The Distribution of College Graduate Debt, 1990 to 2008: A Decomposition Approach

Brad Hershbein and Kevin M. Hollenbeck

October 2013

Prepared for Upjohn/EPI/Spencer Conference on Student Loans, Ann Arbor, MI, October 24–26, 2013

Abstract:

Despite tremendous recent interest on the subject of student debt by both researchers and policymakers, little is known about how the distribution of college graduate debt has been evolving and what factors can explain it. We use National Postsecondary Student Aid Survey data from 1990 through 2008 to document the evolution of college graduate debt profiles. We find that growth in debt over the 1990s was rapid and occurred throughout the distribution; during the 2000s, in contrast, debt grew appreciably only for the top quartile. Employing several decomposition techniques, we exploit the richness of the data to explain these shifts. Over the entire horizon, observable characteristics of students and institutions explain about one-third of the debt increase, though this share tends to be higher around the extensive margin and the median and lower in the right tail. While observables, largely costs, explain a majority of the increase between 1990 and 1996 and again from 2000 to 2008, they explain nothing over the late 1990s. We offer suggestive evidence that this "unobservable" share was supply-side driven due to the advent of both federal unsubsidized Stafford loans and private loans.

Contact Information:

Hershbein is an economist at the W.E. Upjohn Institute for Employment Research, Kalamazoo, MI 49007. Email: hershbein@upjohn.org. Hollenbeck is a vice president and senior economist at the W.E. Upjohn Institute for Employment Research, Kalamazoo, MI 49007. Email: hollenbeck@upjohn.org.

Acknowledgements:

We thank Tim Bartik, Marcus Dillender, Randy Eberts, George Erickcek, Susan Houseman, Marta Lachowska, Michelle Miller-Adams, Chris O'Leary, and Steve Woodbury for helpful comments and suggestions.

Within the past 21 months, there have been almost as many major news articles on the topic of student debt as there were in the preceding 21 years.[1] Headlines have trumpeted the stories of recent graduates with six-figure debt levels and aggregate loan balances exceeding $1 trillion.[2] A growing number of policy organizations and web sites have begun to focus and compile information on student debt.[3] And, perhaps in response, President Obama announced in August of 2013 a major initiative to address growing college costs.[4] Despite this increased attention on debt writ large, surprisingly little is known about how student debt has changed for different types of students or what factors can explain it.

Understanding the patterns and factors underlying debt increase is paramount, both for ensuring that students and their families have a realistic, well-informed picture of college finance, and for guiding policymakers toward debt-amelioration strategies for those who need them most. Considerable focus is spent on average debt levels, because these are easy to update frequently, but this may be misguided. The distribution of debt is so diffuse, as we show, that changes in the mean are not informative for most students. For example, an increase in borrowing among the top 10 percent of borrowers will increase the mean and total accumulated debt—numbers commonly reported—but leave debt levels for 90 percent of students unchanged. This may seem a convenient hypothetical, but it actually closely resembles how debt evolved for college graduates between 2000 and 2008. In short, factors that influence debt for the median borrower may be quite different than those that influence debt for borrowers in the highest decile. The effectiveness of policy proposals meant to address "rising student debt" rests on how they recognize these phenomena.

In this paper, we use the National Postsecondary Student Aid Survey (NPSAS) to investigate the growing student debt of college graduates—those earning a bachelor's degree. Our focus on this segment

[1] From January 1, 2012, through September 16, 2013, we found 504 hits of the search term "student debt" across the *New York Times*, the *Los Angeles Times*, the *Washington Post*, the *Wall Street Journal*, *USA Today*, *Time*, *Newsweek*, and *U.S. News and World Report*. For the same outlets, there were 673 hits between 1991 and 2011, inclusive.

[2] See Martin and Lehren (2012) and Mitchell and Jackson-Randall (2012).

[3] For example, collegerealitycheck.com and collegecompletion.chronicle.com (both by the *Chronicle for Higher Education*), collegecost.ed.gov/scorecard/ and collegecost.ed.gov/catc/ (both by the federal government), projectonstudentdebt.org (by The Institute for College Access and Success), and collegeportraits.org (by the American Association of State Colleges and Universities.)

[4] See http://www.whitehouse.gov/the-press-office/2013/08/22/remarks-president-college-affordability-buffalo-ny.

of students is not because other groups (graduate students, students with sub-baccalaureate degrees, students leaving without any credential) have not experienced rising debt—they have—or because they are unimportant. Rather, the choice is motivated by the recent media interest in college graduates and to limit the analysis to a manageable scale.[5] The data cover the period between 1990 and 2008, and our analysis examines the entire distribution of borrowing among graduates. In addition to documenting how the distribution has changed over time, for all graduates and subgroups, we employ statistical decompositions to apportion the changes by various observable characteristics, such as demographics, attendance patterns, incomes, and costs. Notably, the decompositions allow the role of observable factors to vary at different points of the distribution.

In broad terms, our major findings are that debt profiles increased much more in the 1990s than in the 2000s, with the largest part of this increase occurring in the latter part of the decade. The growth occurred throughout the distribution. Between 2000 and 2008, debt increases were concentrated almost entirely in the top quartile of the distribution and at private institutions, and stemmed largely from the expansion of private borrowing. About one-third of the overall increase in debt at graduation between 1990 and 2008 is explained by observable characteristics of the students and the schools they attend. Interestingly, we find that observables explain more of the increase at the extensive margin (whether a student ever borrowed) and around the median than they do near the top of the distribution. Of the explained share, roughly half can be attributed to college costs alone, although this still implies that costs account for a small fraction of the total increase in borrowing.

When we look at intermediate time intervals, observables explain most of the increase in borrowing—between 50 to 100 percent, or more—from 1990 to 1996 and 2000 to 2008. While cost structure plays an important role, so do other factors, and again there is a greater role for unexplained factors in the upper tail of the distribution. In contrast, observables explain practically none of the debt increase between 1996 and 2000, and this is true throughout the entire distribution.

[5] We hope to investigate other groups in subsequent work.

We investigate several possibilities that might show up as "unexplained" factors in driving debt increases in the upper tail and throughout the distribution in the late 1990s, including movement from informal to formal loans, redistribution of debt from parents to students, variation in interest rates, increases in federal borrowing limits, the introduction of unsubsidized loans (which are not means-tested), and the growing market for private loans. While the first four of these appear to be unimportant, we find suggestive evidence that the latter two may play a prominent role.

We note a number of caveats to our analyses. First, we are not able at this time to examine the distributional changes, if any, that occurred as a result of the Great Recession starting in December 2007.[6] As documented by the College Board (Baum and Payea 2012), aggregate borrowing increased significantly during and after this recession. A second caveat is that the NPSAS data contain information only on current and graduating students. As such, they do not contain data on post-graduation labor market experiences or repayment information, and our paper cannot consider these important outcomes. Third, while the data are quite rich in detail, they do not contain information on previous institutions attended (and the net costs thereof), so differences across students (and over time) in transfer behavior are not captured in the analysis. Finally, the data do not fully document alternative loan sources, including informal loans from friends and family or borrowing against existing assets. Nonetheless, our paper is the first to investigate distributional changes in borrowing over time and link these to changing characteristics of student attendance.

The plan of the paper is as follows. The next section discusses the data source and presents descriptive statistics and distributions from it. We then review the decomposition methods that are used to explain the reasons behind borrowing trends. The results of those decompositions are then presented, followed by a discussion of possible factors that could account for the unexplained portion of the decompositions. Finally, we offer concluding remarks. Two appendices to the paper describe the data

[6] The 2012 wave of NPSAS was released in late August of 2013. While we have requested these data, our analyses are currently limited to 2008 as an endpoint.

processing in detail and provide an overview of the market for student borrowing, including the structure of terms and borrowing limits, from approximately 1980 through today.

The NPSAS and Descriptive Statistics

The National Postsecondary Student Aid Survey (NPSAS) is an approximately quadrennial survey of students attending Title IV institutions (those eligible for federal aid) that is conducted by the National Center for Education Statistics. The nationally representative cross-sectional survey is designed specifically to gather information on how different students pay for higher education.[7] It provides student-level information on financial aid provided by the federal government, the states, postsecondary institutions, employers, and private agencies, along with student demographics and enrollment characteristics. The restricted-use version we employ has incredibly rich detail, including administrative data on student financial aid programs merged from both the Free Application for Federal Student Aid (FAFSA) and the National Student Loan Data System (NSLDS), the central database for all federal loans. Extensive data about family circumstances, demographics, education and work experiences, and student expectations are collected from students through an interview.

