MIT Open Access Articles

The student loan consolidation option

The MIT Faculty has made this article openly available. ***Please share*** how this access benefits you. Your story matters.

Citation: Lucas, Deborah and Damien Moore. "The student loan consolidation option." Journal of Public Economics 174 (June 2019): 1-12. © 2019 Elsevier B.V.

As Published: http://dx.doi.org/10.1016/j.jpubeco.2019.01.011

Publisher: Elsevier BV

Persistent URL: https://hdl.handle.net/1721.1/130139

Version: Original manuscript: author's manuscript prior to formal peer review

Terms of use: Creative Commons Attribution-NonCommercial-NoDerivs License

Massachusetts Institute of Technology

The Student Loan Consolidation Option

Deborah Lucas
MIT

Damien Moore
Moody's Analytics

This version: December 2018

The authors wish to thank John Kolla and Marvin Phaup for help and comments on this project.

Lucas gratefully acknowledges support from the Searle Foundation.

Abstract

The federal government makes credit for higher education widely available at subsidized prices through its student loan programs. A loan provision that at times has significantly increased the size and volatility of the subsidies is the consolidation option, which allows borrowers to convert floating rate loans to fixed rate loans, often on very favorable terms. The option provides a novel setting in which to study how unsophisticated borrowers respond to financial incentives. We develop a model to evaluate the option's cost to the government, and find that the ex-ante cost of the option ranged from 0.8 percent to 6.4 percent of loan principal between 1998 and 2005. Using a sample of 700,000 student loan records, we find that borrowers responded to the time-varying incentives to consolidate although some left sizeable amounts of money on the table; that more indebted borrowers were more likely to optimize; and that there is some evidence of learning over time.[1]

Introduction

The federal government increasingly relies on credit assistance in lieu of other forms of aid to support certain private activities. In 2016, outstanding balances on federal direct and guaranteed loans exceeded $3.6 trillion.[1] The federal student loan programs, which make subsidized federal financing for higher education widely available, comprise about $1.1 trillion of that total.

Government backing of student loans can improve social welfare when imperfections in private credit markets lead to credit rationing or other inefficiencies, or when there are positive externalities from education.[2] However, the effectiveness of the programs remains controversial.[3] In general for a credit subsidy to be cost effective, it should increase access to financing or lower the cost of educational investment for target groups at the lowest achievable cost to taxpayers.[4] Although the federal student loan programs achieve the goal of improving access to credit for higher education, they also include high-cost provisions that are unlikely to be cost effective.

One such provision is the consolidation option, an exotic financial derivative created by a few paragraphs in the Higher Education Act. This provision allows borrowers to convert their floating rate loans to a fixed rate loan having a rate equal to the average floating rate on their outstanding loans (for some borrowers, it also allows maturity extension).[5] For instance, in 2002

---

1 Total direct and guaranteed loans outstanding, reported by the U.S. Office of Management and Budget, includes traditional credit programs that provide direct loans or loan guarantees for mortgages, higher education, small business, agriculture, trade, energy and other private activities, but it excludes the obligations of Fannie Mae and Freddie Mac.

2 Lochner and Monge-Naranjo (2002) provide a general equilibrium analysis of the interaction between borrowing constraints and human capital formation, and show that government policies that weaken credit constraints can lead to substantial increases in human and physical capital.

3 For critiques of the effectiveness of those policies, see for example, De Fraja (2002), Dynarski (2002), Edlin (1993), Hanushek (1989), and Keane (2002). Gale (1991) points out that many federal credit programs probably have a small real effect on the allocation of credit, in many cases simply crowding out private borrowing and lending.

4 Further, the legal exemption of student loans from dismissal in bankruptcy may alleviate market imperfections caused by restrictions on forward contracting of human capital.

5 Subsequent legislation fixed the interest rate on new Stafford loans at 6.8 percent starting in July of 2006 and made other changes that significantly lowered the value of the consolidation option on newly originated loans. The option remains valuable for borrowers with outstanding floating rate loans. It also may help some fixed rate borrowers

borrowers who consolidated their student loans could lock in an interest rate below 3 percent per annum, a rate that was several percentage points below prevailing long-term Treasury rates at that time.

Several considerations suggest that the consolidation option is not a cost effective mechanism for subsidizing investment in education. For one, the option's value is more likely to be recognized by graduating students than by potential freshmen; entering students are unlikely to be aware of the option's worth because it is complicated to value, and lenders generally do not inform borrowers about it until they leave school and enter repayment. Furthermore the benefits vary widely and stochastically across different cohorts of borrowers depending on interest rate conditions and borrower characteristics.[6] The subsidy is not targeted, with the largest benefits accruing to professional students attending high tuition schools.

An analysis of the consolidation option is of research and policy interest for several reasons. The option serves as an important example of how the combination of programmatic complexity and the rules of federal budgeting can obscure the fiscal costs and risks of federal credit programs, and how the tools of modern financial economics can be used to better inform policymakers and the public about them.[7] At the same time the consolidation option provides a novel setting in which to study how individuals respond to financial incentives, and how their responses vary with the size of an incentive and other factors.

---

who would otherwise be ineligible for longer repayment horizons. Consolidation could reemerge as an important determinant of program cost if lawmakers reinstate floating rates, for instance in response to budgetary pressures.

