.1, we present a generalized decision model that allows for a fixed cost of borrowing and credit constraints. When borrowing entails a fixed cost, crowd-out of loans from grants can exceed 100 percent, a testable distinction from the standard lifecycle model. In Section 4.2, we describe sufficient conditions for identifying the intensive-margin response of borrowing and testing for crowd-out exceeding 100 percent among specific student subgroups using estimated treatment effects and population moments. Finally, Section 4.3 outlines the empirical framework that we use to estimate treatment effects and test predictions of the fixed borrowing cost model.

Let $Y = y(G, X, U)$ represent the relationship between outcome Y and grant aid $G = g(Z, X, V)$, student characteristics X, and an error term U, when there is an instrument Z that is conditionally independent of U and V. Let $G^z = g(z, X, V)$ and $Y^z = y(G^z, X, U)$, where z is a realized value of the instrument Z. The instrument may be discrete or continuous in this general formulation. We first focus our discussion on the simplest case of a binary instrument, $Z \in \{0, 1\}$, and then turn to the case of a continuous instrument. In our empirical work, we use both types of instruments and obtain similar results from each individually or combined. Finally, let observed debt B^z correspond to actual debt b^* when it exceeds zero: $B^z = \mathbf{1}[b^* > 0]\, b^*(G^z, X, U)$. Actual debt may be positive or negative, with the latter case representing saving, which is not observed by the econometrician.

In a basic two-period life cycle model of a student's educational investment decisions, variation in grants will reduce a student's optimal debt at a rate between zero and one hundred percent. Suppose an enrolled student maximizes utility $U = u(c_0) + \beta u(c_1)$. Subscripts indicate the period, $\beta \in (0, 1)$ is the discount factor, and $u(\cdot)$ follows standard assumptions for instantaneous utility (strictly increasing, strictly concave, and twice continuously differentiable). Consumption while in school is equal to grants g and endogenously chosen amount borrowed $b \in \mathbb{R}$ plus resources net of fixed schooling costs $\omega \in \mathbb{R}$, i.e. $c_0 = g + b + \omega$. Consumption after school equals earnings W minus the repayment of debt at gross interest rate R: $c_1 = W - Rb$. The first-order condition for optimal debt b^* is $0 = u'(g + b^* + \omega) - Ru'(W - Rb^*)$, which, when differentiated, yields:

$$\frac{db^*}{dg} = -\frac{u''(c_0)}{u''(c_0) + R^2 u''(c_1)} \in (-1, 0). \tag{1}$$

Thus, an exogenous increase in grants will never lead to a greater than 100 percent reduction in the amount borrowed.

A weak test of the basic model is whether $\frac{1}{N} \sum\limits_{i=1}^{N} \frac{db^*}{dg} \in (-1, 0)$. This bounding of average borrowing would be a weak test because equation 1 could hold on average but not for all individuals. A stronger test can be constructed by noting that, under the assumption that the instrument weakly increases grant aid for all students, we can classify students into three mutually-exclusive categories: "always-borrowers" ($A = \{i : B^0 \geq B^1 > 0\}$), "never-borrowers" ($N = \{i : B^0 = B^1 = 0\}$), and "switchers" ($SW = \{i : B^0 > B^1 = 0\}$). This labeling of population groups by their *potential* outcomes is an example of a "principal stratification" (e.g., Frangakis and Rubin 2002).[12] Despite their definition in terms of outcome variables, the strata are not affected by treatment assignment, and hence within-strata comparisons of potential outcomes Y^1 and Y^0 will represent a causal effect.

4.1 Decision model with a fixed cost of borrowing

We model the human capital investment decisions of a student who has enrolled in college, providing evidence later that the decision to enroll is not affected by our quasiexperimental variation in grant aid. The decision model incorporates key features of federal student aid programs in the United States but is simplified so as to highlight the novel empirical implications of a fixed borrowing cost. Appendix C contains proofs. As in the simple life cycle model described above, an individual lives for two periods. She chooses schooling investment s (an index that could include classes taken and effort) and debt b in the first period to maximize lifetime utility, $U = u(c_0) + \beta u(c_1)$.

In the first period, the student has resources equal to her expected family contribution EFC and exogenous income ω, where ω represents the error term in the federal government's estimation of family resources, and can be positive or negative. The student faces costs $T(s|EFC)$ associated with educational investment, which encompass both direct costs $T_d(s|EFC)$ (e.g., tuition and fees) and opportunity costs $T_i(s|EFC)$ (e.g., foregone earnings). $T(s|EFC)$ is twice continuously differentiable, with $T'_d(s|EFC) \geq 0$, $T'(s|EFC) > 0$ and $T''(s|EFC) \geq 0$. While costs may depend on the student's EFC, this fact does not affect the analysis, and we hereafter suppress this notation. The student receives grants $g(s|EFC) = g + h(s|EFC)$ from the government, where the constant component g does not depend on the schooling investment.[13] We assume $h(s|EFC) = 0$ to highlight the effects of variation in g and for expositional clarity and because the income effect appears most relevant for grants like the Pell Grant that are calculated and awarded after an individual has already taken many steps towards enrolling in college. In Appendix C, we present the general model and show that price effects do not alter the model's empirical predictions.

Students choose to borrow and invest in schooling to increase future earnings. In the second period, the student receives earnings $W(s)$ where $W'(s) > 0$ and $W''(s) \leq 0$.[14] Borrowing is subject to multiple interest

[12] We thank an anonymous referee for pointing out this connection.

[13] In our setting, grant aid does not depend on schooling for students who attempt at least 12 credits per semester.

[14] Our model would yield similar predictions if we allowed for heterogeneous costs of schooling investment by letting s enter

rates and potential constraints. The gross market interest rate is $R_m < \frac{1}{\beta}$, but the government subsidizes some student loans by charging the rate $R_{sub} < R_m$.[15] The student receives the subsidized interest rate on all loans up to a limiting amount $b_{sub}^{max} = \min\left\{\bar{b}, T_d\left(s\right) - g - EFC\right\}$, where $\bar{b}$ is a constant. This formulation captures the structure of federal subsidized loans, which can be used to cover unmet need, represented by $T_d\left(s\right) - g - EFC$, up to a fixed limit $\bar{b}$. Additionally, the student can borrow up to the overall federal loan limit $\bar{\bar{b}} > b_{sub}^{max}$.

The student pays a fixed borrowing cost γ if she chooses $b > 0$, which represents discrete monetary, time, and psychic costs of incurring debt. For notational convenience, we define indicator functions $\kappa_0 = \mathbf{1}\left\{b > 0\right\}$ (incurring positive debt), $\kappa_{sub} = \mathbf{1}\left\{b > b_{sub}^{max}\right\}$ (incurring positive unsubsidized debt), and $\xi = \mathbf{1}\left\{T_d\left(s\right) - g - EFC < \bar{b}\right\} = \mathbf{1}\left\{b_{sub}^{max} = T_d\left(s\right) - g - EFC\right\}$ (being bound by the endogenous subsidized borrowing limit) to distinguish between cases.[16]

The student faces budget constraints $c_0 \leq EFC + \omega + g + b - T\left(s\right) - \gamma \cdot \kappa_0$ in the first period and $c_1 \leq W\left(s\right) - R_{sub}b - \kappa_{sub}\left(R_m - R_{sub}\right)\left(b - \bar{b} - \xi\left(T_d\left(s\right) - g - EFC - \bar{b}\right)\right)$ in the second period.[17] Leaving the boundary conditions $0 \leq b \leq \bar{\bar{b}}$ implicit, the student solves:

$$\max_{s,d}\Big\{u\left(\omega + EFC + g + b - T\left(s\right) - \gamma\kappa_0\right) +$$
$$\beta u\left(W\left(s\right) - R_{sub}b - \kappa_{sub}\left(R_m - R_{sub}\right)\left(b - \bar{b} - \xi\left(T_d\left(s\right) - g - EFC - \bar{b}\right)\right)\right)\Big\}$$

Optimal schooling s^* and debt b^* will satisfy two sets of conditions, one for each choice variable:

$$\begin{aligned} \mu_b := \quad & u'\left(c_0\right) - \beta\left(R_{sub} + \kappa_{sub}\left(R_m - R_{sub}\right)\right)u'\left(c_1\right) - \lambda \geq 0, \\ & \mu_b b^*\left(b_{sub}^{max} - b^*\right)\left(\bar{\bar{b}} - b^*\right) = 0 \end{aligned} \tag{2}$$

$$\begin{aligned} \mu_s := \quad & -T'\left(s\right)u'\left(c_0\right) + \beta\left(W'\left(s\right) - \xi\kappa_{sub}\left(R_m - R_{sub}\right)\left(T_d'\left(s\right)\right)\right)u'\left(c_1\right) \geq 0, \\ & \mu_s\xi\kappa_{sub} = 0 \end{aligned} \tag{3}$$

directly into the period utility functions, as in Cameron and Taber (2004), or by letting ability vary across students, as in Lochner and Monge-Naranjo (2011).

[15] If students could earn $R_m > R_{sub}$ on their savings then they could theoretically engage in arbitrage for the full subsidized loan amount. However, in the years we examine, market interest rates were low and students faced a 1 to 3 percent origination fee on all loans, resulting in $R_{sub} \approx R_m$. In almost every year, the real rate of return to investing a subsidized loan in a 12-month certificate of deposit (CD) was negative. The most a student could earn by investing the maximum subsidized loan in a 60-month CD in any year we examine would be \$221 (see Appendix Tables A.3 and A.4). This is in contrast to the setting Cadena and Keys (2013) examine when investigating potential explanations for students' decisions to forgo subsidized loans. Omitting a third term representing the market rate for savings does not affect the predictions of our model.

[16] Psychic costs could enter utility directly (e.g., as a separate argument added to the concave function of consumption). Because this does not change the general form of the solution, we present the notationally-convenient case with all fixed costs entering the budget constraint.

[17] We assume the regularity condition $W''\left(s\right) \leq -R_m T_d''\left(s\right)$ for all s to ensure global concavity of the problem. We deem this condition reasonable because direct costs are linear or concave in observable schooling investment, depending on a student's course load: tuition is linear in credits attempted for part-time students, while full-time students (attempting 12 to 18 credits) are charged a flat rate. Appendix C shows that a weaker condition would suffice.

The first set of conditions show that optimal debt will equate a student's current and future marginal utilities unless b^* falls at one of the threshold values of zero (where the fixed cost binds), the subsidized loan maximum (where there the interest rate jumps), or at the overall borrowing limit. The second set of conditions show that optimal schooling will equate the current and future marginal utilities of schooling given the choice of debt as long as an additional dollar of debt would not also raise the interest rate on some borrowing. An example that satisfies all of these conditions is a student whose remaining need is greater than the subsidized loan limit ($\xi = 0$) and whose optimal borrowing is not constrained at either zero or the maximum total loan amount. For this student, $\mu_b = \mu_s = 0$ and $T'(s) = (R_{sub} + \kappa_{sub}(R_m - R_{sub}))^{-1} W'(s^*)$. In such cases, s^* equates the present discounted marginal costs and benefits of schooling and will not depend on income in either period. In such cases, schooling does not respond to a marginal increase in grant aid. Students who do not face borrowing constraints will not increase their schooling in response to unconditional increases in grant aid, just as wealth does not affect the human capital investment of unconstrained individuals in the model of Lochner and Monge-Naranjo (2011).

For a given level of EFC, students can be ordered in terms of additional resources ω. A partition of this spectrum defines the different cases a student may fall into, labeled groups A through F. Table 2 summarizes students' initial choices of debt and responses to grant aid in each case. Group A is made up of students with resources sufficiently high that they choose to save (i.e., $b^* < 0$). Group F describes students who have so few resources that they would prefer to borrow more than the maximum allowable government loan $\bar{\bar{b}}$ but cannot. For groups between these extreme cases, the optimal level of debt is weakly decreasing in resources. As long as the fixed cost of borrowing $\gamma > 0$, there will be some minimum level of debt that students are unwilling to take on, which we denote as $\underline{b}$.

Though we distinguish six distinct groups of students, they fall into two general types: those choosing corner solutions for debt ("threshold borrowers") and those choosing interior solutions. Groups A, C, and E choose interior levels of debt, and, as a result, the amount they borrow responds to changes in grant aid. Grant aid does not increase the educational attainment of these students. Conversely, threshold borrowers arrive at a corner solution for borrowing due to the presence of the fixed cost (Group B), kinks in the interest rate schedule (Group D), or credit constraints (Group F), and thus respond to additional grant aid by increasing schooling.

Panel A of Figure 1 displays the budget set, borrowing, and consumption choices of Groups B, C, and D. Group A students (not shown) locate to the left of the discontinuity in the budget set caused by the fixed cost, while Group B students arrive at a corner solution and neither borrow nor save. Likewise, Group D members borrow at the subsidized maximum, arriving at a corner solution caused by the interest rate kink, and Group F members (not shown) borrow at the federal maximum, represented by the discontinuity in the far right portion of the budget constraint. Students in Group C locate between the fixed cost discontinuity and the interest rate kink, while those in Group E (not shown) locate between the interest rate kink and discontinuity due to the overall borrowing limit.

Panels B through D of Figure 1 illustrate the impact of grant aid on students' borrowing decisions. Because grant aid weakly reduces desired loans, students in Groups A or B are never-borrowers. Students in Group B remain at the threshold caused by the fixed cost following a marginal increase in grant aid and complete more schooling in order to raise the ratio of future income to current income (Panel B). Members of Groups D and F are also constrained at a positive amount of student loans and have a similar response to grant aid.[18] Students in Groups C through F are always-borrowers and continue to borrow after a marginal increase in grant aid.