The survey waves are reasonably consistent over time, which is important for our analyses of the debt burdens of graduating seniors from five waves: 1990, 1996, 2000, 2004, and 2008 (the most recent available).[8] The richness of the data is important because our goals are to understand why student debt is growing and for whom. Having cross-sections with large sample sizes and spanning almost two decades allows us to examine the growth in student debt over the entire distribution of college graduates. This allows far more nuanced analyses than are possible by examining means or population totals. (More information about the NPSAS and how we process the data for analysis can be found in Appendix A.)

Descriptive Measures

[7] It also forms the sampling frame for two longitudinal surveys: the Beginning Postsecondary Students study (BPS), which tracks first-time students, and the Baccalaureate and Beyond study (B&B), which tracks bachelor recipients. We use the NPSAS rather than the B&B because the former is available earlier and more often.

[8] The NPSAS was also fielded in 1993. However, an error in the survey led to many bachelors degree graduates not being asked the relevant questions on cumulative borrowing, forcing us to exclude this wave.

Table 1 summarizes changes in the distribution of student borrowing at graduation. There were substantial increases over the last two decades in both the rate of borrowing for bachelor's degree earners and in the real levels of borrowing. Between 1990 and 2008, the fraction of graduates who borrowed increased by 20 percent, from about 55 percent to over 66 percent, with the sharpest increase occurring between 1996 and 2000.

The next panel of the table lists quantiles of borrowing for all graduates, including those who did not borrow at all, in constant dollars. Mean levels have more than doubled over the 18-year horizon, with more than $5,000 of the total $9,500 increase coming between the classes of 1996 and 2000. Since 2000, however, average debt has increased more modestly. This trend of rapid debt increase over the 1990s and milder increases over the 2000s is apparent through at least the 75th percentile. Only in the extreme right tail, above the 95th percentile, has borrowing continued to grow as quickly as it did in the 1990s.

The last panel shows that, among borrowers, the median level of borrowing more than doubled between 1990 and 2000, from about $10,400 to just under $22,000. However, that level remained stable between 2000 and 2008. On the other hand, the mean level of borrowing increased between 2000 and 2008 as the individuals in the upper tail of the borrowing distribution significantly increased their levels of borrowing. Even so, while there have been media suggestions of individuals graduating with six-figure levels of debt, these data suggest that such instances are quite rare, as the 99th percentile of borrowers did not reach that level of borrowing in any of the waves (see also Kantrowitz 2012).

Table 1 thus illustrates two facts that are not well known in either the academic or popular press. First, debt at graduation increased much faster between 1990 and 2000 than it did during 2000 and 2008, and this was true throughout the distribution. Second, the increase in borrowing in the later period was entirely concentrated in the top quartile; the bottom 75 percent of graduates of the class of 2008 had roughly the same debt as the class of 2000 and 2004. These facts can perhaps be more directly seen in Figure 1, which displays the cumulative borrowing distributions of graduates from each wave of the NPSAS in constant dollars. The remainder of this paper seeks to gain understanding of the factors that

shifted the borrowing distribution so dramatically between 1990 and 2000, and the factors that shifted the upper tail of the distribution between 2000 and 2008.

Our first analysis is entirely descriptive and is meant to isolate changes in debt among certain subgroups. We examine four of these: dependent versus independent students; public versus private, not-for-profit institutions[9]; graduates who took four or fewer years to degree versus those who took five or more years[10]; and, for dependent students, those whose family income is above versus below the median. Figure 2 displays the cumulative borrowing distributions for dependent and independent students. As would be expected, the proportion of dependent students with loans is smaller than the proportion of independent students with loans, and the levels of borrowing are much smaller for dependent students in all years of the data. While the distributions for dependent students are qualitatively quite similar to the overall distributions (figure 1), the distributions for independent students show a relatively smooth, monotonically increasing pattern over time. That is, each wave's distribution (first-order) stochastically dominates the preceding wave, which is not at all true for dependent students, whose debt profile in 2004 is smaller than in 2000 for the middle segment of the distribution. Nonetheless, both graphs show substantially larger increases in debt over the 1990s than over the 2000s.

Figure 3 shows the comparison between public and private, not-for-profits. It is not surprising that both the propensity to ever borrow and borrowing levels are greater in all time periods at the latter institutions. What is less well known is the remarkable similarity in the debt profiles at publics across the 2000, 2004, and 2008 waves. Aside from a slight increase between 2004 and 2008 in the upper tail, balanced by a modest decrease between 2000 and 2004 in the middle, the distributions almost lie on top of one another. Although much has been made of the decline in state-level appropriations to public universities during the Great Recession (Lewin 2013), these appropriations had actually fallen by about 15 percent per full-time equivalent student between 2000 and 2004, and still remained below 2000 levels

[9] Unfortunately, private, for-profit institutions are too small a group to examine in all but the 2008 wave. Even then, fewer than 4 percent of bachelor's recipients graduated from a private, for-profit institution.
[10] Bound, Lovenheim, and Turner (2010) document rising time to bachelor's degree and attempt to explain the factors behind it.

at the start of the recession (Baum and Ma 2012a, Figure 12B; Quinterno and Orozco 2012, Figure 6). Despite these reductions (and concomitant tuition increases), debt of graduating students changed little, especially relative to the large increases over the 1990s. On the other hand, while debt also increased little over the 2000s for the bottom 60 percent at privates, it increased substantially for the top 40 percent, with the size of increase rising with the distribution quantile. Above the 80th percentile, the debt increase between 2000 and 2008 was comparable to that between 1990 and 2000.

Turning to Figure 4, we look at debt distributions by time to degree, although data limitations restrict the analysis to the 1996 and later waves. Students who take longer to finish accumulate more debt, almost mechanically, and as in the previous groups, debt increases faster between 1996 and 2000 than it does over the following eight years. However, for students who graduate on time, debt actually *fell* between 2000 and 2008 through the 80th percentile while rising, often considerably, above that quantile. Among the students who took longer, debt increased modestly, but monotonically throughout the distribution.

Finally, we compare dependent bachelor's recipients by family income in Figure 5. In general, borrowing levels are not that dissimilar across the income groups; although students from wealthier families have more resources, they also tend to graduate from more expensive schools. The standard pattern of fast debt increases during the 1990s is present here, but the most striking trend is that this increase is mostly concentrated between 1990 and 1996 for the lower income group and (more than) entirely concentrated between 1996 and 2000 for the upper income group. A related point is that for the latter time period the borrowing rate increased only mildly for the poorer students; however, it increased by 25 percentage points, from 30 percent to 55 percent, for wealthier students. Taken together, these trends imply that changing factors between 1996 and 2000 had a disproportionate impact on student borrowers whose family income was above the median and affected borrowing at all levels, not just the top.

A few themes from these comparisons stand out. First, the large increase in debt that occurred throughout the distribution between 1990 and 2000 (Figure 1) is common to all of the subgroups

examined, suggesting that behavioral or policy differences, and not composition effects, are more likely to be the prime suspect. Second, "traditional" graduates (dependents who finish on time) experienced debt increases over the 2000s only in the top portion of the distribution, while "non-traditional" students' debt increases were more uniform.

The next section describes the decomposition methods we employ to unpack systematically changes in the empirical distributions.

Decomposition Techniques

To examine the factors behind increases in borrowing, we employ three different econometric decomposition techniques: (a) Oaxaca-Blinder, (b) semi-parametric reweighting, (c) and recentered influence functions. While the first of these techniques is common in the literature, it is not suitable for decomposing statistics other than the mean. The second and third techniques, while not as well known, allow for the decomposition of the entire distribution of borrowing. We briefly describe each of these techniques, including strengths and weaknesses, below.

Oaxaca-Blinder

This technique, independently published by Oaxaca (1973) and Blinder (1973), linearly decomposes the average difference in outcomes across groups into differences in observable characteristics and differences in structural factors. Formally, let $Y_i = X_i\beta_i + \varepsilon_i$ for i=*A*,*B*. Then:

$$E[Y_B - Y_A] = \{E[X_B] - E[X_A]\}\beta^* + \{E[X_B](\beta_B - \beta^*) + E[X_A](\beta^* - \beta_A)\}. \quad (1)$$

The left-hand term represents the average difference between groups *A* and *B*. The first term in braces on the right shows the difference in average observables between the groups, multiplied by a common or reference coefficient vector, β^*. Since X_i is observable for both groups, this component is considered to be what is "explained" by observables. The second term in braces is the difference between the group-specific coefficient vector and the reference coefficient vector, scaled by the observables, for both groups. Since deviations from the reference coefficient are generally not known, this component is considered to

be structural or "unexplained."[11] It is common in many economic applications (notably, wage discrimination) to set the reference coefficient β^* to the estimates of either β_B or β_A. In this case, one of the terms in the second pair of braces drops out, and the "unexplained" portion is the (scaled) deviation of one group's estimated coefficient vector from the other's.[12] In other cases, the reference coefficient is set to a weighted average of β_A and β_B, where the weights depend on the application (see Jann 2008 for an overview).