[6] Being able to lock in a below market interest rate when short-term rates are low might provide borrowers with a partial hedge against economic downturns. However, that benefit is already present for holders of floating rate loans. It is not clear that a recession is a reason to lock in a lower rate permanently. Furthermore, the option's value as a hedge is likely to be small due to the weak correlation between interest rates and individual income shocks. Other alternatives, such as the income contingent repayment schemes that have been made more generous in recent years, are a more targeted way to protect borrowers against income risks.

[7] For example, Falkenheim and Pennacchi (2003) and Pennacchi and Lewis (1994) also use option pricing methods to determine the value of government liabilities. Lucas (2012) surveys the literature on applications of financial economics to valuing government policies.

Evidence on borrower behavior, based on a sample of 700,000 borrower records extracted from the Department of Education's National Student Loan Data System (NSLDS), is used to calibrate the behavioral assumptions embedded in the valuation model. The data (described in Appendix A) suggest that borrowers respond to the time-varying incentives to consolidate, with higher rates of consolidation occurring in years when the option is most valuable.[8] Further, the sensitivity of exercise behavior to interest rates increases with the amount of indebtedness. There also is some evidence of learning: The share of eligible borrowers exercising the option, controlling for the value of doing so, increases over time. Nevertheless, many borrowers do not maximize the value of the option, and we provide estimates of how much money was left on the table at different points in time by comparing observed behavior with the near-optimal strategy over the time period considered of consolidating as soon as possible.

We find that the consolidation option has at times been extremely costly to the government. For instance, ex ante, the cost in 2004 was more than $4 per $100 of new loans originated. Ex post, the option had a cost to the government of $27 billion over the period 1998 to 2005. Furthermore, based on the statistical properties of interest rates, the high costs that were realized were fairly probable and hence represented a predictable risk.

The remainder of the paper is organized as follows: Section 2 briefly describes the federal student loan program and the consolidation option. Section 3 outlines the valuation model and critical modeling assumptions. A Probit analysis of data from the NSLDS summarizes consolidation take-up rates and their sensitivity to borrower characteristics and market conditions, and take-up rates are compared to approximately optimal behavior with respect to interest rate conditions. Section 4 reports cost estimates and their sensitivity to model assumptions and policy variations. Section 5 discusses policy implications and conclusions. Appendices contain more

8 Although many recent studies have provided evidence of suboptimal financial behavior by individuals, few have reported on whether inefficiency decreases with the amount at stake. An exception is Calvet et al. (2006), who find more efficient portfolio allocations on the part of wealthier households in Swedish data.

detailed descriptions of the valuation model, supporting assumptions, and the statistical analysis of borrower behavior.

2. The Federal Student Loan Program

The federal student loan program is one of the largest credit programs operated by the U.S. government. In 2016, outstanding student loans exceeded $1100 billion. This was up from $126 billion in outstanding loans in 1998, reflecting the rapid growth of education costs and strong appetite for consumer debt during that period.

The federal student loan program is complex, and this analysis takes into account its institutional features that significantly affect the cost of the consolidation option. Up until April 2009, the government ran two competing student loan programs – the Federal Direct Student Loan (DL, or direct loan) program, and the older Federal Family Education Loan (FFEL, or guaranteed loan) program.[9] Schools had a choice of which program to adopt. In the former, the government lends funds directly to qualifying students. In the latter, it guarantees loans originated by private lenders against losses from default and pays lenders any shortfall between the rate charged to students and a promised minimum. Loan terms vary with the purpose of the loan; there are loans for undergraduates, graduate students, various types of professional students, and parent loans. Some students qualify for more highly subsidized loans on the basis of income. The terms offered to a given borrower are generally quite similar under the direct and guaranteed programs, but the cost to the government is not (Lucas and Moore, 2010). Loan terms are primarily set by statute under the Higher Education Act, and to a lesser extent by administrative policies.

[9] After that date all new originations are through the direct loan program.

Most loans made under the direct and guaranteed programs are so-called Stafford loans.[10] On Stafford loans originated between October 1998 and July 2006 -- the focus of this analysis -- students pay a variable interest rate based on a three-month Treasury rate that resets annually, plus a spread. The spread varies with the repayment status of the loan: It equals 1.7 percentage points when the student is in school, in the six-month grace period after leaving school, and in periods of deferment; the spread equals 2.3 percentage points otherwise.[11] The interest rate is capped at 8.25 percent. Loan principal is amortized over the term of the loan, starting when the loan enters repayment.

2.1 The Consolidation Option

The consolidation option has three main features. First, it allows borrowers to combine two or more outstanding loans into a new loan. Because each year's borrowing is counted as a separate loan, the option is available to all students that borrow in more than one academic year, including both their undergraduate and graduate studies. Second, borrowers pay a fixed interest rate on the consolidation loan equal to a weighted average of the rates prevailing on the loans they consolidate. An option to switch a loan from a floating rate to a fixed rate is termed a "swaption" in financial markets. Third, some borrowers also receive an extension option, which allows them to extend the maturity of their loans beyond what is otherwise permitted.[12]

---

[10] The Stafford loan program includes most types of federally-backed student loans, including direct and guaranteed loans to undergraduate, graduate and professional students. The other major student loan program, PLUS, provides loans to parents of undergraduate students and to graduate students that have exceeded Stafford borrowing limits. Upon consolidation, Stafford and PLUS loans are reclassified as Consolidation loans, but the conversion does not affect the amount of credit outstanding.