Responses within and between the groups are all continuous except for students switching between Groups B and C (Figure 1 , Panel C). Students with b^* close to $\underline{b}$ may be induced to switch to $b = 0$ by a small increase in grant aid, Δg. These switchers do not find it worthwhile to pay the fixed cost of borrowing and instead smooth consumption between periods by reducing schooling. An example of this counterintuitive result is the student with $1,000 in unmet need who chooses to work more while in school to avoid borrowing, even if it means she takes fewer classes. Thus, in terms of the general framework, increasing grant aid weakly increases schooling and weakly decreases debt among always-borrowers and never-borrowers, while switchers decrease schooling and may reduce borrowing by more than the amount of the grant aid increase.

4.2 Identifying intensive-margin borrowing responses

A fixed cost of borrowing can generate a borrowing response to additional grant aid that exceeds 100 percent, a result at odds with the simple life cycle model's prediction that $1 of grant aid will reduce loans by, at most, $1. However, unconditional average effects of grant aid on borrowing can still be small when a large number of students never borrow, even if some students' intensive-margin borrowing responses to grant aid are large. Thus, in this subsection, we highlight sufficient conditions for identifying the intensive-margin response of borrowing and testing for crowd-out exceeding 100 percent among specific student subgroups.

To identify intensive-margin borrowing responses, we assume two monotonicity conditions: a first-stage condition, $G^1 \geq G^0$, and a second-stage condition, $B^1 \leq B^0$. The first condition is the standard monotonicity assumption for identifying local average treatment effects via two-stage least squares (Imbens and Angrist 1994). This assumption requires that the instrument weakly increases grants for each member of

[18]Students remain at their respective borrowing thresholds by keeping debt constant, except in the case of the Group D students with unmet need below the exogenous limit on subsidized loans ($\xi = 1$). Grants reduce these students' unmet need and consequently their subsidized borrowing limit. As a result, these students adjust loans so as to remain at the interest rate kink but otherwise behave like other threshold borrowers, increasing schooling as grant aid rises.

the population, which is consistent with our empirical setting. The second monotonicity condition restricts borrowing responses to the set of students who would borrow even when grant aid is low, thereby increasing the minimum effect per responding student that would be necessary to explain the observed average effect.[19] In the model described in Section 4.1, there are no students who increase borrowing when grant aid increases, which satisfies the second monotonicity condition.

Let $D^z = \mathbf{1}[B^z > 0]$ represent the extensive margin of borrowing when the instrument $Z = z$. When the instrument is binary and the above monotonicity conditions hold, there will be three groups: always-borrowers ($A = \{i : D^1 = D^0 = 1\}$), never-borrowers ($N = \{i : D^1 = D^0 = 0\}$), and switchers ($SW = \{i : D^1 = 0, D^0 = 1\}$). We label the population shares of these groups π_A, π_N, and π_{SW}, respectively.

Consider a population-wide change in the value of the instrument from $Z = 0$ to $Z = 1$. The average effect on borrowing is:

$$\begin{gathered} E\left[B^1 - B^0\right] = \\ \pi_A E\left[B^1 - B^0 | A\right] + \pi_N E\left[B^1 - B^0 | N\right] + \pi_{SW} E\left[B^1 - B^0 | SW\right] = \\ \pi_A E\left[B^1 - B^0 | A\right] + \pi_{SW} E\left[-B^0 | SW\right] \end{gathered} \tag{4}$$

The last expression is simplified by the fact that borrowing is always zero for never-borrowers. Equation (4) allows us to define "loan crowd-out among would-be borrowers" as:

$$\frac{E\left[B^1 - B^0\right] / E\left[G^1 - G^0\right]}{\pi_A + \pi_{SW}} = \frac{\pi_A}{\pi_A + \pi_{SW}} \frac{E\left[B^1 - B^0 | A\right]}{E\left[G^1 - G^0\right]} + \frac{\pi_{SW}}{\pi_A + \pi_{SW}} \frac{E\left[-B^0 | SW\right]}{E\left[G^1 - G^0\right]} \tag{5}$$

Equation (5) is a function of the responses of the students whose borrowing can respond to the instrument-induced change in grant aid – those who would have borrowed if not for the additional grant aid (i.e., always-borrowers and switchers) – and it gives the per-grant-dollar reduction in loans among these students. Under the further assumption that the instrument induces equal changes in the grant amount among each type of student, a reasonable assumption in our setting, the expression is precisely equal to the weighted average responses by always-borrowers and switchers. A finding that crowd-out among would-be borrowers is greater than one in absolute value would imply that this is true for at least one of the two types of would-be borrowers, a violation of the prediction of the standard life cycle model that no student would reduce borrowing by more than the amount of the grant increase.

To empirically estimate loan crowd-out for would-be borrowers, estimates of $E\left[B^1 - B^0\right]$, $E\left[G^1 - G^0\right]$, π_A, and π_{SW} are required. With these four quantities, we can point identify the weighted average of the

[19] The monotonicity conditions can hold in either direction as long as the direction is the same for all students. We have written the first-stage condition in the direction expected given our instrument and the second in the direction predicted by theory when borrowing is the outcome of interest.

responses among two groups, the always-borrowers and switchers. We therefore set-identify the possible combinations of intensive-margin responses of each of the two groups; if the average effect for the two groups exceeds -1 in magnitude then it must be that the effect exceeds -1 for at least one of the groups. Identification relies on both monotonicity conditions because if either condition fails, then any change in borrowing could be explained by asymmetric borrowing responses to bi-directional changes in grants (failure of the first stage monotonicity condition) or bi-directional responses of borrowing to grant aid (failure of the second stage condition).

This framework can be described in terms of set identification, as in Kline and Tartari (2016). Potential outcomes for borrowing (D^0 and D^1) define four groups of students. Three of these groups have been described, and monotonicity rules out the existence of the fourth group, "defiers" $D = \{i : D^1 = 1, D^0 = 0\}$. Each group would have an average amount borrowed under each value of the instrument as well as a population share, a total of twelve parameters. The number of parameters can be reduced using the definitions of the groups, which dictate that borrowing is zero for four combinations of instrument value and group. The fact that the population shares must sum to one reduces the number of parameters to be estimated to seven. However, we only obtain four pieces of information by estimating the overall borrowing rate and average amount borrowed under each value of the instrument. The parameter space would be under-identified without further assumptions. The monotonicity conditions rule out an entire group of students ("defiers") and thus eliminates two more parameters. The parameters for each individual group are still not point-identified, but their values can be bounded (e.g. neither the share of switchers nor the share of always-borrowers can exceed the total share of the population that borrows when $Z = 0$), and combinations of these parameters such as loan crowd-out among would-be borrowers can be point-identified.

With a continuous instrument, the generalization of the monotonicity conditions would be $\frac{d}{dz} g(z, X, V) \geq 0$ and $\frac{\partial}{\partial g} b(g, X, U) \leq 0$. Consider an instrument with lower and upper bounds $Z \in \{z_l, z_u\}$. Under the monotonicity conditions, the principal stratification of students categorizes each student by the largest value of the instrument for which she would borrow, which we label $\bar{z}$ with density $f(\bar{z})$. We would then define always-borrowers as $A = \{i : \bar{z} \geq z_u\}$, never-borrowers as $N = \{i : \bar{z} < z_l\}$, and switchers as $SW = \{i : z_l \leq \bar{z} < z_u\}$. Loan crowd-out due to a change in z from an initial value z^0 is now given by:

$$\int\left[\left.\frac{d}{dz}\right|_{z^0}E\left(B\right)\right]f\left(\bar{z}\right)d\bar{z}=$$
$$E\left[\left.\frac{d}{dz}\right|_{z^0}\int\limits_{-\infty}^{z_l}\left(B\right)f\left(\bar{z}\right)d\bar{z}\right]+E\left[\left.\frac{d}{dz}\right|_{z^0}\int\limits_{z_l}^{z_u}\left(B\right)f\left(\bar{z}\right)d\bar{z}\right]+E\left[\left.\frac{d}{dz}\right|_{z^0}\int\limits_{z_u}^{\infty}\left(B\right)f\left(\bar{z}\right)d\bar{z}\right]= \tag{6}$$
$$\pi_A E\left[\left.\frac{d}{dz}\right|_{z^0}B|A\right]+E\left[\left.\frac{d}{dz}\right|_{z^0}\int\limits_{z_l}^{z_u}\left(B\right)f\left(\bar{z}\right)d\bar{z}\right]$$

The above expression is similar to the discrete case given in equation (4) in that it expresses the observed response in terms of the response of always-borrowers and switchers, with a key difference being that the set of students who are switchers varies with the level of the instrument. With a coarse instrument, it is only possible to approximate $\frac{d}{dz}$, and the quality of the approximation would vary with the coarseness of Z because the set of switchers would increase with the distance between z and z^0, an instrument support condition described by Kasy (2014). Our continuous instrument has full support around z^0, which collapses the set of switchers to those with $\bar{z}=z^0$ and allows loan crowd-out among would-be borrowers to be defined analogously to the discrete case.

Thus, under first-stage and second-stage monotonicity conditions, the intensive-margin response of borrowing to grant aid can be identified for a set of students, namely those students who would borrow in the absence of an increase in grant aid ("would-be borrowers"). Our empirical setting and identifying variation, described in the following subsection, ensures that the first-stage monotonicity condition should hold. The second-stage monotonicity condition is predicted by both a simple life-cycle model (as in equation (1)) and by the more general decision model presented in Section 4.1.

4.3 Empirical strategies: RD and RK

We focus our empirical analyses on federal Pell Grants - by far the largest source of grant aid in our setting. We identify the effect of Pell Grant aid on students' educational investment decisions using variation induced by the kink and the discontinuity in the federal Pell Grant Program's formula. In a given year, a student's Pell Grant is a function of the federally set maximum and her EFC, the latter of which is determined by a complex federal formula. At the year-specific threshold level of EFC (efc_t^0), the slope of the Pell Grant schedule changes from -1 to 0 (the kink), and the level changes from the minimum award to 0 (the discontinuity). Figure 2 displays the empirical distribution of first-year students' Pell Grant aid, where EFC is standardized to represent distance from the year-specific eligibility threshold.[20]

[20]The empirical distribution of Pell Grant aid for returning students is similar (Appendix Figure B.3).

In terms of the general set-up outlined at the beginning of this section in which grant aid $G = g(Z, X, V)$, the instrument will be the indicator for having an EFC below the Pell Grant eligibility threshold, $\mathbf{1}\left[\widetilde{EFC} < 0\right]$, where $\widetilde{EFC} = EFC - efc^0$, in the case of the RD design, and the interaction between $\widetilde{EFC}$ and $\mathbf{1}\left[\widetilde{EFC} < 0\right]$ in the case of the RK design. The key identifying assumptions for the RK design are that at the eligibility threshold: (1) the direct marginal impact of EFC on Y is continuous and (2) the conditional density of EFC with respect to V is continuously differentiable (Card et al. 2015). These encompass the required identifying assumptions in the RD design (Hahn et al. 2001).[21] Furthermore, the second assumption generates testable predictions concerning how the density of $\widetilde{EFC}$ and the distribution of observable characteristics should behave in the neighborhood of the eligibility threshold. We show in Section 4.4 that these predictions are borne out in our empirical setting. The first stage monotonicity condition - that no students increase borrowing when receiving more grant aid - ensures that our estimates represent a local average treatment effect.

$\widetilde{EFC}$ imperfectly predicts a given student's Pell Grant. Thus, we estimate the effect of Pell Grant aid via two-stage least squares (2SLS). Consider the following first stage and reduced form equations, where i indexes students, t indexes years, c indexes entry cohorts, and j indexes colleges, $\widetilde{EFC}_{it} = EFC_{it} - efc_t^0$ is a standardized measure of the distance a student's EFC falls from the Pell Grant eligibility threshold in a given year, $f(\cdot)$ and $h(\cdot)$ are flexible functions of $\widetilde{EFC}$ (allowed to vary on either side of the eligibility threshold), and $\mathbf{X_{it}}$ is a vector of predetermined demographic characteristics:

$$G_{it} = f\left(\widetilde{EFC}_{it}\right) + \beta_1 \mathbf{1}\left[\widetilde{EFC}_{it} < 0\right] + \beta_2 \widetilde{EFC}_{it} \times \mathbf{1}\left[\widetilde{EFC}_{it} < 0\right] + \boldsymbol{\eta}\mathbf{X_{it}} + \delta_{jc} + \nu_{ijt} \tag{7}$$

$$B_{ijt} = h\left(\widetilde{EFC}_{it}\right) + \pi_1 \mathbf{1}\left[\widetilde{EFC}_{it} < 0\right] + \pi_2 \widetilde{EFC}_{it} \times \mathbf{1}\left[\widetilde{EFC}_{it} < 0\right] + \boldsymbol{\phi}\mathbf{X_{it}} + \alpha_{jc} + \epsilon_{ijt} \tag{8}$$

We choose the degree of polynomial in $\widetilde{EFC}$ that minimizes the Akaike Information Criterion (AIC), although we show that our estimates are robust to higher or lower order polynomials and local linear regression. In practice, we use both the kink and the discontinuity as instruments for Pell Grant aid, while also showing that our results are robust to using either only the kink or only the discontinuity for identification. Under the assumption that the instruments weakly increase grant aid for all students, 2SLS estimates will represent local average treatment effects.