In practice, Oaxaca-Blinder (O-B) decomposition is straightforward to implement. OLS regressions are performed on groups *A* and *B*, separately, and equation (1) is calculated using estimates of β_A and β_B from these regressions. The technique thus permits detailed decompositions that allow the contribution from each element of X_i (or β) to be estimated, as well as their sum. It is worth noting, however, that the ε_i terms cancel out as a result of the expectations operator (and the standard OLS assumptions). For this reason, O-B decomposition is valid only for the conditional mean function.

In this paper, we use the O-B decomposition to investigate changes in the extensive margin of having ever borrowed at the time of college graduation as well as mean borrowing levels. Our choice of reference coefficients is the set from the earlier time period in the comparison, although we consider the sensitivity of our results to other sets of base coefficients.

Semi-parametric Reweighting

Proposed by DiNardo, Fortin, and Lemieux (1996) in their analysis of changes in the wage distribution, this technique reweights observations from one group so that the joint distributions of the X_i are similar for both groups. By dealing with the joint distributions of the X_i, this reweighting technique overcomes the linearity restriction of O-B and allows the construction of a counterfactual distribution, not

[11] Improper specification of the model, such as omitted variables, would show up in this component.
[12] For example, in gender wage discrimination, the reference coefficient is often set to that of men because the exercise is meant to examine how much of the gender gap can be attributable to differences in characteristics under the assumption that the return to those characteristics is the same as men's (the "explained" part) and how much is attributable to "discrimination" (the "unexplained" part, but see footnote 11).

just the counterfactual mean. Thus, quantiles and other distributional statistics such as variances or Gini coefficients can be compared.[13]

Semi-parametric reweighting is implemented by performing a logit or probit regression on the pooled sample, with the dependent variable being equal to one if an observation is in group *A* and equal to zero if it is in group *B*.[14] The right-hand side variables include all of the elements of X, and in some cases interaction terms as well. Fitted values, $\hat{p}$, from this regression are used to construct propensity weights, $\frac{\hat{p}}{1-\hat{p}}$, for group *B*; weights for group *A* are set equal to 1. If data are sample weighted, the propensity weights can be multiplied by the sampling weights to create composite weights. Distributional statistics for the two groups can be compared by using these composite weights for each group.

We employ this approach to compare the cumulative distributions of borrowing across time periods while controlling for the joint distribution of observables. However, a shortcoming of the reweighting is that it does not easily allow attribution to a specific (marginal) component of *X*. While it is possible to perform the reweighting multiple times, leaving out one element of *X* each time, to isolate the contribution of that particular *X* element on the borrowing distribution, doing so is somewhat cumbersome and tedious for a non-trivially dimensioned vector of *X*. This drawback motivates the third decomposition technique.

Recentered Influence Functions

This technique, suggested by Firpo, Fortin, and Lemieux (2007), is an extension of O-B for statistics beyond the mean, particularly unconditional quantiles. For any quantile *q*, define the recentered influence function as:

$$RIF_q = Y_q + \frac{q}{f(Y_q)} - \frac{1}{f(Y_q)} * \mathbf{1}(Y \leq Y_q), \qquad (2)$$

[13] It is also more robust to functional form violations, such as when Y is a nonlinear function of X. Because this functional form need not be known, the technique is semi-parametric. On the other hand, it is more sensitive to the common support requirement, that sets of covariate realizations are common to both groups.

[14] This setup will make the distribution of X_B resemble the distribution of X_A; switching the coding of the dependent variable will cause the reverse.

where Y_q is the value of Y at quantile q, $f(Y_q)$ is the density of Y at q and needs to be estimated, and $\mathbf{1}(Y \leq Y_q)$ is an indicator function that equals one if, for a given observation, Y is less than or equal to Y_q and zero otherwise. Note that RIF_q takes on only two values determined by whether Y exceeds Y_q. The *RIF* has the interesting property that $E[RIF_q] = Y_q$.[15] Firpo, Fortin, and Lemieux (2007) show that performing O-B on the *RIF* can recover decompositions at the unconditional quantiles of Y.

While it is easy to estimate Y_q in a sample, it is more challenging to estimate $f(Y_q)$. The density is commonly estimated using kernel density methods, and these are somewhat sensitive to bandwidth choices, particularly for distributions, like those for cumulative borrowing, that are not uni-modal and roughly symmetric. On the other hand, the similarity to O-B allows the marginal contribution of specific elements of *X* to be analyzed much more easily than in the case of semi-parametric reweighting.

The Observables

The usefulness of each of the decomposition methods depends on the set of observed variables. Fortunately, the NPSAS data are especially rich. In addition to core demographics such as age, dependency status, sex, ethnicity, and marital status, the set of controls include parents' education, the student's state of permanent residence, the region of the school, whether the student is in-state, whether attendance is full- or part-time and full- or part-year, the type of institution attended, a set of majors, and whether the student changed schools during the last year.[16] In addition to these variables, all of which are binary or categorical, we further include a quartic in expected family contribution (EFC) interacted with dependency status, a quartic in list tuition (or cost of attendance), a quartic in grants, and full interactions of the cost and grant measures.[17] We have chosen to include costs and grants with interactions separately,

[15] This follows because fraction *q* of the population (and sample analogue) has $Y \leq Y_q$, by definition of *q*.

[16] We also experimented with using time to degree, which was available in every wave except 1990. Somewhat surprisingly, the inclusion of this variable, conditional on the others, had a negligible effect on the decompositions. In order to include the 1990 wave, we thus chose to exclude time to degree in the presented results.

[17] Expected family contribution measures a student's family's ability to pay for college expenses and is based on family structure, income, and certain assets (excluding home equity). The relationship between borrowing and EFC is likely non-monotonic, as low EFC levels increase the likelihood of receiving grant aid, and very high EFC levels are associated with a lower need to borrow. The cost of attendance is broader than tuition and fees and also includes room and board, books, travel, and other expenses.

instead of a simple polynomial in net cost (e.g., tuition less grants) because the former approach is more flexible and allows behavioral considerations (such as a response to nominal instead of net prices) while still nesting the more traditional assumption of net cost.[18]

Despite this detail, the data are not quite ideal. The cross-sectional design of NPSAS limits what is observed about attendance history. While it would be useful to know the cost and financial aid structure and attendance intensity for all years before graduation, we see only the final year and must use these data as a proxy for the entire undergraduate experience. Subject to this caveat, summary statistics of these variables, by wave, are presented in Appendix Tables 1A (for continuous measures) and 1B (for categorical measures). We will often refer back to these measures in our discussion of the decomposition results.

Decomposition Results

Oaxaca-Blinder and Ever Borrowed

We first seek to explain the sharp increase in the propensity of having ever borrowed that occurred between 1996 and 2000. As this is a binary outcome and we are interested in the mean change, we use Oaxaca-Blinder and focus across different time intervals that span the borrowing spike. In each of the time intervals, the reference coefficients are set equal to those from the earlier period. The composition effects thus capture changes in the joint means of the observables assuming that the relationship between the observables and borrowing was the same in the later period as it was in the earlier period. This implies that decompositions with different starting periods are not strictly comparable, but they may still be informative. For ease of presentation, we group the observable variables into more aggregate categories[19]: (1) age and dependency status; (2) sex, marital status, and ethnicity; (3) both parents' education level (including missing); (4) state of permanent residence, region of school, and in-

[18] A polynomial in net cost implies parameter restrictions in the more flexible model. When we test these restrictions in the relationship between borrowing and cost and aid structure, they are sharply rejected at all conventional significance levels.

[19] When the constituent variables diverge in their effects, we make note of it.

state status; (5) institutional sector; attendance intensity, and major; (6) expected family contribution; and (7) tuition and grants and interactions.