[11] For "subsidized" Stafford loans, which account for approximately 50 percent of the value of outstanding Stafford loans, borrowers pay no interest while they are in school, in grace or in other periods of payment deferment. The subsidized share of new loans is declining and is expected to continue to decline.

[12] Borrowers in the direct program with loans originated between 1998 and 2006 receive little incremental benefit from the extension option, as they can extend loan maturity on their direct Stafford loans with or without consolidating. Many borrowers in the guaranteed program benefited from the extension option, because only borrowers with balances above $30,000 can extend their loans without also consolidating. Consolidation allows all borrowers to extend their repayment period according to the rules governing the direct program. (After June 2006, the extension rules for the direct program are the same as for the guaranteed program.)

Borrowers have the opportunity to consolidate loans for as long as their original Stafford loans remain outstanding. They also can reconsolidate their previously consolidated loans with new Stafford loans. However, because the new fixed interest rate the borrower obtains is a weighted average of the fixed rates on the previously consolidated loans and the floating rate on the new loans, borrowers have the valuable right to lock in the floating rate on each original loan only once.[13]

Consolidation also affects payments from the government to lenders in the guaranteed program. When a guaranteed loan is consolidated by a guaranteed lender (either by the original lender or by a consolidator), the guaranteed rate of return that the lender receives is reduced by at least 0.75 percentage points annually,[14] and lenders pay a small fee to the government at the time of consolidation. The savings to the government from lower annual payments to guaranteed lenders are partially offset by the longer maturity of some consolidation loans and hence a longer period of borrowers receiving a subsidized interest rate. Furthermore, because the government incurs a lower level of administrative costs on direct loans than it makes in lender payments on guaranteed loans, consolidation of guaranteed loans into direct loans generates savings and vice versa when the consolidation of a direct loan is by a guaranteed lender.

2.2 Benefits to students

The consolidation option provides students with several distinct benefits. For a given maturity, the market rate on a floating rate loan is generally lower than the corresponding fixed rate because yield curves tend to slope upward. By allowing students to convert a floating rate into a long-term fixed rate, students can often lock in a favorable spread relative to prevailing

---

13 Students who have not exhausted their loan limits can lower their borrowing rate further by taking new student loans to pay off older consolidation loans in favorable interest rate conditions. We do not take this possibility into account in estimating the value of the consolidation option.

14 Lenders receive a special allowance payment (SAP) equal to the rate on three-month commercial paper (CP) + 2.34 for Stafford loans, plus any floor income (a payment when rates fall below a floor on loans issued prior to June 2006). Consolidation lenders get CP + 2.64 - 1.05 = CP + 1.59. The difference is at least 0.75 percent, and may be higher at times due to a rate floor.

fixed rates in the market. Conversely, when floating rates are high, students can choose to defer consolidation. Hence, the floating-to-fixed-rate conversion option has significant value, which diminishes over time as the loan amortizes and approaches its maturity date.

The term extension is also of value for qualifying borrowers in the guaranteed program for two reasons: it extends the period over which they can borrow at a subsidized rate, and the smaller monthly payment may alleviate liquidity constraints.[15] The annual interest rate subsidy can be approximated by the difference between the estimated fair market rate, which is inferred from the private student loan market, and the government rate.[16] By extending the maturity of the loan, the credit and interest rate subsidies are received over a longer period, increasing their present value. We refer to this component of government cost as the "extension cost." The extension cost varies positively with the credit risk of the borrower. This suggests that all else equal, students with relatively poor credit scores have a greater incentive to consolidate.[17]

Extending loan maturity also lowers monthly payments. For students facing liquidity constraints, a lower monthly payment might be preferred even if it entails a higher interest rate.[18] Guaranteed lenders marketing consolidation loans emphasize this feature, which lowers monthly payments by about 40 percent. Indeed, extensions motivated by liquidity constraints can mitigate the extension cost for two reasons: First, lower payments may help some borrowers avoid default

15 On a fair value or market value basis "unsubsidized" Stafford loans also provide significant interest rate subsidies. Official estimates of program cost understate market value costs because they do not take account the associated market risk, which taxpayers bear. Administrative costs are accounted for separately, which also complicates the assessment of total costs (CBO, 2004, Lucas and Moore, 2010).

16 The private student loan market developed to serve students that have exceeded federal borrowing limits. Private lenders offer rates that vary with borrower credit score and educational institution. To the extent that private loans are taken disproportionately by professional students (e.g., law, medicine, business), they may not be representative of the average Stafford loan credit quality, although the available evidence suggests similar loss rates on federal and private loans (Lucas and Moore (2010)).

17 Empirically, this effect is not observed, probably because other incentives favor consolidation by borrowers with high credit ratings. Lower risk borrowers tend to attend the more expensive institutions and carry the largest loan balances. Hence, they have a large incentive to consolidate.

18 Gross and Souleles (2002) document the propensity of liquidity-constrained consumers to borrow at high interest rates when credit limits on credit cards are increased.

and, second, the desire for lower payments may cause some borrowers to consolidate even when interest rate conditions are unfavorable.