Table 3 displays first stage estimates of the impact of the kink and discontinuity on Pell Grant aid by student level, where $f(\cdot)$ and $h(\cdot)$ are quadratic functions of $\widetilde{EFC}$, allowed to vary on either side of the

[21] Card et al. (2015) impose the additional identifying assumption that the right and left limits of $Pell\,(EFC)$are equal at the eligibility threshold. This assumption is clearly violated in our case, as there is both a discontinuity and kink. However, Card et al. (2015) show that under the assumption of locally constant treatment effects (e.g., $\frac{\partial Y}{\partial Pell}$ does not vary in the neighborhood of the Pell Grant eligibility threshold) this assumption can be relaxed without affecting identification.

eligibility threshold. On average, barely-eligible first-year students receive approximately \$390 in Pell Grant aid, and for every dollar decrease in $\widetilde{EFC}$, Pell Grant aid increases by approximately \$0.76. Point estimates for the sample of returning students are similar. To estimate the effect of additional grant aid on attainment outcomes, we replace B_{ijt} with Y_{ijt} in equation (8), where Y_{ijt} represents one of several attainment outcomes (reenrollment in the following academic year, credits attempted, credits earned, or cumulative GPA).[22]

Estimating loan crowd-out among would-be borrowers requires estimates of $E\left[B^1{-}B^0\right]/E\left[G^1{-}G^0\right]$ and $\pi_B = \pi_A + \pi_{SW}$. Let π_B^0 denote the share of would-be borrowers when $G^0 = 0$, and let π_B^{min} denote the share of would-be borrowers when $G^0 = minPell$, the baseline grant amounts for the RD and RK designs, respectively. At the discontinuity, we estimate π_B^0 as the limit of the borrowing rate as $\widetilde{EFC}$ approaches the threshold from the right (i.e. using only the sample of Pell-ineligible students) by estimating:

$$D_{it} = \pi_B + \widetilde{EFC}_{it} + \omega_{it} \tag{10}$$

We estimate $E\left[B^1{-}B^0\right]/E\left[G^1{-}G^0\right]$ with $\hat{\tau}_{RD}$, the RD estimate of the impact of a dollar of Pell Grant aid on the dollar amount of loans, and $\frac{\hat{\tau}_{RD}}{\hat{\pi}_B^0}$ provides an estimate of crowd-out among would-be borrowers. For our continuous instrument, the kink in the Pell Grant formula, would-be borrowers are students who would have borrowed if they received the minimum Pell Grant. Thus, we estimate π_B^{min} via equation (10) as the left-hand limit of the borrowing rate (i.e. restricting the sample to Pell-eligible students), and $\frac{\hat{\tau}_{RK}}{\hat{\pi}_B^{min}(minPell)}$ provides a second estimate of crowd-out among would-be borrowers. Finally, to generate a conservative estimate of crowd-out among would-be borrowers, we scale our combined RD/RK estimator by the π_B^0 (which is larger than π_B^{min}).

4.4 Evaluating the RD and RK identifying assumptions

We evaluate the RD/RK identifying assumptions through three exercises. First, we test for discontinuities in the level and slope of the density of CUNY students at the Pell Grant eligibility threshold. Second, we

[22] We also test whether Pell Grant aid has persistent impacts on educational investment, by regressing attainment in year $t+n$ on Pell Grant aid received in year t. We estimate 2SLS models where the second stage takes the form:

$$Y_{ist} = \tau_n \widehat{G}_{it-n} + k_n\left(\widetilde{EFC}_{it-n}\right) + \varsigma \mathbf{X_{it}} + \varphi_{sc} + \epsilon_{istn} \tag{9}$$

Here, τ_n represents the impact of an additional dollar of Pell Grant aid in year $t-n$ on the year t outcome, *vis-à-vis* all other intermediate outcomes affected by Pell Grant aid (including future grants awards). If Pell Grant aid received in one year affects aid received in subsequent years, $\hat{\tau}_n$ represents the Cellini et al. (2010) "intent to treat" effect and we would need to use dynamic regression discontinuity methods to estimate "treatment on the treated" effects. An additional dollar of Pell Grant aid in a student's first year leads to an approximately \$1.19 increase in cumulative Pell Grant aid received three years after entry (Appendix Table B.3). As \$1 of this amount comes from the mechanical impact of first-year Pell Grant aid on cumulative Pell Grant aid, we assess the impact of first-year Pell Grant aid on future Pell Grant aid by testing the hypothesis that cumulative impacts on Pell Grant aid equal \$1. Since we do not reject this hypothesis ($p = 0.140$), we conclude that Pell Grant aid in year t does not have affect Pell Grant aid in future years and that equation (9) will also produce estimates of the "treatment on the treated" impact of first-year Pell Grant aid.

similarly test for discontinuities in the probability of attendance conditional on submitting an application to the CUNY system. Finally, we test for discontinuities in the level and slope of the distribution of observable characteristics, including gender, race, immigrant status, family AGI, parental education, and dependency status.

As shown in Panel A of Figure 3, in the sample of first-year students, the density function's level and slope are continuous through the Pell Grant threshold.[23] The lack of discontinuities in the density function provides suggestive evidence that Pell Grant generosity does not influence students' initial enrollment decisions. However, finding no change in the density of CUNY students is not sufficient to rule out a more complicated story that includes both an increase in CUNY enrollment and some potential CUNY students "upgrading" to more expensive, non-CUNY schools upon receiving Pell Grant aid. Therefore, we also test for discontinuities in the density of applications and the probability of attendance, conditional on applying, at the threshold. We match CUNY applicant data to FAFSA and enrollment information for the fall 2007 through fall 2010 applicant cohorts. We observe each applicant's ranking of up to six CUNY institutions and to which institution (if any) she ultimately matriculated. Panel B of Figure 3 displays the density of applications and the conditional probability of enrollment by $\widetilde{EFC}$. The probability of enrollment conditional on application is continuous. There is a slight decrease in the number of applications to the left of the threshold, but the size of this discontinuity is not statistically significant. Thus, we maintain the assumption that the probability of submitting an application does not depend on whether a student receives Pell Grant aid, which is consistent with the fact that students do not learn of their Pell Grant award until after they have applied and been admitted to a CUNY institution.

We formally estimate the change in the probability of enrollment conditional on application; OLS estimates of equation (8) and 2SLS estimates are displayed in Table 4. We include a linear term in $\widetilde{EFC}$, which minimizes the AIC in all first stage and reduced form specifications. The first three columns present estimated impacts on enrollment in a given CUNY institution, conditional on application, which enables us to test for "upgrading" from community colleges to four-year schools within the CUNY System. Each observation represents a prospective student by application combination. The fourth column presents estimates of the impact of Pell Grant aid on the probability of enrollment in any CUNY institution, and each observation represents a prospective student.[24]

We estimate that a \$1,000 increase in Pell Grant aid leads to an insignificant 0.003 percentage point increase in the probability of enrollment in a given two- or four-year institution (Panel B, Column 1). The

[23] Appendix Figure B.4 displays the density of returning CUNY students in our sample.

[24] Estimates from the first stage equation (7) are displayed in Panel A of Appendix Table B.4. As our sample of applicants does not include prospective students in the fall 2006 cohort, the estimated change in Pell Grant aid at the eligibility threshold is slightly larger than our estimate displayed in Table 3 (\$476 versus \$389).

95 percent confidence interval excludes effects as small as an additional $1,000 of Pell Grant aid leading to a 0.8 percentage point increase in the probability of enrollment, which represents a 4.8 percent increase relative to ineligible students' mean enrollment rate.[25] These estimates may understate the impact of Pell Grant aid on college going if, for instance, Pell Grant aid increases the probability of enrollment in a CUNY community college while also inducing students who would have attended a four-year CUNY institution to enroll in a more expensive school. Therefore, we divide our sample of applicants into community college applicants and four-year college applicants. As shown in Columns 2 and 3, we find no evidence that, conditional on submitting an application, Pell Grant aid affects the probability of enrollment in a CUNY community college or four-year institution. In the community college sample, the point estimate from our main specification suggests that an additional $1,000 in Pell Grant aid results in an insignificant 1.1 percentage point (5 percent) increase in the probability of enrollment. In the four-year college sample, our point estimate suggests that a $1,000 Pell Grant award leads to an insignificant 0.5 percentage point (3.5 percent) decrease in enrollment among applicants. These small, insignificant effect sizes are consistent with evidence against Pell Grant induced "upgrading" in the nationally representative sample examined by Turner (2014) and among Tennessee high school graduates that Carruthers and Welch (2015) study.

Our fourth specification examines impacts on enrollment in any CUNY institution as a function of Pell Grant aid. The point estimate from our main specification is significant at the 10 percent level but small in magnitude, suggesting that a $1,000 increase in Pell Grant aid leads to a 1.4 percentage point increase in the probability of enrollment in any CUNY institution. Our 95 percent confidence interval allows us to rule out enrollment impacts greater than 2.2 percent. Finding no evidence of consistent effects of Pell Grant aid on enrollment, we proceed by analyzing impacts of Pell Grant aid on the population of CUNY enrollees, for whom we observe borrowing and educational attainment.

We find little evidence of discontinuous changes in the level or slope of the distribution of observable characteristics (Figure 4). Appendix Table B.5 contains corresponding point estimates from regressions of these characteristics on the kink and discontinuity, degree program and school by year fixed effects, and a polynomial in $\widetilde{EFC}$, allowed to vary on either side of the Pell Grant eligibility threshold. We estimate a statistically significant, negative relationship between AGI and Pell Grant eligibility, although when we include higher order polynomials in $\widetilde{EFC}$, these estimates are no longer significant. Outside of AGI, only one of the 24 point estimates is statistically distinguishable from zero. We control for these characteristics in our main specification, while also showing that our estimates are robust to excluding these controls.

[25] Estimates from specifications that only use the discontinuity or kink in Pell Grant aid for identification are similar but less precise (Appendix Table B.4, Panels B and C).

5 The Impact of Pell Grant Aid on Educational Investment

Our model predicts that Pell Grant aid reduces unconstrained students' borrowing, with crowd-out exceeding 100 percent when the fixed cost binds. Predicted effects on educational attainment vary, with "threshold borrowers" increasing schooling, unconstrained students not altering their schooling, and students who stop borrowing in response to additional Pell Grant aid ("switchers") reducing schooling.

5.1 Pell Grant aid reduces borrowing

Figure 5 displays first year students' mean loans (Panel A) and probability of borrowing (Panel B) by distance from the Pell Grant eligibility threshold. At the eligibility threshold, average loan aid and the probability of borrowing fall discontinuously, and the relationship between borrowing and EFC also changes discontinuously, indicating that on average, students reduce borrowing upon receiving additional grant aid. Returning students' borrowing follows similar patterns (Appendix Figure B.5).

We quantify the contemporaneous impact of Pell Grant aid on borrowing separately for new and returning students via OLS and 2SLS (Table 5). Panel A presents reduced form estimates of Pell Grant eligibility and generosity on student loan aid from equation (8). Panel B displays 2SLS estimates of the impact of Pell Grant aid on debt using both the kink and discontinuity as excluded instruments. An additional dollar of Pell Grant aid induces first-year students to reduce borrowing by $0.43. Returning students respond to an additional dollar of Pell Grant aid by forgoing $0.51 in loan aid. Appendix Table B.6 contains estimates from separate IV-RD and IV-RK models, which largely support our assumption of locally constant treatment effects. Point estimates using only the discontinuity as an instrument are larger in magnitude than estimates obtained from using only the kink, but among new students, these estimates are not statistically distinguishable. Among returning students, differences between estimates obtained from IV-RD and IV-RK models are marginally significant, with $p = 0.098$.[26]

Few CUNY students borrow (Table 1). The magnitude of the estimated impact of Pell Grant aid on loans suggests large responses among a subset of students. As described in Section 4.2, we can identify loan crowd-out among would-be borrowers by using extensive-margin borrowing responses to quantify the number of never-borrowers, always-borrowers, and switchers. Crowd-out of loans exceeds 100 percent among would-be borrowers in even the most conservative pair of estimates. As shown in Table 6, an additional dollar of Pell Grant aid crowds-out over $1.80 in loan aid among first-year would-be borrowers.[27] Returning would-be

[26]Appendix Table B.3 displays estimated impacts of an additional dollar of Pell Grant aid in a student's first year on cumulative student loan debt three years after entry. Pell Grant aid has persistent effects on borrowing; an additional $1,000 of Pell Grant aid in a student's first year reduces cumulative debt by close to $600 three years after entry, a 57 percent decrease from the sample mean.

[27]Standard errors are generated via block bootstrap that allows for clustering at the institution by year level.

borrowers reduce loans by close to \$1.90 for every dollar increase in Pell Grant aid. We can reject the hypothesis that crowd-out for borrowers falls below 100 percent at the 99 percent level. Estimates of crowd-out using only either the discontinuity or the kink as instruments are larger in absolute value but less precise (Appendix Table B.7).

A fixed cost of borrowing can rationalize our finding of crowd-out exceeding 100 percent and also offers predictions about the distributions of loans and changes in the distribution in response to changes in grant aid. As shown in the decision model in Section 4.1, when borrowing entails a fixed cost, a small increase in grant aid may cause a student to switch from borrowing hundreds or thousands of dollars to borrowing nothing. This is because there is a range $(0, \underline{d})$ in which the amount of debt that would maximize utility in the absence of the fixed cost would produce only a small utility gain over not borrowing and hence, would not be worth paying the fixed cost. As a result, few students should borrow small amounts and crowd-out will be largest among students who would otherwise take up small loans.