The first panel of Table 2A looks at the 1990 to 2000 time period, during which the borrowing rate increased by 9 percentage points. The decomposition shows that about 4 percentage points (45 percent) of the increase was due to observable factors, with most of the effect concentrated in expected family contribution (the mean of which fell in this time period) and tuition (which rose). The remaining share of the increase was due to changes in the coefficients relating the observables to the outcome, but the factor-specific estimates are too imprecise to isolate changes in marginal relationships. When examining the longer interval from 1990 to 2008, again 45 percent of the increase is explained, with a slightly greater role for tuition.[20]

The next two panels use 1996 as the base year. While there are slightly larger increases in the percentage borrowing relative to the 1990 base year, the share of the increases attributable to the covariates was smaller, between 20 and 30 percent. Furthermore, the explanatory share in this horizon did not load so heavily on costs, but was more diffuse. On the other hand, the estimates on the coefficients for expected family contribution are quite large and statistically significant, suggesting that for a given ability to pay, students were becoming more likely to borrow. However, because the 1996 sample size is relatively small (see Appendix Table 4), these coefficient estimates are less reliable, and we treat them cautiously.

Table 2B shows that if we use the broader cost of attendance measure instead of tuition (but leave other variables the same), the picture is similar. A slightly larger share is accounted for by the observables, between 50 and 55 percent for the 1990 base, and 30 to 40 percent for the 1996 base, suggesting that increases in non-tuition expenses also increased borrowing. In summary, roughly half of the long-term increase in the borrowing rate was due to observable factors, with cost increases explaining the lion's share. This leaves a substantial fraction due to structural changes, and more so if 1996 is used as

[20] There is also a slightly positive role from demographics (an increasing minority share) and an offsetting role from location (regional changes in graduates from the Northeast to the South).

the base instead. Note that this pattern is consistent with the large increases in the debt profile between 1996 and 2000. The behavioral explanation for this trend is a topic to which we will return.

Semi-parametric Decompositions of the Distributions

What would the cumulative distribution of borrowing in 2008 look like if the distribution of covariates was the same as it was in 1990? Figure 6 answers this question by plotting the CDFs from 1990 and 2008 (as in Figure 1) against just such a counterfactual distribution. Reweighting the 2008 distribution shows that just over half of the increased propensity to borrow (the change in density at 0) can be explained by changes in the covariates—very similar to the Oaxaca-Blinder findings despite a quite different methodology. Moving up the distribution, changes in observables explain approximately half of borrowing levels up to the median, but they become less and less relevant in higher quantiles. This result is not entirely unexpected: the rise in borrowing limits, both federal and private (see Appendix B), are not included in the set of covariates since they vary only over time and not in the cross-section, but we would anticipate that their effect would be concentrated heavily in the upper tail of the distribution.

The four panels of Figure 7 use the same reweighting approach across shorter time intervals. By looking at different time horizons, it is possible to locate when, and where in the distribution, structural changes were more important in affecting borrowing than covariates. The top left panel reweights the covariate distribution in 1996 to resemble the distribution in 1990. Interestingly, the counterfactual shows a much larger reduction in borrowing rates than actually took place, indicating that structural or policy changes increased the fraction of graduates who borrowed.[21] For levels, the counterfactual for 1996 generally gets more than halfway to the 1990 distribution between the 60^{th} and 90^{th} percentiles, and is basically identical to the 1990 distribution for the top decile. Observables clearly explain the bulk of the debt increase between 1990 and 1996, and we will subsequently analyze which observables were important to this change.

[21] Indeed, as Appendix B documents, unsubsidized Stafford loans, which are not means-tested, became available in 1992, and this would be expected to increase borrowing at the extensive margin.

In contrast, the top right panel of Figure 7 illustrates a negligible role for observables between 1996 and 2000. Virtually all of the debt increase throughout the entire distribution is due to unexplained or structural factors. Consistent with the Oaxaca-Blinder analysis, the importance for behavioral changes is much greater in the late 1990s than earlier that decade.[22] In fact, the bottom left panel, looking at changes between 1990 and 2000, is almost a composite of the previous two panels, with the counterfactual distribution approaching the halfway point between the actual 1990 and 2000 distributions through the 50th percentile, and roughly a quarter of the way from the 2000 to 1990 distribution at higher quantiles. In some ways, this is reassuring, as it suggests that the earlier results are not just due to small sample issues in the 1996 wave. The last panel focuses on the change between 2000 and 2008, when debt profiles increased relatively little. Here, reweighting the covariates accounts for *all* of the change up to the 80th percentile, about half of the change between the 80th and 90th percentiles, and almost none in the top decile.

From these decompositions, it appears that changes in observables were responsible for much of the observed shifts in borrowing between 1990 and 1996 and again from 2000 to 2008. The exceptions are that observables over-explain the lower tail in the early 1990s and under-explain the upper tail in the 2000s. Moreover, observables seem to have no explanatory power during the late 1990s. What policy or behavioral explanations fit with these patterns is a topic we return to in the next section. Before that, however, we turn to recentered influence function decompositions in order to gauge which set of observables mattered most.

Recentered Influence Function Decompositions

While the reweighting-based decompositions give a useful graphical overview, the RIF method allows for a greater level of detail of the importance of specific factors. Table 3A presents the decompositions over the 1990 to 2008 period for four statistics: the mean level of borrowing (which uses the Oaxaca-Blinder decomposition), and the 50th, 75th, and 90th percentiles of the distributions of debt. Of

[22] Note that while the Oaxaca-Blinder results showed a small role for observables at the extensive margin during this time period, the estimate was not statistically different from zero.

the mean difference in borrowing levels of $9,440, only $1,740—or about 18 percent—is explained, a result quite consistent with Figure 6. Although the largest component is due to cost structure, the coefficient is imprecise.[23] Smaller, but more precise, effects are due to changes in expected family contribution and attendance patterns.[24]

At the median and 75th percentile, about one-third of the borrowing increase is explained, which is somewhat less than in Figure 6 for the median and somewhat more for the higher quantile, but not statistically different in either case. At both quantiles, the factors that were important at the mean are still relevant, as well as a slight role for age due to graduates being more likely to finish in their mid-to-late 20s. At the 90th percentile, observables explain essentially nothing, as before, and this appears to be due to the coefficient on cost structure turning negative and canceling out the mostly positive effects from other factors.

Table 3B repeats the analysis using the broader cost of attendance measure. The results are similar except for a slightly greater share explained by observables, and this is due entirely to a higher loading on cost structure (which is now statistically significant). Non-tuition costs like room and board matter, particularly in the top decile, where the coefficient is now positive and nearly 30 percent of the increase in borrowing is now explained.

Looking at the bigger picture, between 30 and 40 percent of the debt increase between 1990 and 2008 at both the median and 75th percentiles is explained by observables, and half of this share (15 to 20 percent of the total increase) is due to changes in the cost structure alone (mostly tuition). At the top decile of borrowing, changes in the tuition and grant structure led to students borrowing *less*, but this reverses once non-tuition expenses are accounted for.[25] A small share, between 5 and 7 percent, of the overall rise in borrowing is due to the increased financial resources among dependent students *and*

[23] Standard errors are calculated by bootstrapping the entire RIF procedure (with 100 replications) to account for estimation error in the density function.
[24] The attendance coefficient is largely driven by a shift from non-doctoral to doctoral public institutions (see Appendix Table 1B).
[25] For evidence of the growing role of amenities in driving students' attendance decisions, see Jacob, McCall, and Stange (2013).

decreased financial resources of independent students. This apparent paradox results from fewer dependent students qualifying for need-based grant aid and more independent students not having sufficient outside income to pay for college expenses. In all cases, a large share remains unexplained, with the largest change in coefficients loading on cost structure. That is, not only did costs rise, but for a given set of costs and grants, students borrowed more than they did in the past, conditional on the other observables.[26]

We now break down the decompositions by time period, looking at the 1990s and 2000–2008 periods separately. For the former, a quick glance reveals that, as in the reweighting-based decomposition, observables explain quite little (less than 15 percent) of the debt increase over the 1990s, and this is true throughout the distribution and whether the tuition (Table 4A) or cost of attendance measure (Table 4B) is used. The reduction in explanatory power relative to that of the longer period stems from the smaller load on cost changes, which is now negligible. This pattern is consistent with relatively small increases in both list tuition and cost of attendance in this time frame (Appendix Table 1A). While *structural* (coefficient) effects on costs are generally still present, the coefficients representing the constant term, which likely capture omitted variables such as policy reforms, are quite large and statistically significant, especially in upper quantiles.