The fact that borrowers face liquidity and other constraints complicates the evaluation of whether they are behaving rationally, and of what the option is worth to them. Hence, we are careful to distinguish between the cost to the government, which we assume is only affected by standard financial considerations, and the benefit to students, which we do not attempt to quantify precisely.

3. Modeling Consolidation Cost

The cost of the consolidation option to the government, and its sensitivity to program rules, economic conditions, and behavioral assumptions, are evaluated using the valuation model that is described briefly in this section and presented in more technical detail in Appendix B. We begin by discussing the critical factors affecting option value, which include the properties of interest rates and borrower behavior.

3.1 Level, Persistence, and Volatility of Interest Rates

The value of the consolidation option depends on the term structure of interest rates, the persistence of interest rates over time, and rate volatility. The cost of consolidation is increasing in the slope of the yield curve. Because the term structure is upward-sloping most of the time, consolidation benefits students under normal interest rate conditions. Because interest rates exhibit a high degree of persistence, the value of the option varies with interest rate conditions at origination. As is generally true for options on interest rates, higher interest rate volatility increases option value.

The well-known model of Cox, Ingersoll, and Ross (CIR) (Cox, et al., 1985) is used to generate a stochastic term structure of interest rates. We employ a two-factor version of the model, following Jagannathan, Kaplin, and Sun (2001) among others. The interest rate model feeds into the stochastic valuation model that is used to determine the cost to the government of the option. It affects both cash flows and discount rates.[19] The model posits that the instantaneous (i.e. short-term) rate follows a stochastic mean-reverting factor process, with volatility proportional to the square root of each factor. As described in Appendix B, the interest rate model is calibrated to match the historical volatility of the short-term Treasury rate and also, depending on the question at hand, to match the historical or contemporaneous shape of the yield curve.

3.2 Borrower Consolidation Behavior

The cost of consolidation to the government depends on when and whether borrowers choose to consolidate. In a frictionless world without borrowing constraints or information frictions, a value-maximizing borrower's decision to consolidate would be strongly dependent on interest rate conditions and rate volatility. As shown in Figure 1, observed consolidation rates are correlated with interest rate conditions, but, as we explain later, borrowers are less responsive to interest rate conditions than a simple optimizing model would predict.[20]

The consolidation option became available in 1994, but it attracted little attention for many years because interest rates were relatively high and often close to the rate cap. In the mid 2000's, as short-term Treasury rates fell to historical lows, the rate of consolidation rose and in some years exceeded 25 percent. The intense marketing by lenders of consolidation loans at that time probably contributed to increased recognition of the option's value and to rising consolidation rates.

[19] The results of the analysis would likely be similar under any of the alternative interest rate models commonly used in valuing interest contingent options.

[20] Defining the consolidation rate consistently over time is complicated by the restriction on in-school consolidations prior to 2005. The restriction implies that a loan originated in a particular year could be ineligible for consolidation for many years while the student completes his or her education. We approximate the loans eligible for consolidation in each year before 2005 as a weighted average of loans from past years, with a linearly increasing weight from 0 percent in the current year to 100 percent in loans 5 years old or older. The graph includes re-consolidations of previously consolidated loans, increasing the reported rates relative to the new consolidations that are the basis of the estimates of cost.

Other options available to borrowers – default, prepayment, deferment and forbearance – all affect the value of consolidation by altering the period of time over which a reduced interest rate is obtained, and so may affect the decision to consolidate. In our benchmark model, we assume that exercise behavior on those other options is unaffected by the consolidation decision. Historically, default rates on student loans were low, averaging less than 2.5 percent per year over the period analyzed (see Moore and Lucas, 2010). Consolidation loans have experienced slightly lower annual default rates on average than Stafford loans. Notably, in the direct program where there is no incremental extension option associated with consolidation, default rates are not lower. This is consistent with the intuition that term extensions reduce defaults by liquidity-constrained borrowers. However, borrowers that consolidate may be less likely to default for other reasons, for example, because they are more financially savvy than non-consolidating borrowers. Furthermore, the availability of income-based repayment, deferment and forbearance provide alternative, and often superior, means for borrowers to avoid default. To account for the possible effect of consolidation on default rates, in the sensitivity analysis we consider a scenario where consolidation reduces default rates and hence the estimated cost of consolidation to the government (see Table 4). However, to the extent that borrowers use the other loan options strategically -- for instance foregoing prepayment or default and seeking loan extensions through deferments and forbearance when consolidation loans are most valuable to them – our base case estimates will understate subsidy costs.

An upper bound on the cost of the consolidation option to the government could be found by assuming that borrowers follow a consolidation strategy that maximizes the theoretical value of the option with respect to interest rates, and that subsidized loans are extended to the maximum maturity upon consolidation. That is, at every point in time, borrowers compare the present value of the future payments if they consolidate immediately with the present value of all future payments if they consolidate at an optimal time in the future (as in a standard American option valuation).

In reality, borrower behavior, and what would constitute optimal behavior, is more complex. As for refinancing home mortgages, the decision to exercise the option is also affected by other factors such as the desire to lower monthly payments. In addition, borrowers may simply fail to make optimal choices for a variety of informational or psychological reasons that have been noted in the literature such as inattention, poor computational skills, or high subjective discount rates. Whatever the cause, those suboptimal choices lower the cost to the writer of the option. Hence, we use a behavioral model of consolidation choice, rather than an optimizing model, in calculating government costs.