As shown in Figure 6, which displays the distribution of loans for first-year borrowers subject to the exogenous subsidized loan limit of \$3,500, students are unlikely to take up small amounts of debt. Estimates of the quantile treatment effects of grant aid on loans reinforce the fixed cost interpretation by showing that the impact of Pell Grant eligibility on borrowing is larger for quantiles corresponding to small positive amounts of debt. Figure 7 displays CDFs of loan amounts for first-year (Panel A) and returning students (Panel B) within \$1,000 of the Pell Grant eligibility threshold. The horizontal distance between the curves provides a reduced-form estimate of the corresponding quantile treatment effect of Pell Grant eligibility. Among first-year students, Pell Grants do not affect borrowing below the 76th quantile because three-quarters of students borrow nothing irrespective of their Pell Grant eligibility. Differences in borrowing between eligible and ineligible students are also small at the highest quantiles, but at intermediate quantiles, a less than \$1,000 Pell Grant corresponds to a \$3,000 difference in borrowing. The patterns for quantiles in which ineligible students borrow but eligible students do not borrow suggest heterogeneous values of $\underline{d}$ that may reach into the thousands of dollars. These patterns can only provide an unbiased estimate of $\underline{d}$ if Pell Grants do not induce students to switch quantiles and the assumption of rank-invariance holds.[28]

[28] A Tobit model that treats $\underline{d}$ as constant across students (i.e., equal to zero or some other small amount) will fail to capture this feature of the distribution. In order to fit the low share of students with loans, such a model offers the unrealistic prediction that the average student's latent desired loan is large and negative (i.e., that the average student wishes to *save* tens of thousands of dollars). An alternative model that allows students to face heterogeneous fixed borrowing costs can rationalize the upward sloping density between small loan values and \$2,000. Alternative economic models with discrete adjustments, such as a labor supply model with adjustment frictions, could generate crowd-out rates exceeding 100 percent but would not predict reductions specifically among small loans without strong assumptions. Indeed, crowd-out of loans is likely *reduced* to the extent that grants induce students to make discrete reductions in work hours rather than responding only through borrowing. Likewise, borrowing frictions at positive levels of debt, such as mental anchoring at loan limits, would not generate the extensive-margin response to Pell Grant aid that we find.

5.2 Robustness of the estimated impact of Pell Grant aid on borrowing

We perform a number of robustness tests to show that our borrowing results, including our finding of crowd-out exceeding 100 percent among would-be borrowers, are not driven by factors outside of the discontinuities in Pell Grant aid at the eligibility threshold. Our estimates are robust to allowing Pell Grant aid to affect receipt of other grant aid, restriction of the sample to avoid mechanical effects due to changes in subsidized loan eligibility limits, exclusion of covariates, alternative bandwidths and polynomials, and local linear regression.

Table 7 presents results from three robustness tests. First, given the finding of Turner (2014) that public colleges supplement Pell Grant aid with additional institutional grants, we want to rule out the possibility that our estimated borrowing responses exceed 100 percent because other sources of grant aid increase with Pell Grant aid.[29] To do so, we replace the dependent variable in equation (7) with the sum of Pell and other grant aid; results are consistent with those generated by our main specification (Panel A). Second, we show that estimates from models that exclude all covariates besides the quadratic in $\widetilde{EFC}_{it}$ are consistent with our main results (Panel B). Finally, we address the concern that Pell Grant aid may mechanically decrease some students' borrowing by reducing subsidized loan eligibility. Subsidized loans cannot exceed unmet need (equal to the cost of attendance minus EFC and grant aid). With sufficiently low unmet need, crossing the Pell Grant eligibility threshold will reduce subsidized loan eligibility by an amount that is, at most, equal to the minimum Pell Grant.[30] We show that mechanical changes in subsidized loan eligibility do not drive our results by exploiting the fact that subsidized loans are also capped by a particular dollar amount for each year of schooling (e.g. $3,500 for first-year students). We limit our sample to students whose cost of attendance minus EFC and non-Pell grant aid is greater than the subsidized loan cap for their year in school plus the minimum Pell Grant for the calendar year, a population for whom the Pell Grant could not affect eligibility for subsidized loans, and we obtain similar results (Panel C).[31]

[29] Appendix Figure B.6 displays the distribution of total grant aid by distance from the Pell Grant eligibility threshold and the contributions of Pell, TAP, and other grant aid to this amount. We find no evidence that Pell Grant aid affects TAP Grant aid (Appendix Table B.8, Panel A). Pell Grant aid is positively correlated with other grant aid (a category that includes institutional, federal, and non-TAP New York State aid) for first-year students who receive an additional $0.08 in other grant aid for every dollar of Pell Grant aid (Appendix Table B.8, Panel B).

[30] There are several reasons to doubt the importance of this "mechanical" effect. First, while grant aid affects some students' subsidized loan eligibility, it does not alter any student's total loan eligibility. Thus, such effects would only reduce borrowing among students who maximize their subsidized loans and avoid unsubsidized loans. Second, the effect on subsidized loan eligibility itself would not generate loan crowd-out rates in excess of 100 percent unless students chose to reduce loans by more than the mechanical offset of 100 percent. Finally, it is irrelevant for the regression kink design because every dollar increase in Pell Grant aid results from a dollar decrease in EFC, leaving unmet need and subsidized loan eligibility unaffected. Appendix Table B.7 shows that our main results hold when we only use variation from the kink for identification.

[31] Panel A of Appendix Figure B.7 displays the distribution of subsidized loan aid among the full sample of first-year students while Panel B displays the distribution for students whose subsidized loan eligibility is not mechanically affected by Pell Grant aid. The kink and discontinuity are of a similar magnitude in both groups. Since students do not take up unsubsidized loans until their eligibility for subsidized borrowing is exhausted, overall, fewer students take up unsubsidized loans and the resulting discontinuity in unconditional unsubsidized loan aid at the Pell Grant eligibility threshold is also smaller (Appendix Figure B.7, Panel C). When we limit our sample to students whose subsidized loan eligibility is not mechanically affected by Pell Grant aid, estimated crowd-out of subsidized loans still significantly exceeds 100 percent (Appendix Table B.9).

Second, we wish to rule out the possibility that our estimates are driven by our choice of bandwidth or polynomial in $\widetilde{EFC}$. We estimate 2SLS models in which we focus on students with EFCs within \$4,000, \$3,000, \$2,000, and \$1,000 of the Pell Grant eligibility threshold. Within each window, we allow for up to a fourth degree polynomial in $\widetilde{EFC}$ and list the degree of polynomial that minimizes the AIC. Finally, we employ the goodness-of-fit test suggested by Lee and Lemieux (2010), by testing the joint significance of \$200 $\widetilde{EFC}$ bin dummies added to our main specification (brackets contain p-values from this test). This exercise also directly tests for discontinuities in borrowing away from the Pell Grant eligibility threshold. Appendix Table B.10 displays these estimates for first-year students. Our estimates are robust to smaller windows and higher order polynomials. Although precision decreases, the point estimates increase in magnitude when we reduce the size of the window around the Pell Grant eligibility threshold. All but two of the 16 point estimates suggest that crowd-out among would-be borrowers significantly exceeds 100 percent.

Finally, we estimate the impact of Pell Grant aid on first-year students' borrowing via local linear regression. We use the Imbens and Kalyanaraman (2012) optimal bandwidth (Panel A), the Fan and Gijbels (1996) rule-of-thumb bandwidth (Panel B), and the coverage error optimal bandwidth chosen by the procedure proposed in Calonico et al. (2014), Calonico et al. (2016a), and Calonico et al. (2016b) (Panel C). In all cases, we use a uniform (rectangular) kernel and cluster standard errors at the institution by year level. In Appendix Table B.11, we report first-stage, reduced form, and 2SLS estimates (using either the kink, the discontinuity, or both as instruments) as well as estimated crowd-out for would-be borrowers. In the case of 2SLS specifications, we use the bandwidth chosen in the reduced form specification. In each case, our estimates are less precise but consistent with those obtained from our parametric specification, and crowd-out among would-be borrowers significantly exceeds 100 percent in every specification.

5.3 Impacts on educational attainment

We generate 2SLS estimates of the impact of an additional \$1,000 in Pell Grant aid on educational outcomes, including reenrollment (the probability of remaining enrolled in the following academic year), effort (academic and remedial credits attempted), attainment (academic credits earned), and performance (GPA) (Table 8).[32] On average, Pell Grant aid has small and insignificant impacts on all outcomes except for credits attempted by first-year students, where our estimate suggests that an additional \$1,000 of Pell Grant aid induces students to attempt an additional 0.5 credits (an approximately 2 percent increase). This increase does not translate into a significant increase in academic credits earned. We can rule out first-year impacts of an additional \$1,000 of Pell Grant larger than a 1 credit (6 percent) increase.

[32] Appendix Figure B.8 displays the reduced form relationship between $\widetilde{EFC}$ and first-year students' attainment.

We also test whether Pell Grant aid affects longer-run attainment. As we find no evidence that first-year Pell Grant aid affects Pell Grant aid received in future years, we estimate the impact of first-year Pell Grant aid on future outcomes. The second column of Table 8 displays estimates of first-year Pell Grant aid on reenrollment, cumulative credits attempted, and cumulative credits earned three years after entry. An additional $1,000 of Pell Grant aid in a student's first year leads to an insignificant increase of 0.2 cumulative credits. We can rule out impacts on cumulative credits that are larger than a 2.6 credit (6 percent) gain three years after entry, suggesting that, on average, Pell Grant aid does little to increase the longer-run educational attainment of CUNY students. Estimated impacts of Pell Grant aid on contemporaneous attainment of returning students, shown in Column 3, are quite similar.

Our finding that the average impact of Pell Grant aid on educational attainment is not significantly different from zero may mean one of two things. First, it could be that students near the Pell Grant eligibility threshold are generally unresponsive to grant aid. Alternatively, there could be heterogeneous and offsetting treatment effects across groups of students, as suggested by our decision model in Section 4.1. Grants induce switchers to stop borrowing, resulting in a decrease in educational attainment. Conversely, threshold borrowers increase their educational attainment in response to grant aid.

We provide suggestive evidence of heterogeneous educational effects consistent with this predicted pattern using an alternative identification strategy that exploits the panel nature of our data. We limit our sample to students that we observe at least twice within the $4,000 window around the Pell Grant eligibility threshold and estimate a model that includes student fixed effects, a flexible function of EFC, controls for time-varying student characteristics, and year, level, and years since entry fixed effects. Thus, we take advantage of within-student variation in Pell Grant aid over time, rather than across-student variation in Pell Grant within a given year. To account for concerns of reverse causality (e.g., that students who take less than a full course load receive a prorated Pell Grant), we instrument for a student's actual Pell Grant with predicted Pell Grant for a full-time, full-year student. We then examine the effect of Pell Grant aid conditional on borrowing, the effect of Pell Grant aid conditional on not borrowing, and the correlation between borrowing and attainment.

The results of the student-fixed-effects attainment regressions appear in Table 9. The first row shows that conditional on not borrowing, Pell Grant aid significantly increases credits attempted and earned, suggesting that an additional $1,000 in Pell Grant aid leads to an additional 0.5 credits attempted and 0.4 credits earned. The second row of Table 9 shows positive but insignificant effects of Pell Grant aid when students borrow. The last row shows that when students switch from borrowing to not borrowing they attempt and complete significantly fewer credits (a reduction of approximately 2 credits attempted and 1 credit earned). These patterns are consistent with our theoretical predictions that Pell Grant aid will increase attainment of those whose borrowing is unaffected by additional grant aid (significantly so for nonborrowers, empirically) while decreasing the attainment of those induced to stop borrowing. The coefficient on the indicator for not having loans gives a partial correlation that will underestimate the causal impact of borrowing on attainment if unobservable financial shocks cause students to both borrow more and take fewer credits. The true nature of the endogeneity is unknown implies that stronger identifying assumptions are required for a causal interpretation, and thus, we take these results to be suggestive.

5.4 Characterizing the Fixed Cost of Borrowing

The fixed cost faced by CUNY students may be driven by some combination of psychic costs caused by debt aversion, hassle and administrative costs caused by paperwork and other requirements, and opt-in costs caused by the default offer of zero. In this subsection, we provide evidence that among non-CUNY schools it is only in those employing a zero default offer that crowd-out exceeds 100 percent. If CUNY students face psychic costs of borrowing similar to those faced by public school students enrolled in other schools that make nonzero loan offers, debt aversion cannot be the only driver of the fixed cost. Finally, we show that within the CUNY system, borrowing patterns do not vary with implied time and hassle costs.

To test whether non-CUNY students have similar borrowing responses to Pell Grant aid when their schools employ a default loan offer of zero, we use nationally representative data from the 2012 NPSAS. We limit the NPSAS sample to first- through third-year public school students who have an EFC within $4,000 of the Pell Grant eligibility threshold and estimate a model in which our instruments and Pell Grant aid are fully interacted with an indicator for whether a school employs a nonzero default loan offer. Due to the small number of observations in the $4,000 EFC window in the NPSAS, we only use a subset of the control variables from our main specification, including a linear term in $\widetilde{EFC}$, dependency status, AGI, and degree-program (interacted with the packaging indicator, as most bachelor's degree seeking students outside of the CUNY system attend schools that package loans). 2SLS estimates and estimated impacts of Pell Grant aid on would-be borrowers are displayed in Table 10.[33] This approach requires the identifying assumption that the probability of enrolling in a community college with a zero (versus nonzero) default is continuous and smooth through the Pell Grant eligibility threshold. The small sample size of the NPSAS limits our ability to generate a strict test, but we deem the assumption reasonable given that (a) most community college students attend the nearest school, and (b) the fact that past research does not find significant impacts of Pell Grant aid on college enrollment or school quality (Kane 1995; Rubin 2011; Carruthers and Welch 2015; Turner 2014).