Before turning to the 2000s, it is helpful to look at the 1990s more granularly, as the reweighting decompositions showed substantial differences in the role of observables between the two. Appendix Tables 2 and 3 examine the 1990–1996 and 1996–2000 periods separately. Reassuringly, the RIF estimates are quite consistent with the reweighting decompositions, down to over-explaining the increase in borrowing at the median in the early 1990s.[27] For the earlier period, the change in age composition, cost structure, and EFC all contribute significantly toward greater debt at the median and above, although

[26] Recall that the cost structure observables includes quartics in list prices and grants *and* their pair-wise interactions. These relationships are highly nonlinear and non-monotonic, so it is not quite accurate to say that "costs" rose; we are trading off accuracy for expediency.

[27] That borrowing increased less than predicted at this quantile is due to a change in the coefficients on marital status; specifically, married graduates borrowed less than single graduates in 1990, but this relationship reversed in 1996. Since singles increased as a share of graduates between 1990 and 1996 (Appendix Table 1B), borrowing was predicted to have increased.

the explanatory share falls from about one-half to about one-third when moving from the 75th to the 90th percentile, as the roles of cost structure and age diminish at the top of the distribution.[28] In contrast, between 1996 and 2000, the role of age has diminished to negligible levels, and while attendance patterns contribute slightly toward greater borrowing, this covariate is outweighed by the reversal in the relationships of EFC and costs. In short, observables explain more of the debt increase between 1990 and 1996 because the observables that tend to matter the most, cost and EFC, changed more in this period than they did between 1996 and 2000. *The apparent puzzle is that the fastest debt growth occurred in a period when list cost barely budged and net cost actually fell.*

The contribution of changes in observables in the latter period, on the other hand, is considerable, and much of this is due to increases in cost. Over the 2000s, the RIF approach substantially over-explains the shift in borrowing at the median and 75th percentiles, and more so when cost of attendance is used instead of tuition. Changes in cost often account for half or more of the observable share, although parental education, attendance patterns, and EFC also matter. At the 90th percentile, where the reweighting decomposition found only a modest role for observables, the RIF technique places more weight on them, explaining at least two-thirds of the debt increase, with cost and EFC having the largest impact.[29]

In summary, the RIF decompositions qualitatively resemble the reweighting-based decompositions: a large role for observables in the early to mid 1990s and again between 2000 and 2008, but almost no explanatory power for them in the late 1990s. In both techniques, observables explain less at the very top of the borrowing distribution. The RIF analysis, however, shows that the most important observed factor contributing toward greater borrowing is cost, and this is particularly true over the 2000s, when costs increased relatively quickly. EFC also matters consistently across time periods and quantiles.

[28] The explanatory shares of the cost structure are slightly lower when using cost of attendance rather than tuition, reflecting a reduction in non-tuition expenses that partially offset tuition hikes over this period (see Appendix Table 1A).

[29] The difference between the decomposition techniques over this time period may stem from the functional form limitation of the RIF (i.e., linearity). However, isolating the role of specific observables through sequential reweighting yields similar relative magnitudes as RIF.

Smaller, but still meaningful effects are found for age composition over the 1990s and attendance patterns and parental education over the 2000s, with all of these more prominent in the middle of the debt distribution. Core demographics and geography, despite changing a great deal over time, do not seem to be related to the shift in student debt.

Discussion

Although much of the increase in debt over the 2000s can be explained by changes in student and institutional characteristics, the NPSAS data point to structural, behavioral, or policy shifts underlying the majority of debt increases in the 1990s, particularly the late 1990s, when debt profiles increased faster than in any other period. In this section, we review several possible explanations and weigh the evidence for each.

Formal Loans Crowding Out Informal Loans

The analysis above considers formal loans in which a promissory note has been signed and repayment shows up in credit reports. Informal loans from friends and family also occur, although it is unclear whether these are actually intended to be repaid (with or without interest). A possible "structural" explanation for debt increase is that formal loans have displaced informal loans over time. The NPSAS asked about informal loans only through 2000, but that covers the period where observables have little explanatory power. Figure 8 plots CDFs of borrowing for 1990, 1996, and 2000, both with and without informal borrowing. The dashed lines reflect the distributions of formal borrowing from Figure 1, while the solid lines add in informal borrowing. If displacement were occurring, we would see the difference between total and formal borrowing shrink over time. In fact, we see the opposite: while informal borrowing is rare in 1990, it expands by 1996, and is of a similar magnitude in 2000. We can thus rule out this story.

Parents Transferring Loan Burden to Children

Our analysis also has focused on debt in the student's name and thus has excluded borrowing directly by parents in the form of PLUS loans. While the terms of student-level loans are more generous than parent-level PLUS loans (see Appendix B), parents are often in a better financial position with which to make repayment. However, if parents have become less willing or able to borrow for their children than in the past, the transference of the burden could explain increases in student-level borrowing. Figure 9 shows that this is not the case. The dashed-line CDFs in the figure again show the distributions of borrowing taken from Figure 1, while the solid lines add in cumulative PLUS borrowing of parents (data for 1990 are unfortunately unavailable). Rather than decreasing over time, PLUS borrowing has increased substantially and become more pronounced farther down the distribution. Thus, the intergenerational transfer explanation, at least through PLUS loans, does not work, either.[30]

Change in Interest Rates

The prevailing interest rates for student loans (federal and other) have varied over time, and basic economic theory implies that lower interest rates should increase borrowing, all other things equal. Figure 10 presents time series for the interest rates on the predominant federal loan program, the Stafford loan, as well as the U.S. prime rate, a benchmark for private loans and, at times, federal loans as well. Prior to the early 1990s, Congress periodically set a fixed rate for Stafford loans (sometimes creating a huge subsidy, as in the early 1980s). Between 1992 and 2006, Stafford rates were tied to market rates before being fixed again.[31] Although interest rates drop by nearly 4 percentage points in the early 1990s, they are relatively flat over the late 1990s, and they drop considerably again over the early 2000s recession. This pattern is not consistent with large increases in borrowing during the late 1990s and a flat, or even falling, borrowing profile between 2000 and 2004. It would appear that the borrowing decisions of bachelors

[30] We unfortunately cannot observe direct transfers from parents to students. However, since parents are borrowing more through PLUS loans, it is very unlikely that the transfer motive has decreased.

[31] In late summer of 2013, new legislation passed that will again tie federal borrowing rates to market rates. Although there was much attention before that legislation that subsidized rates would double to the unsubsidized rate for the 2013–2014 school year, it was generally unmentioned that the rates had differed only since 2009. Note also that, historically, Stafford rates were similar to or below the U.S. prime rate *except* during and after the Great Recession.

graduates are insensitive to the cost of loans, and interest rates therefore cannot explain the observed borrowing patterns.[32]

Borrowing Eligibility

Another set of possible factors deals with increased eligibility to borrow (on both extensive and intensive margins). As Appendix B shows, major changes in the student loan market took place over the 1990s, including the introduction of unsubsidized Stafford loans (which are available to higher-income families than subsidized Stafford loans), increases in the federal statutory borrowing limits, and the development of private loans. In contrast, there was little structural change in the loan market between 2000 and 2008 (although there was considerable change *after* the 2007–2008 school year).

Except for a brief period in the late 1970s and early 1980s, federal loan eligibility was means-tested and subject to a family income threshold (with the government paying interest while the student was enrolled) until 1992. That year, unsubsidized loans first became available. While interest accumulates on these loans from the time of disbursement, students have access to them regardless of family income. Their availability opened up a large segment of the student population to federal loans, so it would not be surprising if increases in the debt distribution followed. However, Figure 5 shows that borrowing increased much faster in the late 1990s than between 1990 and 1996 for wealthier students—the group that would be expected to benefit most from unsubsidized loan eligibility. Why didn't their borrowing increase immediately after 1992? Three reasons suggest that unsubsidized loans became more important in the late 1990s despite becoming available for the 1992–1993 school year. First, total unsubsidized loan volume was quite small initially. In 1992–1993, unsubsidized loan disbursements for *all postsecondary students* totaled $440 million (in 2011 dollars), just 1.9 percent of aggregate loan volume. In 1993–1994, disbursements had increased to $2.7 billion , a 8.9 percent share, and in 1994–1995, the numbers jumped to $9.5 billion, a 26.7 percent share. From this point, the share increased slowly, to 32.5 percent by 1999–2000 (Baum and Payea, 2012, Figure 6). Their impact thus would have been muted for the 1996

[32] There is surprisingly little work on the elasticity of credit demand with respect to price in the higher education context, and it would be a fruitful area for future research.

graduating cohort relative to the 2000 cohort. Second, the Survey of Consumer Finances shows a pronounced jump between 1995 and 1998 (but not between 1992 and 1995) in both the median and mean values of educational loans among families, and these increases were concentrated among families whose heads were college educated, were in managerial or professional occupations, and had income in the second-highest quintile (Survey of Consumer Finances, 2012). This pattern is consistent with increased loan volume for the types of households that would benefit most from unsubsidized loans occurring several years after the program's introduction. Third, and perhaps most tellingly, the income distribution among graduates with a Stafford loan their senior year increased sharply (relative to the income distribution of all graduates) between 1996 and 2000, but not between 1993 and 1996.[33] This is shown in Figure 11, with the dashed lines representing the income distributions for Stafford borrowers and the solid lines those for all graduates.