Figure 1: Student Interest Rate versus New Consolidation Loans as a Percentage of Unconsolidated Loans

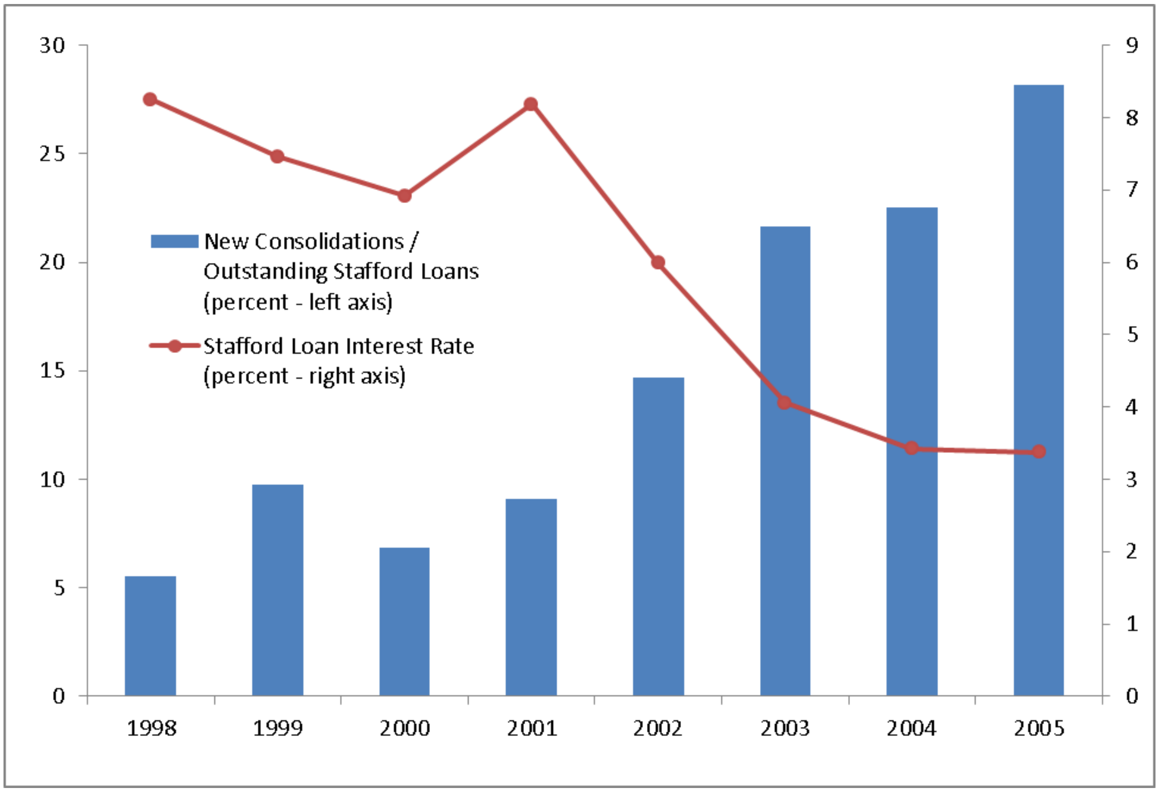

Program rules influence consolidation behavior in several significant ways that are incorporated into the model. Until 2005, borrowers in the guaranteed program could not consolidate loans during the in-school period, creating long and variable lags between origination and consolidation. Historically, the delay could be as short as a year, or in excess of eight years for a loan taken by a freshman who goes on to pursue graduate studies. Temporary relaxation of this rule caused a wave of in-school consolidation activity in 2005.

Two factors provide incentives for borrowers to consolidate their loans early on. For loans issued prior to July 2006, a loophole in the law allowed loans that are consolidated in school or during a grace period to bear rates that are permanently 0.6 percent lower than if consolidation were postponed to a normal repayment period.[21] Further, early consolidation maximizes the principal balance upon which the subsidy is based, because the balance declines as the loan amortizes.

Whether a given borrower would optimally choose to consolidate to take advantage of the extension option depends on unobservable borrower characteristics such as credit score and liquidity constraints. As for the interest rate option, even if the theoretically optimal policy were apparent, actual exercise behavior is likely to be more complex. A fundamental difficulty in estimating the cost of the extension option is that the portion of extensions to attribute to consolidation is uncertain. All students under the direct loan program, and students with high loan balances in the guaranteed program, have the extension option independent of whether they consolidate. Yet historically, most extensions occurred with a consolidation. An unobservable but critical quantity is the percentage of loans that would have been extended in the absence of the consolidation option. The base case counterfactual assumption is that a constant fraction of eligible borrowers extend without the option (see Appendix B.3). When consolidation is available, we assume that, consistent with observed behavior, all consolidating borrowers extend to the maximum allowable maturity. In the sensitivity analysis, we consider the possibility that the extension option involves no incremental cost because the propensity to extend is the same with or without consolidation.

[21] The grace period applies to students while in school, and in the first six months after leaving school. Allowing extension of the grace-period discount has proven to be costly, as shown in the sensitivity analysis.