Pell Grant aid leads to significantly larger reductions in borrowing for students attending schools that,

[33]See Appendix Table B.12 for first stage and reduced form impacts of Pell Grant eligibility and generosity.

like CUNY, do not package loans (-0.61 versus -0.40 in schools with nonzero default offers, $p = 0.080$), even though borrowing rates in such schools are substantially lower than loan take-up in packaging schools. Furthermore, crowd-out for would-be borrowers only significantly exceeds 100 percent in schools with a zero loan default offer. Estimated crowd-out for would-be borrowers attending zero-default schools (-1.54) is only slightly smaller than our estimates generated from the CUNY sample, supporting our hypothesis that opt-in costs due to the loan default choices account for a substantial portion of the fixed cost faced by CUNY students. To the extent that other costs associated with borrowing (e.g., entrance counseling, filling out forms, debt aversion) are common to students across schools with different default loan offers, we can rule out such explanations as factors driving the fixed cost faced by CUNY students.

Although the default loan offer made by all CUNY institutions is zero, in the years we examine, schools differ in the process through which students opt into borrowing. While all CUNY schools require a short supplemental loan application, four institutions allowed students to submit this application online, while the remaining 13 required students to submit the application to their institution's financial aid office in person.[34] If hassle and time costs substantially contribute to the fixed borrowing cost, we would expect lower crowd-out among students who could submit an application online. Under the identifying assumption that schools that offer an online loan application do not differ from those that do not in ways that also influence borrowing decisions, we can test for heterogeneity in the impact of Pell Grant aid on borrowing along this dimension by fully interacting an indicator for having an online loan application with Pell Grant aid and our instruments. We find no evidence that the impact of Pell Grant aid on borrowing varies by the availability of an online loan application (Appendix Table B.13). Estimates of crowd-out among would-be borrowers are very similar for students attending colleges that offer an online application and colleges that do not (-1.72 versus -1.87). Given our finding of crowd-out exceeding 100 percent for NPSAS students attending institutions with zero loan offer defaults outside of the CUNY system and that we find no evidence that borrowing patterns vary by implied time and hassle costs, we conclude that loan opt-in costs drive the fixed cost faced by CUNY students.

5.5 Generalizability

We conduct several exercises to illustrate the generalizability of our findings outside of the CUNY system. Compared to students in the nationally representative NPSAS sample, CUNY students are more likely to be nonwhite, classified as dependent students, first or second generation immigrants, and are less likely to have a college educated parent. Given the differences in the demographic characteristics of CUNY students and public school students as a whole, we would like to determine the extent to which our results depend

[34]Information from archived institution websites pertaining to student loans was obtained through web.archive.org.

on the distinct characteristics of the CUNY population. Appendix Table B.14 presents estimated impacts of Pell Grant aid on borrowing and credits earned by first-year students from models where Pell Grant aid and our instruments are fully interacted with indicators for whether the student is white, a first- or second-generation immigrant, classified as dependent, or has a college educated parent. We test the equality of overall impacts on borrowing and impacts for would-be borrowers between groups defined by these characteristics. Unconditional crowd-out among immigrants is smaller than for non-immigrants ($p = 0.053$), but none of the estimated impacts on would-be borrowers are statistically distinguishable along any of these dimensions. We find evidence of heterogeneous impacts of Pell Grant aid on attainment by dependency status, with a \$1,000 increase in Pell Grant aid leading independent students to earn a statistically significant 3 additional credits, while impacts on credits earned by dependent students are negative and insignificant. We can reject the equality of estimated impacts on credits earned by dependency status with $p = 0.016$.

The possibility of nonlinear grant aid treatment effects raises another potential concern for the generalizability of our results. For example, the minimum Pell Grant may not be large enough to improve educational outcomes, but larger grants could be more effective. Similarly, a marginal increase in grant aid may have different effects on students with different base amounts of grant aid. To address these concerns, we first test whether our main estimates vary by TAP grant aid receipt. Approximately 74 percent of students in our sample receive a TAP grant equal to an average of \$1,815. Second, we test whether the estimated impact of Pell Grant aid varies with the minimum Pell Grant.[35] We test for heterogeneity by "high minimum Pell" years (2009 and 2010) in which, on average, the minimum Pell Grant was more than twice as large as in lower minimum Pell Grant years (2007, 2008, and 2011). Appendix Table B.15 displays these estimates. Neither TAP grant receipt nor the size of the minimum Pell Grant significantly interact with Pell Grant aid, providing support for our assumption of linearity in the impacts of Pell Grant aid.

Finally, our estimates of crowd-out for would-be borrowers are driven by "switchers" who are induced to cease borrowing by a small increase in grant aid. If switchers have unique characteristics, our crowd-out estimates may not be generalizable to the broader population of Pell-eligible and near-eligible students. Although switchers cannot be identified individually, it is possible to identify their average characteristics under the assumption of monotonicity (Abadie 2003) - that increases in Pell Grant aid do not cause any students to switch from not borrowing to borrowing. Appendix Table B.16 shows the mean value of various characteristics among switchers and compares these to the sample mean of students near the Pell Grant eligibility threshold.[36] Switchers do not significantly differ from other students in any observable characteristics.

[35]The minimum Pell Grant was \$400 in 2007 and 2008, \$890 in 2009, \$976 in 2010, and \$555 in 2011 (all in nominal terms).
[36]We thank an anonymous referee for this suggestion.

6 Conclusion

In this paper, we take advantage of the nonlinearities in the Pell Grant Program's formula to estimate the impact of need-based grant aid on borrowing and educational attainment. Our main findings - that Pell Grant aid reduces borrowing and has small to zero average impacts on educational outcomes - is consistent with traditional models of educational investment under credit constraints. We observe very few CUNY students exhausting their federal loan eligibility, suggesting most students do not face traditional borrowing constraints.

However, among students who would borrow had they not received additional Pell Grant aid, a \$1 increase in Pell Grant aid leads to a greater than \$1 reduction in loans, which is inconsistent with traditional life-cycle models in which the marginal cost of borrowing is continuous in the amount borrowed. To explain this irregularity, we extend the traditional credit constraints framework to allow for a fixed cost of borrowing. Our model predicts that grant aid may actually reduce these students' educational attainment, offsetting the expected improvements among students constrained by the fixed borrowing cost or loan limits, and provide suggestive evidence in support of this effect, potentially explaining our finding of no aggregate impact of Pell Grant aid on attainment. The possibility of a subsidy having the opposite of the desired effect when the desired activity is related to a lumpy decision may prove applicable in many other settings, such as retirement savings programs with default saving rates or the taxation of pollutants with adjustment costs in production.

We provide evidence that the fixed cost faced by CUNY students is driven by opt-in costs. Our findings are likely relevant for a substantial portion of federal student aid recipients. Pell Grant recipients are more likely to attend a community college than any other type of institution, and the majority of community colleges follow the CUNY practice of a zero default loan offer. If the fixed cost we identify is driven by the default loan offer being set to zero, such schools can reduce or eliminate the fixed cost with little or no expense by switching to a nonzero default or requiring students to make an active borrowing decision.

In 2016, outstanding student loan debt exceeded \$1.26 trillion (Federal Reserve Bank of New York 2016). While Dunlop (2013) and Wiederspan (2016) show that access to federal loan aid increases community college students' attainment, in general, there is limited evidence concerning the impact of federal loan aid on student outcomes. Furthermore, while estimated returns to higher education suggest that borrowing to finance college is optimal (Avery and Turner 2012), student loan debt may impose costs that alter students' behavior when they enter the labor force or while students are still making educational investments (e.g., Field 2009; Rothstein and Rouse 2011). Imposing a fixed borrowing cost may reduce welfare by reducing educational attainment or increasing other types of debt, but it may enhance welfare if student debt distorts future decisions. We leave welfare analysis and estimation of these interesting parameters to future work.

References

Abadie, Alberto, "Semiparametric Instrumental Variable Estimation of Treatment Response Models," *Journal of Econometrics*, 2003, *113* (2), 231–263.

Avery, Christopher and Sarah Turner, "Student Loans: Do College Students Borrow Too Much - Or Not Enough?," *Journal of Economic Perspectives*, 2012, *26* (1), 165–192.

Bettinger, Eric, "How Financial Aid Affects Persistence," in Caroline M. Hoxby, ed., *College Choices: The Economics of Where to Go, When to Go, and How to Pay For It*, University of Chicago Press, 2004, chapter 5, pp. 207–237.

Bettinger, Eric P., Bridget Terry Long, Philip Oreopolous, and Lisa Sanbonmastu, "The Role of Simplification and Information in College Decisions: Results from the H&R Block FAFSA Experiment," *Quarterly Journal of Economics*, 2012, *127* (3), 1205–1242.

Bound, John, Michael F. Lovenheim, and Sarah Turner, "Why Have College Completion Rates Declined? An Analysis of Changing Student Preparation and Collegiate Resources," *American Economic Journal: Applied Economics*, 2010, *2* (3), 129–157.

Cadena, Brian C. and Benjamin J. Keys, "Can Self-Control Explain Avoiding Free Money? Evidence from Interest-Free Student Loans," *Review of Economics and Statistics*, 2013, *95* (4), 1117–1129.

Calonico, Sebastian, Matias D. Cattaneo, and Max H. Farrell, "Coverage Error Optimal Confidence Intervals for Regression Discontinuity Designs," 2016. working paper.

_ , _ , and Rocío Titiunik, "Robust Nonparametric Confidence Intervals for Regression-Discontinuity Designs," *Econometrica*, 2014, *82* (6), 2295–2326.

_ , _ , Max H. Farrell, and Rocío Titiunik, "`rdrobust`: Software for Regression Discontinuity Designs," 2016. working paper.

Cameron, Stephen V. and Christopher Taber, "Estimaton of Educational Borrowing Constraints Using Returns to Schooling," *Journal of Political Economy*, 2004, *112* (1), 132–182.

Card, David, David S. Lee, Zhuan Pei, and Andrea Weber, "Inference on Causal Effects in a Generalized Regression Kink Design," *Econometrica*, 2015, *83*, 2453–2483.

Carruthers, Celeste K. and Jilleah G. Welch, "Not Whether, but Where? Pell Grants and College Choices," 2015. Working paper.

Cellini, Stephanie Riegg, Fernando Ferreira, and Jesse Rothstein, "The Value of School Facility Investments: Evidence from a Dynamic Regression Discontinuity Design," *Quarterly Journal of Economics*, 2010, *125* (1), 215–261.

Chetty, Raj, John N. Friedman, Søren Leth-Petersen, Torben Heien Nielsen, and Tore Olsen, "Active vs. Passive Decisions and Crowd-Out in Retirement Savings Accounts: Evidence from Denmark," *Quarterly Journal of Economics*, 2014, *129* (3), 1141–1219.

Choi, James J., David Laibson, Brigitte C. Madrian, and Andrew Metrick, "Optimal Defaults and Active Decisions," *Quarterly Journal of Economics*, 2006, *124* (4), 1639–1674.

Cochrane, Debbie and Laura Szabo-Kubitz, "At What Cost? How Community Colleges that Do Not Offer Federal Loans Put Students at Risk," 2014. Institute for College Access and Success Report.

Deming, David and Susan Dynarski, "College Aid," in Phillip Levine and David Zimmerman, eds., *Targeting Investments in Children: Fighting Poverty When Resources are Limited*, The University of Chicago Press, 2010, pp. 283–302.

Denning, Jeffrey T., "Born Under a Lucky Star: Financial Aid, College Completion, Labor Supply, and Credit Constraints," 2016. working paper.

Dunlop, Erin, "What do Stafford Loans Actually Buy You? The Effect of Stafford Loan Access on Community College Students," 2013. CALDER working paper 94.

Dynarski, Susan M. and Judith E. Scott-Clayton, "The Cost Of Complexity In Federal Student Aid: Lessons From Optimal Tax Theory And Behavioral Economics," *National Tax Journal*, 2006, *59* (2), 319–356.

__ **and** __ , "Complexity and Targeting in Federal Student Aid: A Quantitative Analysis," in James M. Poterba, ed., *Tax Policy and the Economy, Volume 22*, MIT Press, 2008, pp. 109–150.

Fan, Jianqing and Irene Gijbels, *Local Polynomial Modelling and its Applications*, London: Chapman and Hall, 1996.

Federal Reserve Bank of New York, "Quarterly Report on Household Debt and Credit, May 2016," Technical Report 2016. New York, NY: Federal Reserve Bank of New York.

Field, Erica, "Educational Debt Burden and Career Choice: Evidence from a Financial Aid Experiment at NYU Law School," *American Economic Journal: Applied Economics*, 2009, *1* (1), 1–21.

Frangakis, Constantine E. and Donald B. Rubin, "Principal Stratification in Causal Inference," *Biometrics*, 2002, *58* (1), 21–29.

Goldrick-Rab, Sara, Douglas N. Harris, Robert Kelchen, and James Benson, "Need-Based Financial Aid and College Persistence: Experimental Evidence from Wisconsin," 2012. Working paper.