Furthermore, it is possible to use a back-of-the-envelope calculation to quantify how important unsubsidized loans were to the increase in total borrowing between 1996 and 2000. Figure 12 shows the *difference* in senior-year borrowing between 2000 and 1996 by percentile (now on the x-axis). The red, solid line counts all borrowing, and the blue, dashed line nets out unsubsidized Stafford loans. Throughout much of the distribution, the gap between the two years is significantly reduced once unsubsidized loans are taken out of the picture.[34] Indeed, the mean gap in the latter case is only 36 percent of the former; the mean squared deviation, 66 percent. (Excluding the area above the 98th percentile, the two figures are 24 and 23 percent). In levels, the mean gap is reduced by $460, and if this pattern held for previous years of enrollment, it could account for about $2,500, or about half of the mean increase in total debt of $5,200 between 1996 and 2000.[35] However, because of the factors mentioned above, this gap in senior-year borrowing probably overstates the gap for earlier years, when unsubsidized loan volume was

[33] We use the 1993 wave here because the data are for senior-year and not cumulative borrowing. While NPSAS separates survey school year Stafford borrowing by subsidized and unsubsidized status beginning in 1996, it unfortunately does not do so for cumulative borrowing.
[34] The notable exception is above the 95th percentile and is probably due to private loans, which are discussed next. The negative section of the lines reflects the nominal annual borrowing cap for Stafford loans.
[35] This assumes a time to degree of 5.5 years, the average for the class of 2000 in the NPSAS data.

smaller. But the senior year alone can account for 9 percent of the $5,200 difference, and even conservative estimates for previous class years would bring this share to a quarter.

Thus, unsubsidized loans were important for debt increases in the late 1990s, but what about increases in federal borrowing limits? The annual borrowing limit under the Stafford program, by far the largest federal loan program, was fixed—in nominal terms—between July of 1993 and June of 2007. The supplementary limits, which raise the maximum borrowing for independent students and some dependent students, were fixed between July of 1994 and June of 2008. However, the increase in statutory limits in July of 1993 applied only for students in their second or higher class year, and the shift in supplementary limits in July of 1994 applied only for students in their third or higher year. This means that graduates of the class of 1996 experienced the same nominal limits as the class of 2000; in real terms, borrowing limits declined slightly.[36] Consequently, borrowing limits are not behind the debt increase of the late 1990s.

The other major innovation of that time period was the rise of the private loan sector. Too new to be asked about explicitly in the 1996 NPSAS, private loans in that wave must be inferred by netting out institutional and state-level loans from all nonfederal loans. For graduates of that class, 1.0 percent took out a nonfederal loan their senior year, and most of these were institutional or state loans: just 0.3 percent took out an "other" or private loan. By 2000, the numbers had increased to 6.2 and 4.6 percent, and they continued to grow through 2008 before retrenching during the recession. The NPSAS data do not break out cumulative borrowing by nonfederal sector, but one can compare cumulative total borrowing to cumulative federal borrowing, and this is done in Figure 13 for the 1996 and later waves. The dashed lines represent cumulative total borrowing, as in Figure 1, while the solid lines show cumulative federal borrowing only. The difference between the two captures nonfederal borrowing. For 1996, the gap occurs entirely below the 70th percentile and is small, reflecting that most of these nonfederal loans are institutional or state-based and intended for lower income borrowers.[37] By 2000, the gap has widened, especially above the 75th percentile. In fact, about half of the total increase in borrowing in the top quartile

[36] These statements also apply to lifetime borrowing limits.
[37] In 1996, the median recipient of institutional or state loans was below the 30th percentile of the family income distribution of all students; the median recipient of a private loan was at the 70th percentile.

is due to nonfederal—essentially private—loans. (For the same quartile between 1996 and 2008, about three-quarters of the increase is due to private loans, and between 2000 and 2008, all of it is, though much of this latter increase is accounted for by observables.[38])

Broadly speaking, informal loans, PLUS loans, interest rate changes, and statutory borrowing limits are unlikely to be factors behind the debt increase of the late 1990s. The advent of unsubsidized Stafford and private loans, on the other hand, are probable culprits, at least in part.

Other Explanations

One factor for which we have been unable to control is the use of tax credits, which came into existence in the late 1990s. Nicholas Turner (2012) has shown that the value of educational tax credits is largely capitalized into increases in net tuition (i.e., fewer institutional grants), with suggestive evidence that students compensate by borrowing more. However, the timing of tax credit availability precludes them from being a major factor for the borrowing increases observed during the 1990s. Both the Hope and Lifetime Learning credits first became available during the 1998–1999 school year (and tax deductions not until the 2002–2003 school year). As the Hope credit was available only for the first two years of postsecondary education, it could not have benefitted students who graduated during the 1999–2000 school year and took four or more school years to finish (at least 87 percent of the graduates, and probably close to 98 percent[39]). Until 2003, the Lifetime Learning program allowed a (non-refundable) credit of 20 percent of the first $5,000 in tuition and fees, making its maximum value $1,000.[40] But, according to the NPSAS data, only 20 percent of graduates' families claimed the credit in the 1999 tax year, and only about 20 percent of the claimants qualified for the maximum credit. Even if the credit were

[38] That observables can account for most of the change between 2000 and 2008 is likely in large part due to private loans being well established by 2000. An auxiliary analysis that examined the period between 1996 and 2008, with 1996 reference coefficients, found a much smaller role for observables.

[39] The NPSAS data record the year of first postsecondary attendance and the year of degree, and our measure of time to degree is the difference between these two years. Since some students start in the spring and would graduate in the fall after eight semesters (e.g., begin in spring of 1996 and finish in fall of 1999) but with only three calendar years' time elapsed, our measure is biased downward. The Baccalaureate and Beyond 2000 survey shows that only 2.1 percent of bachelor's graduates in the 1999–2000 school year took 36 or fewer months to finish.

[40] The threshold was raised from $5,000 to $10,000 in 2003. In all years, the credit was subject to income phase-outs.

completely capitalized into higher borrowing, it could explain only a tiny portion of the increase in debt in the late 1990s.[41]

Another factor we did not examine explicitly is the use of home equity loans and lines of credit for education financing. There is some evidence that these tools became less important over the 1990s. In the NPSAS, about 4 percent of the class of 1996 claimed to have used them, versus 2.3 percent in 2000. Similarly, the Survey of Consumer Finances shows that in 1989, 0.6 percent of all households were using a home equity loan or line of credit to pay for education expenses; by 1995, the last year in which the purpose can be ascertained, the figure had fallen to 0.4 percent. These levels are probably too small to have had a meaningful impact on student borrowing.

Conclusion

Using NPSAS data, we have shown that the entire debt profile of college graduates grew much faster in the 1990s than in the 2000s, and that this growth was concentrated especially in the late 1990s. Between 2000 and 2008, debt was remarkably stable for the bottom three quarters of the distribution—the increase that took place in the top quartile was driven by graduates at private, not-for-profit institutions and stemmed largely from greater private borrowing.

Statistical decomposition techniques consistently indicate that observable characteristics of students and institutions—such as demographics, geography, attendance patterns, income, and costs—explain about one-third of the overall debt increase across the two decades. However, their explanatory power is stronger at the extensive margin and in the middle of the distribution than in the right tail. Moreover, the observables account for more than half of the increase between 1990 and 1996 and approximately all of it between 2000 and 2008, leaving the late 1990s as the period that remains unexplained. Among the observables, cost tends to be the most important factor, often explaining about

[41] A back-of-the-envelope calculation that subtracts the real value of the credit from cumulative borrowing—among *all* students, not just those claiming the credit—shows that debt would have risen by about $500 less at most points of the distribution between 1996 and 2000. At the mean, this is less than 10 percent. At the 90th percentile, it is about 5 percent.

half of the observable share, but expected family contribution is also important, as is the age structure in the 1990s and parental education in the 2000s.