With these considerations in mind, we estimate the effect of interest rate conditions and loan term extension on consolidation rates using a panel of 700,000 randomly selected borrower records taken from the National Student Loan Data System (NSLDS), administered by the Department of Education. The database comprises multiple linked files containing loan characteristics (e.g., loan program, original maturity and amount), limited borrower characteristics (e.g., school attended, whether undergrad/grad/professional, etc.), and time-series information recording changes in the status of each loan since its origination (in school or grace period, defaulted, consolidated, extended, etc.) . Using the data, we constructed a time series for each borrower in the sample that includes consolidation events, total loan amount outstanding (which determines eligibility for term extension), the consolidation interest rate, and other borrower and loan characteristics.

A Probit regression model is used to estimate the probability that a borrower consolidates in any year.[22] This probability is estimated as a function of a 3-tier categorical indicator of the total loan balance outstanding, (BALCAT: < $20,001; $20,001 to $60,000; >$60,000); the variable interest rate on the original loans (RATE); the length of time the loans have been in repayment (YRS IN REPAY); interactions between these variables; a year dummy (YRDUM) to control for changes in program rules and the availability of information about consolidation; and dummy variables indicating that students borrowed only in the direct program (DIRECT), or that they have loans from both the direct and guaranteed programs (BOTH), controlling for differences in the value of the extension option. Data, variable construction and econometric considerations are described in more detail in Appendix A.

Table 1 below reports the coefficient estimates and standard errors from the Probit regression. Parameter estimates have the expected sign and magnitude, and all of the variables are statistically significant at a 5 percent level except where otherwise indicated. (The reported

22 Knapp and Seaks (1992) use a Probit model similar to ours to analyze the related question of what determines student loan default behavior.

standard errors are cluster robust standard errors with the observation year used as the clustering variable to account for the correlation of consolidation decisions across borrowers within a year.) Consistent with optimizing behavior in the presence of fixed costs, interest rates interact with balance size to produce larger consolidation probabilities among borrowers with larger balances and when rates are lower. For instance, at a variable interest rate of 8 percent, a borrower with a balance of $60,000 or more has a 4 percent probability of consolidating in the year he or she enters repayment, versus a 2 percent probability for a borrower with a balance under $20,000. At a 4 percent interest rate, those probabilities increase to 39 percent and 4 percent, respectively. Figure 2 reports the estimated annual consolidation probability for borrowers with particular characteristics for various interest rates and balance sizes.

We also considered several other specifications not reported here: (1) A Probit augmented with the slope of the yield curve, a long-bond rate, and interactions between the yield curve slope and loan balances; (2) A Logit with the parameters in Table 1; and (3) A Logit augmented as in (1). All specifications had similar explanatory power and produced similar estimated rates of borrower consolidation and subsidy costs.

Table 1: Parameters Estimates and Standard Errors from a Probit Regression of Consolidation Decisions on Covariates.

Parameter	Estimate	Standard Error
BALCAT1	-1.45	0.15
BALCAT2	-0.83	0.19
BALCAT3	-0.63	0.21
RATE * BALCAT1	-0.08	0.02
RATE * BALCAT2	-0.11	0.02
RATE * BALCAT3	-0.13	0.02
YRSINREPAY * BALCAT1	-0.02*	0.01
YRSINREPAY * BALCAT2	0.01*	0.02
YRSINREPAY * BALCAT3	0.03*	0.03
DYRPOST02	0.20	0.10
DYR98TO02	0.03*	0.07
HAS DIRECT LOANS	-0.33	0.12
HAS DIRECT AND GUARANTEED LOANS	1.10	0.10

* Not significant at a 5 percent level

Figure 2: Estimated Consolidation Probabilities by Outstanding Balance and the Consolidation Interest Rate