Hahn, Jinyong, Petra Todd, and Wilbert Van der Klauuw, "Identification and Estimation of Treatment Effects with a Regression-Discontinuity Design," *Econometrica*, 2001, *69* (1), 201–209.

Imbens, Guido and Karthik Kalyanaraman, "Optimal Bandwidth Choice for the Regression Discontinuity Estimator," *Review of Economic Studies*, 2012, *79* (3), 933–959.

Imbens, Guido W. and Joshua D. Angrist, "Identification and Estimation of Local Average Treatment Effects," *Econometrica*, 1994, *62* (2), 467–475.

Kane, Thomas J., "Rising Public College Tuition and College Entry: How Well Do Public Subsidies Promote Access to College?," 1995. NBER working paper 5164.

Kasy, Maximilian, "Instrumental Variables with Unrestricted Heterogeneity and Continuous Treatment," *Review of Economic Studies*, 2014, *81* (4), 1614–1636.

Kline, Patrick and Melissa Tartari, "Bounding the Labor Supply Responses to a Randomized Welfare Experiment: A Revealed Preference Approach," *American Economic Review*, 2016, *106* (4), 971–1013.

Lee, David S. and Thomas Lemieux, "Regression Discontinuity Designs in Economics," *Journal of Economic Literature*, 2010, *48* (2), 281–355.

Lochner, Lance J. and Alexander Monge-Naranjo, "The Nature of Credit Constraints and Human Capital," *American Economic Review*, 2011, *101* (6), 2487–2529.

Madrian, Brigitte C. and Dennis F. Shea, "The Power of Suggestion: Inertia in 401(k) Participation and Savings Behavior," *Quarterly Journal of Economics*, 2001, *116* (4), 1149–1188.

Pallais, Amanda, "Small Differences that Matter: Mistakes in Applying to College," *Journal of Labor Economics*, 2015, *33* (2), 493–520.

Rothstein, Jesse and Cecilia Elena Rouse, "Constrained After College: Student Loans and Early Career Occupational Choices," *Journal of Public Economics*, 2011, *95(1-2)*, 149–163.

Rubin, Rachel B., "The Pell and the Poor: A Regression-Discontinuity Analysis of On-Time College Enrollment," *Research in Higher Education*, 2011, *52* (7), 675–692.

Scott-Clayton, Judith, "Information Constraints and Financial Aid Policy," in Donald E. Heller and Claire Callender, eds., *Student Financing of Higher Education: A Comparative Perspective*, Routledge, 2013.

Seftor, Neil S. and Sarah E. Turner, "Back to School: Federal Student Aid Policy and Adult College Enrollment," *Journal of Labor Economics*, 2002, *109* (5), 336–352.

Turner, Lesley J., "The Road to Pell is Paved with Good Intentions: The Economic Incidence of Federal Student Grant Aid," 2014. Working paper.

U.S. Department of Education, "2011-2012 Federal Pell Grant Program End-of-Year Report," 2013. Washington DC: U.S. Department of Education, Office of Postsecondary Education.

Wiederspan, Mark, "Denying Loan Access: The Student-Level Consequences When Community Colleges Opt Out of the Stafford Loan Program," *Economics of Education Review*, 2016, *51*, 79–96.

Figures and Tables

Figure 1: The Impact of Pell Grant Aid on Debt by Level of Exogenous Resources

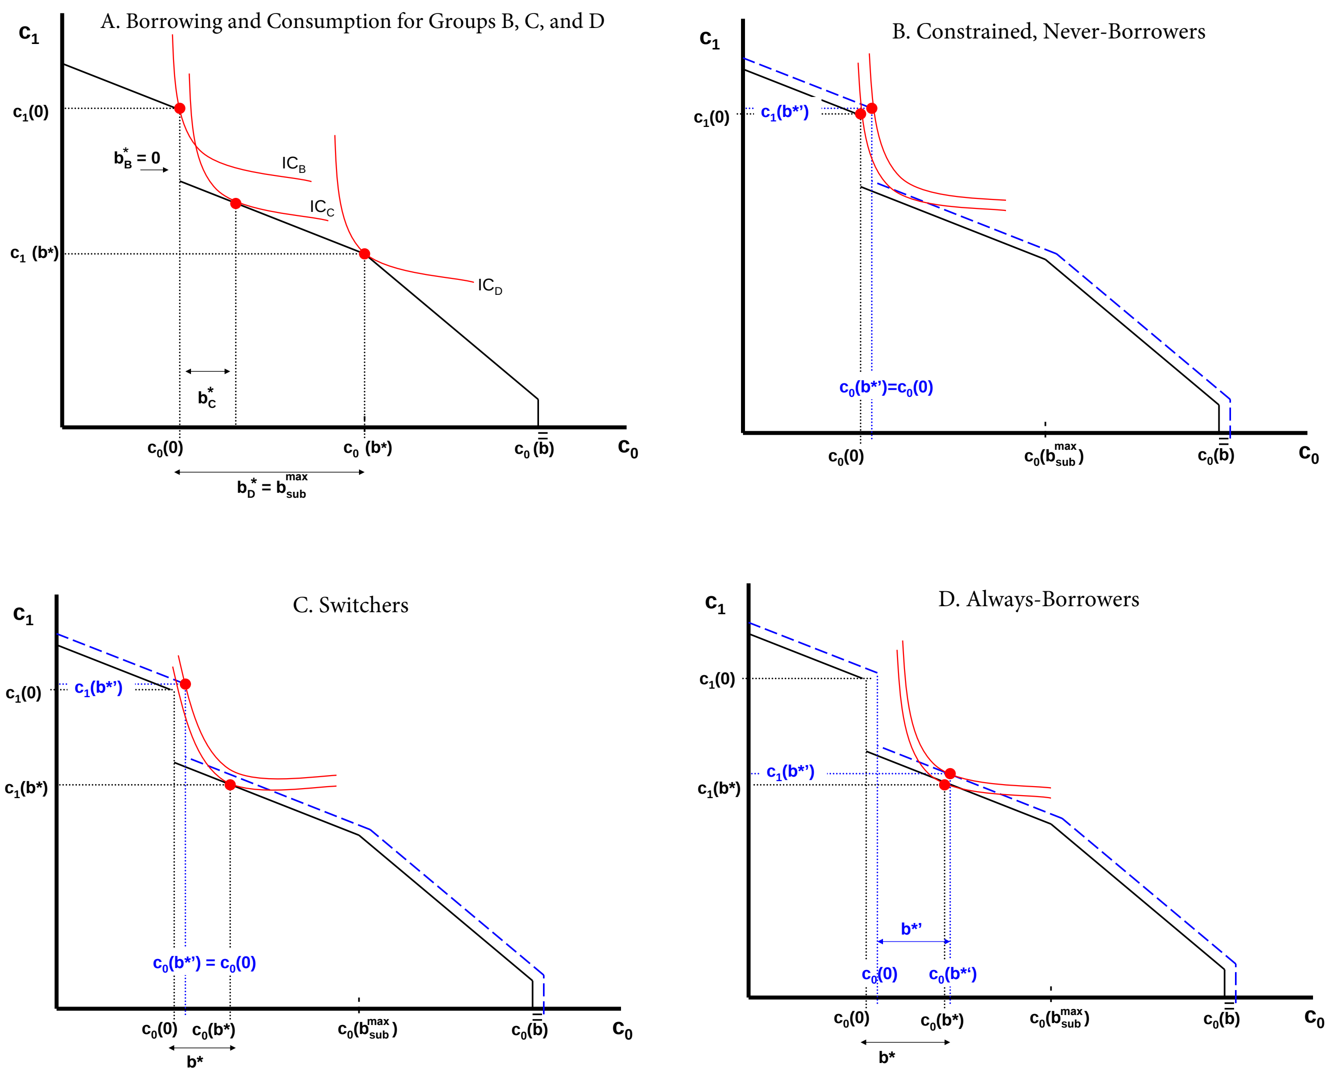

Notes: Black lines represent the student's budget constraint in the absence of the grant aid increase, dashed lines in Panels B through D represent the student's budget constraint upon the receipt of additional grant aid, c_0 is consumption in the first period, and c_1 is consumption in the second period. See Section 4.1 for descriptions of groups. Panel B – $\frac{\partial b}{\partial g} = 0$ and $\frac{\partial s}{\partial g} > 0$. Panel C – $\frac{\triangle b}{\triangle g} < -1$ and $\frac{\triangle s}{\triangle g} < 0$. Panel D – $\frac{\partial b}{\partial g} \in (-1, 0)$ and $\frac{\partial s}{\partial g} = 0$.

Figure 2: The Empirical Distribution of Pell Grant Aid by Distance to Eligibility Threshold

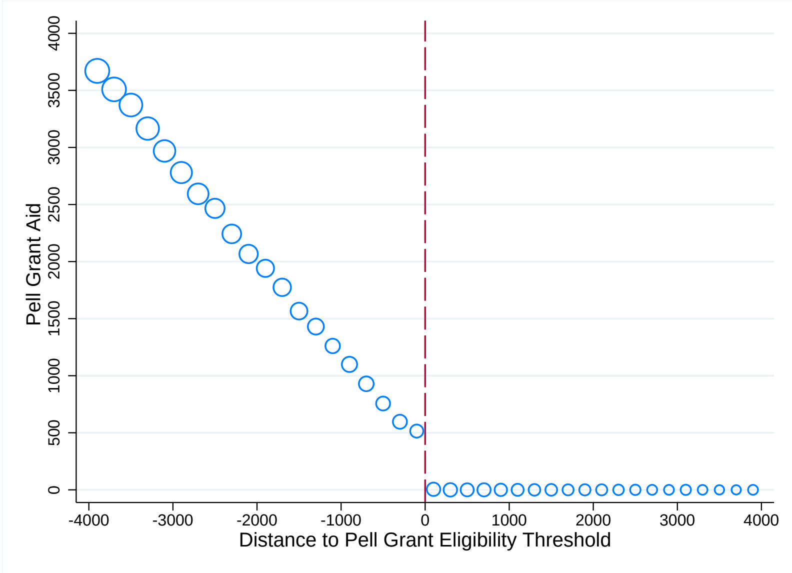

Notes: First-year CUNY undergraduate degree seeking students; 2007 through 2011 entry cohorts. $200 EFC bins. Each circle represents the average Pell Grant aid received by students in the bin. Larger circles represent a larger underlying sample size. All dollar amounts adjusted to represent constant 2012$.

Figure 3: The Density of EFC and Probability of Enrollment at the Pell Grant Eligibility Threshold

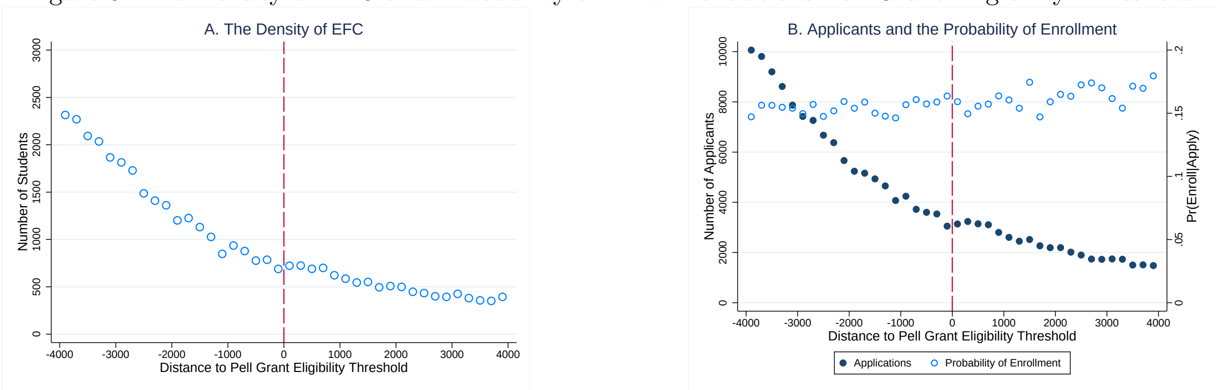

Notes: Panel A includes first-year CUNY undergraduate degree seeking students; 2007 through 2011 entry cohorts. Panel B includes first-year CUNY undergraduate degree seeking applicants; 2008 to 2011 cohorts. $200 EFC bins. In Panel A, each circle represents the total number of students in the bin. In Panel B, solid circles represent the total number of applicants in the bin while hollow circles represent the probability of enrollment conditional on submitting an application for students in the bin. All dollar amounts adjusted to represent constant 2012$.