In trying to unpack the puzzle of the late 1990s, we have ruled out informal loans, PLUS loans, interest rate changes, and statutory borrowing limits as likely explanations. Instead, the evidence is consistent with the "unobservable" share being driven by the advent of unsubsidized Stafford loans and private loans. This would imply that much of the debt increase over the 1990s—a much greater increase than over the 2000s—was primarily due to supply-side and not demand-side factors. Indeed, while nominal costs of college have risen considerably, so has financial aid in the form of grants (Avery and Turner 2012; Greenstone and Looney 2013; see also Appendix Table 1A.) Our findings show that changes in costs account for only 20 percent of the increase in distribution-wide borrowing between 1990 and 2008, and this is after capturing the complex interplay between nominal prices and grants.

Of course, the partial-equilibrium analysis we have undertaken does not account for the possible endogeneity between college costs and financial aid. Recent research on the incidence of federal financial aid indicates that in many cases, schools seem to capture much of the benefit: Nicholas Turner (2012) cannot rule out that schools offset one dollar of student tax benefits with one dollar higher net tuition, Lesley Turner (2012) finds that schools on average reduce institutional grants by 16 cents per dollar of Pell grants (80 cents for selective non-profits), and Cellini and Goldin (2012) conclude that Title IV-eligible for-profits charge tuition 78 percent higher for comparable programs in non Title IV-eligible for profits. Whether these results hold specifically for loans, however, is still an open question where more research is needed.

Finally, we acknowledge that our findings do not capture the effects of the Great Recession. We plan to update the paper to include analysis from the 2012 NPSAS in 2014.

References

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

Baum, Sandy, and Jennifer Ma. 2012a. *Trends in College Pricing (2012)*. The College Board.

Baum, Sandy, and Kathleen Payea. 2012. *Trends in Student Aid (2012)*. The College Board.

Blinder, Alan S. 1973. "Wage Discrimination: Reduced Form and Structural Variables." *Journal of Human Resources* 8(4): 436–455.

Bound, John, Michael F. Lovenheim, and Sarah Turner. 2010. "Increasing Time to Baccalaureate Degree in the United States." NBER Working Paper 15892.

Cellini, Stephanie R. and Claudia Goldin. 2012. "Does Federal Student Aid Raise Tuition? New Evidence on For-Profit Colleges." NBER Working Paper 17827.

Consumer Financial Protection Bureau. *Private Student Loans*. 2012. Report to the Senate Committee on Banking, Housing, and Urban Affairs, the Senate Committee on Health, Education, Labor, and Pensions, the House of Representatives Committee on Financial Services, and the House of Representatives Committee on Education and the Workforce. Washington: Consumer Financial Protection Bureau.

DiNardo, John, Nicole M. Fortin and Thomas Lemieux. 1996. "Labor Market Institutions and the Distribution of Wages, 1973–1992: A Semiparametric Approach." *Econometrica* 64(5): 1001–1044.

Dynarski, Susan, and Judith Scott-Clayton. 2013. "Financial Aid Policy: Lessons from Research." *Postsecondary Education in the United States* 23(1): 67–91.

Firpo, Sergio, Nicole M. Fortin and Thomas Lemieux. 2009. "Unconditional Quantile Regressions." *Econometrica* 77(3): 953–973.

Greenstone, Michael, and Adam Looney. 2013. "Rising Student Debt Burdens: Factors Behind the Phenomenon." The Hamilton Project. July 5. Retrieved on July 19 from: www.hamiltonproject.org/files/downloads_and_links/June_Jobs_Blog_Final_1.pdf

Jacob, Brian, Brian McCall, and Kevin Stange. 2013. "College as Country Club: Do Colleges Cater to Students' Preferences for Consumption?" NBER Working Paper 18745.

Jann, Ben. 2008. "The Blinder-Oaxaca Decomposition for Linear Regression Models." *Stata Journal* 8(4): 453–479.

Joint Economic Committee. 2013. "The Causes and Consequences of Increasing Student Debt." U.S. Congress. June.

Kantrowitz, Mark. 2012. "Who Graduates College with Six-Figure Student Loan Debt?" Accessed on June 1, 2013 at: www.finaid.org/educators/20120801sixfiguredebt.pdf.

Lewin, Tamar. 2013. "Financing for Colleges Declines as Costs Rise." *New York Times*. March 6.

Martin, Andrew, and Andrew W. Lehren. 2012. "A Generation Hobbled by the Soaring Cost of College." *New York Times*. May 12.

Mitchell, Josh, and Maya Jackson-Randall. 2012. "Student-Loan Debt Tops $1 Trillion." *Wall Street Journal*. March 22.

Oaxaca, Ronald. 1973. "Male-Female Wage Differentials in Urban Labor Markets." *International Economic Review* 14: 693–709.

Quinterno, John, and Viany Orozco. 2012. *The Great Cost Shift: How Higher Education Cuts Undermine the Future Middle Class*. Demos.

Senate Committee on the Budget. 2006. "2002 Student Loan Law Takes Effect, Lowers Interest Rates." U.S. Senate Budget Bulletin. August 4.

Survey of Consumer Finances. 2012. *2010 SCF Chartbook*. Federal Reserve Board of Governors. Accessed on 2013 September 23 at: www.federalreserve.gov/econresdata/scf/files/2010_SCF_Chartbook.pdf.

Turner, Lesley. 2012. "The Incidence of Student Financial Aid: Evidence from the Pell Grant Program." University of Maryland Working Paper.

Turner, Nicholas. 2012. “Who Benefits from Student Aid? The Economic Incidence of Tax-Based Federal Student Aid.” *Economics of Education Review* 31: 463–481.

Appendix A: NPSAS Data Details

The National Postsecondary Student Aid Survey (NPSAS) has been fielded eight times: in 1987, 1990, 1993, 1996, 2000, 2004, 2008, and 2012. In each case, the year references the spring semester or the end of the school year. This paper employs the restricted-use 1990, 1996, 2000, 2004, and 2008 waves.[42] The 1987 wave was not used because its sampling frame and question bank were significantly different from subsequent waves. The 1993 wave was not used because of an interviewing error that caused the cumulative borrowing question—the key outcome of interest—to be not asked of some graduates. Finally, the 2012 wave was released too late to be included in the present study.

The longitudinal comparisons in the paper required that the data be harmonized across waves. Although each wave is similar to the preceding one, there have been many variable name changes and some definition (or universe) changes over time. In most cases, it was straightforward to rename variables or recode values for consistency, although this often necessitated losing some detail for categorical responses. The processing of the most important variables for the analysis is described in this section.[43]

The primary outcome variable is the cumulative borrowing of the student from all sources for undergraduate education. The variable is called BORAMT1 in each of the waves we use, although its construction varies somewhat over time. In 2004 and 2008, BORAMT1 is constructed as the greater of the student's self-reported borrowing total, the cumulative federal borrowing total taken from the National Student Loan Data System plus self-reported private borrowing in the survey school year, or self-reported total borrowing in the survey school year. In each case, the borrowing numbers exclude loans in the name of parents or guardians (e.g., PLUS loans), as well as informal loans without a promissory note. This is the definition of cumulative borrowing we use in the paper.

In 2000, there are two versions of the BORAMT1 variable: one that matches the 2004 and 2008 definition, and one that also includes informal loans. We use the first version. In 1996, only the version that includes informal borrowing is available. However, a separate variable (FAMLOAN) asks about cumulative loans from friends and family, although this variable is present only for the portion of the sample that conducted a CATI interview (about two-thirds of the overall sample). We revise the BORAMT1 variable in 1996 by subtracting from it FAMLOAN among the CATI part of the sample and use a revised sampling weight (see below) to correct for the smaller sample. In 1990, historical data from NSLDS was not included. All cumulative borrowing is thus self-reported. After applying the maximizer decision rule used in the 2004 and 2008 waves, PARLOAN is subtracted from BORAMT1, where PARLOAN represents informal borrowing for the 1989–1990 school year. (There is no cumulative informal loan variable in the 1990 NPSAS, so cumulative loan totals, especially near the top of the distribution may be biased upward.)