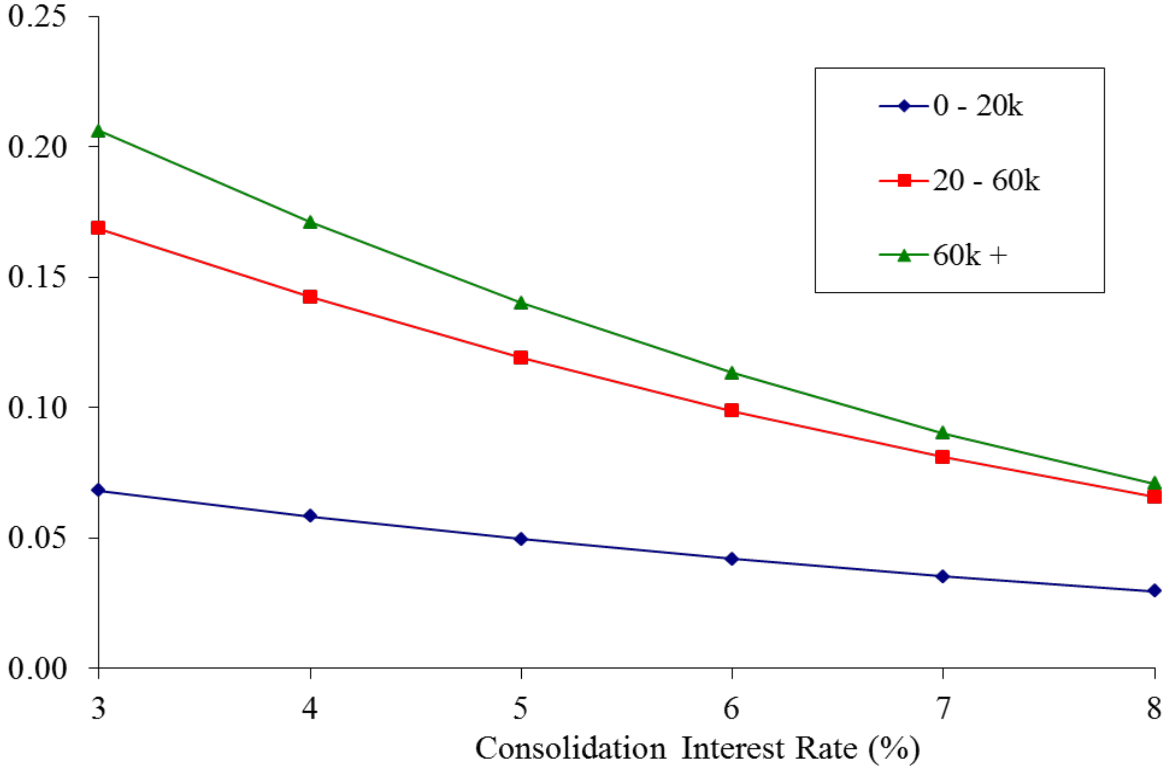

In the prospective cost estimates, we use the original Probit model specification as our model of consolidation behavior. Thus, in repayment, the probability of consolidation is negatively related to the current short-term interest rate. However, even when interest rate conditions are unfavorable a non-negligible fraction of borrowers consolidate their loans. The sensitivity analysis explores alternative behavioral parameterizations with respect to interest rate sensitivity. The typical consolidation behavior summarized by the Probit model also is compared to the implications of a near-optimizing model to evaluate how much of the potential value of the option is captured by borrowers.

3.3 Borrowers Interest Rates

As discussed earlier, the greater the difference between the interest rates charged by private lenders and by the government, the greater the extension cost. The extension cost estimate is based on the typical spread between private and government loan rates and on default and recovery rates estimated from information from the Department of Education, as summarized in Lucas and Moore (2010).

The rate on all federal floating rate student loans is capped at 8.25 percent. The cap affects option value and is reflected in the valuation model. The cap effectively increases the value of the original variable rate loans, particularly in high interest rate periods. At the same time, it reduces the benefit of a consolidation because it limits the worst-case interest rate a borrower would pay without consolidating.

3.2 Valuation Model

The cost of the consolidation option is the difference between the present value of net government cash flows on loans with and without the option. The cost of a loan, with and without the option, is estimated using an arbitrage pricing approach, implemented numerically with Monte Carlo simulation. Projected cash flows are discounted to the date of loan origination along each sample path at risk-adjusted rates.

Cash flows depend on the stochastic path of future interest rates, the program rules, and the behavioral decision rules. Each month a random draw determines the innovation in the short-term interest rate, and the corresponding term structure is derived from the CIR model described earlier. In each simulated period, the fraction of outstanding floating rate loans that consolidate depends on the weighted average rate of interest on the loans in that period and the corresponding consolidation rates predicted by the behavioral model for that level of rates. For the direct program, monthly cash flows to the government depend on whether the student is in school, the current short-term rate if the loan has not yet been consolidated, and a fixed rate based on the short-term rate at the time of conversion if the loan has been consolidated, whether the interest rate cap on the student loan is binding, the average rate of default and prepayment, and an administrative charge.[23] For the guaranteed program, the relevant cash flows are between the government and guaranteed lenders. These cash flows consist of special allowance payments (SAP) that vary with consolidation status, and of compensation to lenders for credit losses.

The discount rates reflect the opportunity cost of providing loan capital to students. There are at least two potential sources of priced risk that affect the discount rate – interest rate risk and credit risk. The approach taken here is to capture the price of interest rate risk by using a risk-neutral representation of the CIR model. For simplicity, we assume that credit risk is orthogonal to interest rate risk.[24] The credit risk premium—the expected excess return that investors demand as compensation for the market risk associated with credit losses--is represented by a spread over risk-neutral rates. The credit risk premium is taken from Lucas and Moore (2010), where it was inferred from data on both federal and private student loans, with adjustments for expected losses and administrative charges. These adjustments yield a credit risk premium of approximately 2 percent per annum.

4. Results

The main focus of the analysis is on prospective, or ex ante, estimates of the cost to the government, which is the decision-relevant metric for program evaluation (rather than realized, or ex post, costs). These ex ante estimates measure the lifetime cost imposed on taxpayers in the year a loan is originated. We estimate these costs for historical interest rate conditions, and examine the sensitivity of the results to parametric and program variations. We also report on ex post realized costs, both as a plausibility check on the more subtle ex ante calculations, and to be

23 The administrative charge on direct loans cancels out when pricing the option as long as the old and new loans have similar maturity.

24 In our data we find default rates on student loans to have a low correlation with interest rates. Other studies (e.g., Collin-Dufresne and Solnik (2001)) find a negative correlation between default-free interest rates and the likelihood of corporate defaults, but we are unaware of published evidence on the correlation for consumer credit.

able to explore the question of to what extent the high ex post costs could have been anticipated by policy analysts.