Figure 4: The Distribution of Baseline Characteristics

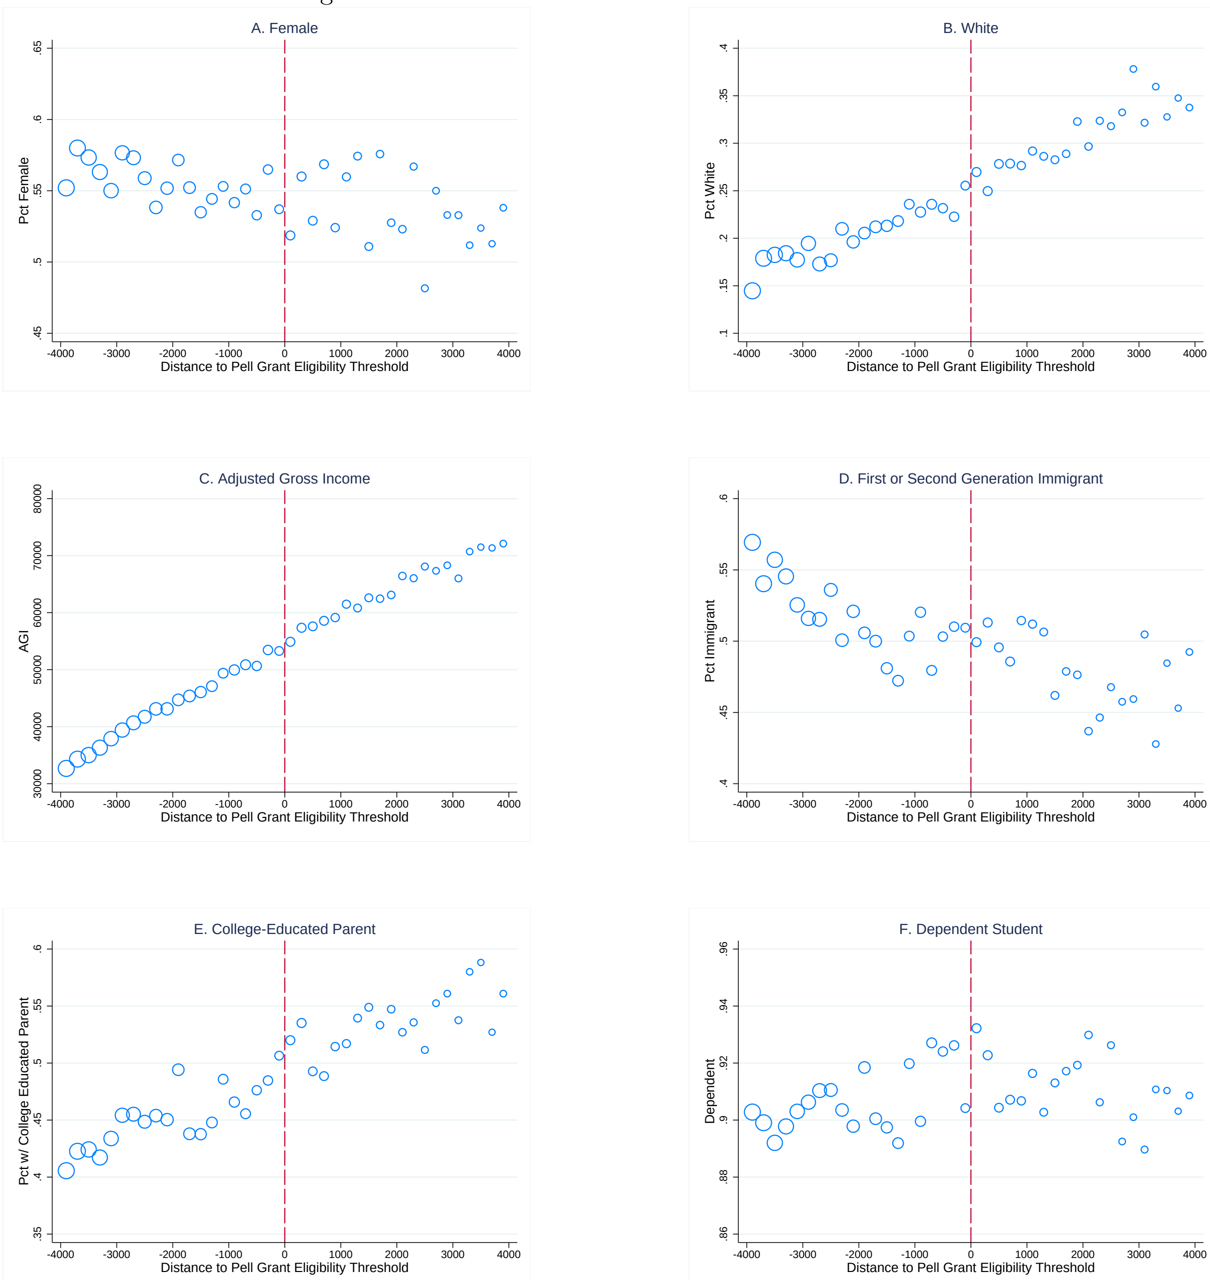

Notes: First-year CUNY undergraduate degree seeking students; 2007 through 2011 entry cohorts. $200 EFC bins. Each circle represents the average characteristic of students in the bin. Larger circles represent a larger underlying sample size. All dollar amounts adjusted to represent constant 2012$.

Figure 5: The Reduced Form Impact of Pell Grant Eligibility and Generosity on Borrowing Outcomes

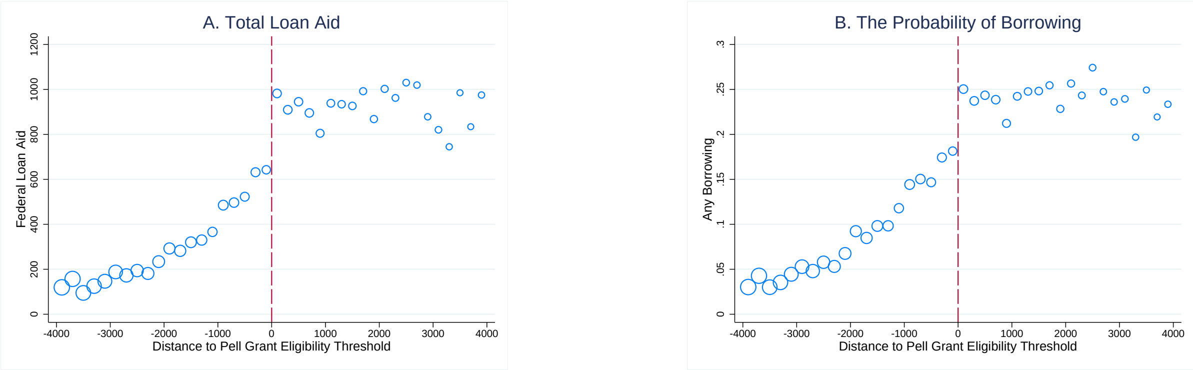

Notes: First-year CUNY undergraduate degree seeking students; 2007 through 2011 entry cohorts. $200 EFC bins. Each circle represents average amount borrowed (Panel A) or the average probability of borrowing (Panel B) for students in bin. Larger circles represent a larger underlying sample size. All dollar amounts adjusted to represent constant 2012$.

Figure 6: The Distribution of Loans

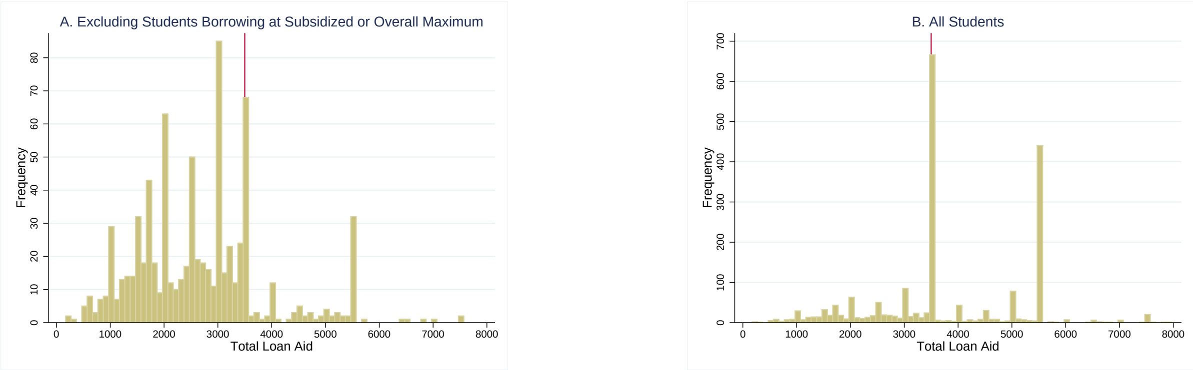

Notes: First-year CUNY undergraduate degree seeking students with EFCs less than $4,000 from the Pell Grant eligibility threshold who are subject to the exogenous subsidized borrowing limit; 2007 through 2011 entry cohorts. $100 bins. Dollar amounts in nominal terms.

Figure 7: Empirical CDFs of Student Loans by Pell Grant Eligibility

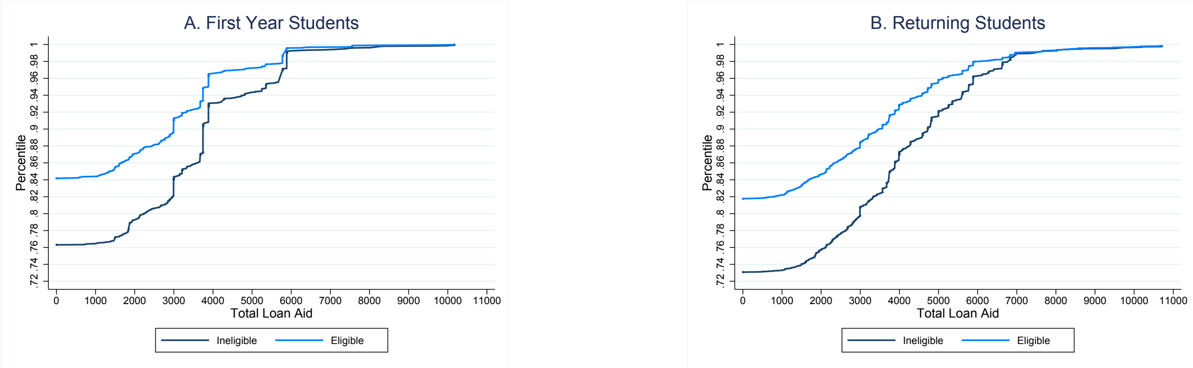

Notes: First-year CUNY undergraduate degree seeking students, 2007 through 2011 entry cohorts (Panel A) or second- and third-year CUNY undergraduate degree seeking students, 2005 through 2010 entry cohorts (Panel B), with EFCs less than $1,000 from the Pell Grant eligibility threshold. Borrowing is zero in percentiles that are not listed. All dollar amounts adjusted to represent constant 2012$.

Table 1: Characteristics of First-Year Students by Pell Grant Eligibility

	(1) Ineligible	(2) Eligible	(3) Full sample
A. Cost of attendance and financial aid			
Expected family contribution (EFC)	$6,451	$2,254	$3,381
Total need (= cost of attendance - EFC)	$6,772	$10,406	$9,430
Total grant aid	$1,012	$4,313	$3,411
Pell Grant aid	$0	$2,394	$1,751
TAP Grant aid	$753	$1,573	$1,352
Unmet need (= COA - EFC - grants)	$5,739	$6,054	$5,969
Any borrowing?	0.24	0.07	0.12
Borrowing at subsidized limit | any loan	0.77	0.60	0.69
Borrowing at overall limit | any loan	0.04	0.01	0.02
Subject to endogenous limit	0.38	0.33	0.34
Subsidized borrowing limit	$2,836	$2,936	$2,909
Total loan aid	$923	$244	$427
Share subsidized	0.73	0.80	0.76
B. Student demographic characteristics			
Female	0.54	0.56	0.55
Dependent student	0.91	0.90	0.91
Black	0.30	0.34	0.33
Hispanic	0.27	0.33	0.32
White	0.30	0.19	0.22
SAT percentile	38.8	32.4	34.2
Foreign-born	0.15	0.19	0.18
Foreign-born parent(s)	0.41	0.47	0.45
Parents' highest education			
Less than high school	0.04	0.06	0.06
High school	0.37	0.40	0.40
College	0.53	0.45	0.47
Parents' resources			
Adjusted gross income	$64,405	$42,522	$48,434
Savings	$6,860	$3,794	$4,618
Student's resources			
Adjusted gross income	$4,422	$3,044	$3,414
Savings	$498	$305	$357
Initial degree program = BA/BS	0.44	0.35	0.37
Number of students	10,231	27,869	38,100

Notes: First-year CUNY undergraduate degree seeking students; 2007 through 2011 entry cohorts. COA = total cost of attendance (equal to tuition and fees, books and supplies, and living expenses). Race and parental education categories may not sum to one due to missing values. Mean SAT percentile calculated for students with nonmissing SAT math and verbal scores ($N = 24,760$). Students with EFC greater than $4,000 from Pell Grant eligibility threshold are excluded. All dollar amounts adjusted to represent constant 2012$.

Table 2: Optimal Borrowing and Educational Investment Decisions by Level of Exogenous Resources

Group	Never borrowers		Switchers	Always borrowers			
	A	B	B/C	C	D	E	F
b^*	$(-\infty, 0)$	0	$\underline{b}$	$(\underline{b}, b_{sub}^{max})$	b_{sub}^{max}	$\left(b_{sub}^{max}, \bar{\bar{b}}\right)$	$\bar{\bar{b}}$
$\frac{\partial b^*}{\partial g}$	$(-1, 0)$	0	$\frac{\Delta b^*}{\Delta g} = \frac{0-\underline{b}}{\Delta g} < -1$	$(-1, 0)$	$\xi \left(\frac{\partial s^*}{\partial g} T_d' \left(s^* \vert EFC\right) - 1\right)$	$(-1, 0)$	0
$\frac{\partial s^*}{\partial g}$	0	$(0, \infty)$	$\frac{\Delta s^*}{\Delta g} = \frac{s_0 - \underline{s}}{\Delta g} < 0$	0	$(0, \infty)$	0	$(0, \infty)$

Notes: Groups are listed in decreasing order of exogenous resources ω, where group A has the highest resources and group F has the lowest resources. Observed debt is bounded from below by 0 and $b^* < 0$ implies saving.