The cleaned BORAMT1 variable is converted to year 2012 dollars using the PCE deflator from the Bureau of Economic Analysis. We then applied the following sample restrictions to focus on our population of interest (the 2008 nomenclature of the variable(s) used for the restriction are in parentheses):

(1) The respondent was enrolled in the fall term of the survey school year at an institution in the 50 U.S. states or the District of Columbia (COMPTO87=1)
(2) The respondent was enrolled at a primarily bachelor's (or higher) degree-granting institution (SECTOR9=3,4,6,7,9)

[42] These are available by license from the Institute of Education Sciences after an application process.
[43] A complete list of variables used is available by request from the authors.

(3) The respondent was an undergraduate during the survey school year (STUDTYP=1)
(4) The respondent earned a bachelor's degree during the survey school year (COLLGRAD=1)
(5) The respondent was not a foreign or international student on an education visa (SAMESTAT≤2)

Additionally, because construction of a consistent, cumulative borrowing amount in the 1996 wave required variables from the interview component of the survey, the estimation sample for that wave was restricted to respondents with a positive interview weight (CATIWT>0). The WTA000 sample weights were used for the 2000 and later waves; the WTB000 sample weight was used for the 1996 wave, and the PSKEEPWT sample weight was used for the 1990 wave.

The NPSAS is also used as the sampling frame for two longitudinal studies, the Beginning Postsecondary Study (BPS), which follows first-time postsecondary students, and the Baccalaureate and Beyond (B&B), which follows bachelors recipients. These two studies alternate NPSAS waves, with BPS being drawn from the 1990, 1996, 2004, and 2012 NPSAS waves, and B&B being drawn from the 1993, 2000, and 2008 NPSAS waves. Since the population of interest for the longitudinal studies is over-sampled in the NPSAS, the effective sample sizes for college graduates is quite large in 2000 and 2008 and somewhat smaller in 1990, 1996, and 2004. Final sample sizes, rounded to the nearest 10 to comply with IES rules on disclosure, are found in Appendix Table 4.

The NPSAS is not the only data set that can be used to track how student debt profiles have changed over the last twenty years. Some longitudinal data sets, such as the National Longitudinal Survey of Youth and the Panel Study of Income Dynamics, ask about educational borrowing, as does the repeated, cross-sectional Survey of Consumer Finances. These latter data sets have the advantages of containing information on other forms of debt as well as income and loan repayment information. However, they have a few disadvantages relative to the NPSAS. First, the sample sizes are much smaller. While the NPSAS typically has several thousand college graduates in each wave (or cohort), the other data sets often have only a few hundred per cohort, and this makes examining entire distributions more difficult. Second, the data on college expenses and financing are not as detailed. The NPSAS benefits from the merge with administrative data on exact loan amounts and family financials, and it surveys students immediately after the school year. The other data sets often gather loan data from retrospective questions, introducing the possibility of recall bias.

While we believe NPSAS is the best data set to look at cumulative borrowing at graduation, the other data sets have comparative advantage when investigating the topics of loan repayment, default, and debt-income ratios, and how and why these have changed over time. These are clearly important topics, but they are beyond the scope of the current paper.

Appendix B: A Primer on Student Loans[44]

Federal Loans

The primary federal borrowing program for undergraduates is the Stafford program, established in 1965 as part of the Higher Education Act that year, and later named after Senator Robert Stafford in 1988. Originally intended for lower income students, all Stafford loans were subsidized until 1992, with the federal government paying the interest while the student was enrolled.[45] Beginning that year,

[44] Much of the material in this appendix, including the timeline of changes and statutory borrowing limits is drawn from www.finaid.org and Dynarski and Scott-Clayton (2013).
[45] Except for a brief period between 1978 and 1982, eligibility for subsidized loans has been means-tested throughout the life of the program.

unsubsidized Stafford loans became available for students regardless of their financial background. These loans accumulate interest while the student is enrolled, although repayment for either type of Stafford loan does not begin until six months after school leaving.

Stafford loans account for the vast majority of federal lending to students. Among graduating seniors in the 2007–2008 school year, for example, Stafford loans accounted for 96 percent of federal borrowing, with about 60 percent of this volume as subsidized loans.[46] While these loans require the student to fill out the Free Application for Federal Student Aid (FAFSA) form, they are not subject to credit checks and do not require a cosigner. There are limits, however, to how much can be borrowed each year and over a student's lifetime, and these limits have changed over time.

The other major undergraduate federal student lending program is the PLUS loan, which consists of loans to the students' parents (or legal guardians) rather than directly to the student herself. These also require the FAFSA form to be filed. Unlike the Stafford loan, PLUS loans are subject to a credit check.
Both the Stafford and PLUS loans were administered under the Federal Family Educational Loan (FFEL) program as well as the direct program. Under FFEL, private lenders made loans to students under the terms set by the federal government and received subsidies to cover interest rate spread and nonpayment. Under the direct program, the federal government acted as the lender. The FFEL program was ended in 2010, with all new loans operating under the direct program. From the point of view of the borrower, there is practically no difference between the two programs, as terms and conditions are identical.

The table below shows the annual limits (in nominal dollars) for the Stafford program.

Undergraduate Stafford Borrowing Limits, by calendar year and class standing

	1st year	2nd year	3rd year	4th year +
1977 to 1986	2,500	2,500	2,500	2,500
1987 to June 1993	2,625	2,625	4,000	4,000
July 1993 to June 2007	2,625	3,500	5,500	5,500
July 2007 to June 2012	3,500	4,500	5,500	5,500

Starting in July of 2008, these limits were increased by $2000 for each class year for dependent students (whose parents were not denied a PLUS loan), but this higher limit was available only as an unsubsidized loan.

For independent students and dependent students whose parents were denied a PLUS loan, the limits in the table above were increased by the amounts shown below, with these higher limits also available only as unsubsidized loans:

Supplemental Undergraduate Stafford Borrowing Limits, by calendar year and class standing

	1st year	2nd year	3rd year	4th year +
1986 to1994	4,000	4,000	4,000	4,000
July 1994 to June 2008	4,000	4,000	5,000	5,000
July 2008 to June 2012	6,000	6,000	7,000	7,000

[46] Almost all of the remaining 4 percent consists of Perkins loans, which are intended for very low-income students and function similarly to subsidized Stafford loans with slightly more generous terms. About 6 percent of graduating seniors in 2008 received Perkins loans, while 51 percent received Stafford loans.

In addition to these annual loan limits, there are also cumulative lifetime limits on Stafford borrowing. Through 1986, this aggregate limit was $12,500. In 1987, the limit was raised to $17,500, and in October of 1992, it was raised again to $23,000. For subsidized loans, this (nominal) cap is still in place. For dependent students (whose parents were not denied a PLUS loan), the lifetime Stafford limit was increased to $31,000 in July 2008, but amounts beyond $23,000 must be unsubsidized loans. For independent students (or dependent students whose parents were denied a PLUS loan), the limit was increased to $46,000 in July 1994 and to $57,500 in July 2008, but again the balance beyond $23,000 must be unsubsidized loans.

The PLUS loan for parents was initially capped at $4,000 per year (and $20,000 per student lifetime), but this was changed in 1993, with the annual limit set to the net cost of attendance (list price minus grants) and an unlimited lifetime amount.

Private Loans

Private educational lending (not to be confused with FFEL loans, above) was practically nonexistent until the late 1990s, when attendance costs grew enough relative to the federal borrowing limits to create a market for additional lending. Private loans do not require a FAFSA but often do require credit checks and/or a cosigner. Their interest rates and fee structures are generally more variable than federal loans, with terms that are often worse except for the most creditworthy students. While there are no statutory borrowing limits, borrowing is functionally limited by the net cost of attendance as well as creditworthiness.

Among graduating seniors in 2008, 20 percent took out a private loan that year, and these loans comprised 34 percent of total borrowing. (These numbers were up from 5 percent and 10 percent, respectively, in 2000.) Private loan volume decreased dramatically after 2008 in the wake of the Great Recession as private capital dried up, but it is expected to grow again as the economy improves as long as education costs continue to grow faster than federal borrowing limits.

Other Loans

In addition to the federal and private lending programs, some states and educational institutions themselves have lending programs. These programs, however, are very small relative to federal and private loans. Less than 2 percent of graduating seniors in 2008 received loans from states or educational institutions, and these loans comprised less than 2 percent of the total loan volume for the same set of students.

Finally, there are informal loans in which students borrow directly from their parents or other relatives and friends. As these loans do not appear on credit reports and it is uncertain whether they are expected to be repaid, NPSAS stopped collecting data on them after 2000.