Table 3: Impact of Pell Grant Eligibility on Pell Grant Aid

	(1) First-year students	(2) Returning students
Pell Grant eligible	388.69	349.15
	(27.60)**	(24.19)**
× Distance from threshold	-0.761	-0.752
	(0.020)**	(0.019)**
Observations	38,100	46,744

Notes: OLS estimates of the impact of Pell Grant eligibility on Pell Grant aid. Sample includes first-year CUNY undergraduate degree seeking students, 2007 through 2011 entry cohorts (column 1) or second- and third-year CUNY undergraduate degree seeking students, 2005 through 2010 entry cohorts (column 2). Each column contains estimates from a separate regression. Clustered standard errors (institution by year) in parentheses; ** p<0.01, * p<0.05, + p<0.1. All regressions include controls for age, family AGI, and indicators for race (white versus nonwhite), dependency status, parents' highest level of education (college, high school, less than high school, or missing), level of attendance (for federal loan eligibility purposes), degree program, school by cohort fixed effects, years since entry fixed effects, and a quadratic in student expected family contribution ($\widetilde{EFC}_{it} = EFC_{it} - efc_t^0$, where efc_t^0 is the threshold for Pell Grant eligibility in year t), allowed to vary on either side of the eligibility threshold. Students with EFC greater than $4,000 from Pell Grant eligibility threshold are excluded. All dollar amounts adjusted to represent constant 2012$.

Table 4: The Relationship between Pell Grant Aid and the Probability of Enrollment

	Enrollment | application			
	(1) Pooled	(2) Community colleges	(3) 4-year colleges	(4) Any enrollment | any application
A. OLS estimates: impacts on Pr(Enroll)				
Pell Grant eligible	-0.002	0.011	-0.005	0.002
	(0.004)	(0.007)	(0.005)	(0.010)
× Distance from threshold	-0.000002	-0.000006	-0.0000004	-0.000009
	(0.000002)	(0.000004)	(0.000002)	(0.000005)+
Mean | Pell Grant ineligible	0.162	0.236	0.141	0.629
Observations	161,841	39,056	122,785	38,971
B. 2SLS estimates: impacts on Pr(Enroll)				
Federal grant aid ($1k)	0.003	0.009	0.0002	0.014
	(0.003)	(0.006)	(0.003)	(0.007)+
95% CI	[-0.003, 0.008]	[-0.003, 0.021]	[-0.005, 0.006]	[-0.0001, 0.028]
Observations	161,841	39,056	122,785	38,971

Notes: CUNY undergraduate degree seeking applicants; fall 2007 through fall 2010 cohorts. The sample in Columns 1 through 3 includes one observation per prospective student-application. The Column 4 sample includes one observation per applicant. Each column within a panel contains estimates from a separate regression. Clustered standard errors (institution by year) in parentheses; ** p<0.01, * p<0.05, + p<0.1. All regressions also include school by year fixed effects, $\widetilde{EFC}_{it}$, an indicator for Pell Grant eligibility $\mathbf{1}[\widetilde{EFC}_{it} < 0]$, $\widetilde{EFC}_{it} \times \mathbf{1}[\widetilde{EFC}_{it} < 0]$, and school ranking fixed effects (Columns 1 through 3 only). Excluded instruments are $\mathbf{1}[\widetilde{EFC}_{it} < 0]$ and $\widetilde{EFC}_{it} \times \mathbf{1}[\widetilde{EFC}_{it} < 0]$. Students with EFC greater than $4,000 from Pell Grant eligibility threshold are excluded. All dollar amounts adjusted to represent constant 2012$.

Table 5: The Impact of Pell Grant Aid on Borrowing

	(1) First-year students	(2) Returning students
A. OLS estimates		
Pell Grant eligible	-224.45	-273.99
	(53.76)**	(55.88)**
× Distance from threshold	0.295	0.335
	(0.075)**	(0.073)**
Observations	38,100	46,744
B. 2SLS estimates: combined RD/RK		
Pell Grant aid	-0.428	-0.508
	(0.092)**	(0.095)**
Observations	38,100	46,744

Notes: Column 1 sample includes first-year CUNY undergraduate degree seeking students; 2007 through 2011 entry cohorts. Column 2 sample includes second- and third-year CUNY undergraduate degree seeking students; 2005 through 2010 entry cohorts. Each column contains estimates from a separate regression. Clustered standard errors (institution by year) in parentheses; ** p<0.01, * p<0.05, + p<0.1. See Table 3 notes for a list of additional controls. Excluded instruments for Panel B regressions are $\mathbf{1}[\widetilde{EFC}_{it} < 0]$ and $\widetilde{EFC}_{it} \times \mathbf{1}[\widetilde{EFC}_{it} < 0]$. F-stat from test of significance of excluded instruments: 917 (first-year student sample), 854 (returning student sample). Students with EFC greater than $4,000 from Pell Grant eligibility threshold are excluded. All dollar amounts adjusted to represent constant 2012$.

Table 6: The Impact of Pell Grant Aid on Borrowing Among Would-be Borrowers

	(1) First-year students	(2) Returning students
Pell Grant aid	-1.821	-1.892
	(0.281)**	(0.250)**
H_0: crowd-out > -1, *p* -value	0.002	<0.001
Observations	38,100	46,744

Notes: 2SLS estimates of the impact of an additional dollar of Pell Grant aid on borrowing among would-be borrowers (see Section 5 for estimation details). Column 1 sample includes first-year CUNY undergraduate degree seeking students; 2007 through 2011 entry cohorts. Column 2 sample includes second- and third-year CUNY undergraduate degree seeking students; 2005 through 2010 entry cohorts. Each column contains estimates from a separate specification. Bootstrapped standard errors (clustered at institution by year); ** p<0.01, * p<0.05, + p<0.1. Students with EFC greater than $4,000 from Pell Grant eligibility threshold are excluded. All dollar amounts adjusted to represent constant 2012$.

Table 7: The Impact of Pell Grant Aid on Borrowing: Robustness Tests

	(1) First-year students	(2) Returning students
A. Accounting for other grant aid		
Pell + other grant aid	-0.386	-0.523
	(0.079)**	(0.097)**
Crowd-out | borrower	-1.617	-1.933
	(0.232)**	(0.253)**
H_0: crowd-out >-1, *p*- value	0.004	<0.001
Observations	38,100	46,744
B. Excluding covariates		
Pell Grant aid	-0.415	-0.540
	(0.096)**	(0.099)**
Crowd-out | borrower	-1.766	-2.059
	(0.306)**	(0.247)**
H_0: crowd-out >-1, *p*- value	0.006	<0.001
Observations	38,100	46,744
C. Eliminating mechanical effect on subsidized loan eligibility		
Pell Grant aid	-0.438	-0.558
	(0.144)**	(0.157)**
Crowd-out | borrower	-1.634	-1.696
	(0.418)**	(0.423)**
H_0: crowd-out >-1, *p*- value	0.064	0.050
Observations	23,762	24,191

Notes: 2SLS estimates of the impact of an additional dollar of Pell Grant aid on borrowing for all students and would-be borrowers. Column 1 sample includes first-year CUNY undergraduate degree seeking students; 2007 through 2011 entry cohorts. Column 2 sample includes second- and third-year CUNY undergraduate degree seeking students; 2005 through 2010 entry cohorts. Each point estimate within a panel contains estimates from a separate specification. Clustered and bootstrapped standard errors (institution by year) in parentheses; ** p<0.01, * p<0.05, + p<0.1. In Panel A, we instrument for total grant aid rather than Pell Grant aid. In Panel C, we limit the sample to students for whom receipt of the Pell Grant does not or would not affect eligibility for subsidized loans (COA less EFC and other grant aid is greater than the sum of the minimum Pell Grant and the statutory maximum subsidized loan). Panel A and B regressions also include controls specified in Table 3 notes. Excluded instruments are $\mathbf{1}[\widetilde{EFC}_{it} < 0]$ and $\widetilde{EFC}_{it} \times \mathbf{1}[\widetilde{EFC}_{it} < 0]$. Students with EFC greater than \$4,000 from Pell Grant eligibility threshold in their first year are excluded. All dollar amounts adjusted to represent constant 2012\$.

Table 8: The Impact of Pell Grant Aid on Educational Attainment

	First-year students		Returning students
	(1) Current	(2) Cumulative	(3) Current
A. Reenrollment			
Pell Grant aid ($1k)	0.012	-0.002	0.015
	(0.020)	(0.023)	(0.015)
	[-0.03, 0.05]	[-0.05, 0.04]	[-0.01, 0.04]
Mean | Pell Grant ineligible	0.79	0.67	0.71
Observations	38,100	32,271	46,744
B. Credits attempted (academic + remedial)			
Pell Grant aid ($1k)	0.490	0.539	0.450
	(0.266)+	(1.190)	(0.341)
	[-0.03, 1.01]	[-1.79, 2.87]	[-0.23, 1.11]
Mean | Pell Grant ineligible	25.5	59.8	24.5
Observations	38,100	32,271	46,744
C. Credits earned (academic only)			
Pell Grant aid ($1k)	0.212	0.223	0.492
	(0.410)	(1.233)	(0.372)
	[-0.59, 1.02]	[-2.19, 2.64]	[-0.24, 1.22]
Mean | Pell Grant ineligible	17.6	44.7	19.8
Observations	38,100	32,271	46,744
D. Cumulative grade point average			
Pell Grant aid ($1k)	-0.025	--	0.003
	(0.035)	--	(0.032)
	[-0.09, 0.04]	--	[-0.06, 0.07]
Mean | Pell Grant ineligible	2.65	--	2.75
Observations	34,203	--	44,231

Notes: 2SLS estimates of the impact of an additional $1,000 of Pell Grant aid on the outcome specified in a given panel in the current year of school (Columns 1 and 3) or three years after entry (Column 2). The sample in Columns 1 and 2 includes first-year CUNY undergraduate degree seeking students; 2007 through 2011 entry cohorts. Column 3 sample includes second- and third-year CUNY undergraduate degree seeking students; 2005 through 2010 entry cohorts. Each column within a panel represents estimates from a separate regression. Clustered standard errors (institution by year) in parentheses; ** p<0.01, * p<0.05, + p<0.1. The corresponding 95 percent confidence interval is displayed below each point estimate in brackets. "Reenrollment" indicates the probability of re-enrolling the following year. Credits attempted represents academic and remedial course credit equivalents. Students do not earn credits for remedial courses, but CUNY converts the hours students spend in such courses into credit equivalents. Grade point average is measured on a four-point scale. See Table 3 notes for a list of controls. Excluded instruments are $\mathbf{1}[\widetilde{EFC}_{it} < 0]$ and $\widetilde{EFC}_{it} \times \mathbf{1}[\widetilde{EFC}_{it} < 0]$. Students with EFC greater than $4,000 from Pell Grant eligibility threshold are excluded. All dollar amounts adjusted to represent constant 2012$.

Table 9: The Impact of Pell Grant Aid and Borrowing on Educational Attainment: Estimates from Student Fixed-Effects Models

	(1) Credits attempted	(2) Credits Earned
Pell Grant Aid ($1k)		
x **1**[No loans]	0.457	0.395
	(0.214)*	(0.225)+
x **1**[Any loans]	0.285	0.225
	(0.198)	(0.208)
1[No loans]	-1.820	-1.014
	(0.189)**	(0.194)**
Observations	48,496	48,496

Notes: 2SLS estimates of the impact of an additional $1,000 of Pell Grant aid on the specified outcome. First-, second-, and third-year CUNY undergraduate degree seeking students with an EFC within $4,000 of the Pell Grant eligibility threshold and at least two years of enrollment; 2006 through 2010 entry cohorts. Each column contains estimates from a separate regression. Clustered standard errors (by student) in parentheses; ** p<0.01, * p<0.05, + p<0.1. Credits attempted include both academic and remedial credit equivalents. All regressions include student fixed effects, a cubic in EFC, age, AGI, level, and year and years since entry fixed effects. The excluded instruments are predicted Pell Grant aid assuming full-time, full-year attendance fully interacted with borrowing indicators. All dollar amounts adjusted to represent constant 2012$.

Table 10: The Impact of Pell Grant Aid on Borrowing by Default Loan Offer: NPSAS Sample

	1. 2SLS Estimates		2. Crowd-out | would-be borrower	
Zero loan default?	*N*	*Y*	*N*	*Y*
Pell Grant aid	-0.397	-0.605	-0.534	-1.536
	(0.123)**	(0.152)**	(0.159)**	(0.291)**
H_0: crowd-out > -1, *p*- value	--	--	--	0.033
	[0.080]		[0.008]	
Observations	5,410		5,410	

Notes: First-, second-, and third-year NPSAS undergraduate students attending public schools; 2010 through 2012 cohorts. Clustered standard errors (by school) in parentheses; ** p<0.01, * p<0.05, + p<0.1. 2SLS impacts of Pell Grant aid on total loan aid. All regressions include $\widetilde{EFC}$, an indicator for whether the school makes a zero loan default offer, degree program indicators interacted with zero loan default indicator, dependency status, AGI, and years since entry fixed effects. Excluded instruments are $\mathbf{1}[\widetilde{EFC}_{it} < 0]$ and $\widetilde{EFC}_{it} \times \mathbf{1}[\widetilde{EFC}_{it} < 0]$, fully interacted with the zero loan default indicator. Students with EFC greater than $4,000 from the Pell Grant eligibility threshold and those attending community colleges for which we were unable to the default loan amount are excluded. Sample sizes rounded to the nearest 10 per NCES requirements. All dollar amounts adjusted to represent constant 2012$.