4.1 Prospective Costs

Our prospective cost estimates take into account consolidation behavior and program rules that affect the time between origination and consolidation and whether a consolidation occurs at all. Under the base-case assumptions, and under interest rate conditions prevailing in each year from 1998 to 2005, the cost of the consolidation option per $100 of Stafford loans originated is shown in Table 2. The table also shows the current three-month T-bill and ten-year Treasury bond rates, which explain much of the difference in cost across years. The cost of the option is largest when the short-term Treasury interest rates are lowest and the yield curve is more steeply upward sloping. Most of the cost arises from the interest rate option; only 10 percent of the cost is attributed to the incremental extension of loan term that consolidation affords.

Table 2: Prospective Consolidation Costs and Interest Rates (1998 – 2005)

Origination Year	3-month T-Bill Rate (percent) (1)	Ten-year T-bond Rate (percent) (2)	Consolidation Cost ($ per $100) (3)
1998	4.98	5.50	1.69
1999	4.57	5.90	1.58
2000	5.69	6.10	1.23
2001	3.49	5.28	2.19
2002	1.70	4.93	2.98
2003	0.92	3.33	4.74
2004	1.27	4.73	3.48
2005	2.97	4.00	2.99

Going forward these costs are expected to be much lower than during the period studied, primarily because of subsequent legislative changes. Interest rates on Stafford loans originated after July 2006 reverted to a fixed 6.8 percent, with consolidation interest rates on those loans as high as 6.875 percent. Although policy makers appear not to have explicitly discussed the impact of various legislative changes on the value of the consolidation option, the rule changes effectively eliminate the option's value for loans originated after 2005. Even so, understanding the cost of this option may prove useful as policymakers contemplate changes to lending terms in the future.

4.2 Sensitivity analysis

We examine the sensitivity of the prospective cost estimates for each origination year to several of the key economic, behavioral, and programmatic parameters: initial interest rate conditions, the volatility of interest rates, the credit risk spread in the discount rate, consolidation behavior, and extension behavior. We also estimate the cost of the grace-period loophole, which allows borrowers a permanent rate decrease of 0.6 percent if they consolidate during the grace period, and the cost of allowing immediate consolidation instead of forcing delay until the borrower enters repayment. All costs are expressed as a present value per $100 of original Stafford loans at origination.

4.2.1 Interest Rates and Spreads

As discussed earlier, the cost of the consolidation option increases with a reduction in short-term interest rates and an increase in the volatility of interest rates. The cost also is affected by the credit spread, but it has two partially offsetting effects: a higher credit spread makes the extension option more valuable because it increases the annual credit subsidy, but it also increases the discount rate, which decreases the present value of the stream of future subsidies. Table 3 shows that moderate changes in the credit spread (which is 2 percentage points in the base case), have only small effects on prospective costs. The column labeled "High Volatility" reports the results of increasing the interest rate volatility parameter of the first factor in the CIR model, σ_1, from 0.2 to 0.3, approximately corresponding to a 50 percent increase in the volatility of the instantaneous interest rate.

Table 3: Prospective Cost Sensitivity to Interest Rate Model ($ per $100 of Stafford originations)				
Origination Year	Benchmark Model (1)	Discount Rate T+1% (2)	Discount Rate T+3% (3)	High Volatility (4)
1998	1.69	1.40	1.76	1.95
1999	1.58	1.36	1.57	1.86
2000	1.23	1.00	1.49	1.47
2001	2.19	1.99	2.35	2.48
2002	2.98	3.01	3.07	3.44
2003	4.74	4.52	4.87	5.07
2004	3.48	3.27	3.23	3.42
2005	2.99	2.77	3.25	3.46

4.2.2 Borrower Behavior

The extent to which borrowers respond to the economic value of the option is estimated with considerable uncertainty. We compare the results of the base case, in which consolidation rates are quite sensitive to interest rates, with the alternative of fixing the consolidation rate at close to its historical average. Comparing the first two columns of Table 4, a fixed consolidation rate increases the cost to the government in some years and decreases it in others. The reason that a fixed consolidation rate is not uniformly less costly is that the benchmark rule was not chosen to be optimal but rather to match observed behavior.

Another uncertainty is the extent to which the option to consolidate increases the propensity to extend loan maturity. Recall that even without consolidation, all borrowers under the direct loan program, and borrowers with combined balances in excess of $30,000 in the guaranteed program, can request a term extension without consolidating. However, consolidation increases the incentive to extend maturity, because the rate reduction obtained with consolidation increases the value of extension. The empirical evidence does not offer much guidance on the size of this effect because the legislation that enabled borrowers to extend loan maturity took effect at about the same time that the consolidation option became highly valuable. In the base case we assume, in the counterfactual scenario where extension is available but consolidation is not, that each year borrowers will extend the term of their loans with a constant probability.[25] An alternative set of cost estimates, shown in Column 3 of Table 4 is more conservative in that we assume borrowers extend their loans in the counterfactual case as frequently as they consolidate them (where the rules governing term extension allow this). Comparison of the first and third column of Table 4 shows that raising the probability of extensions in the absence of consolidation lowers the cost by approximately 30 percent.

If consolidation reduces the incidence of default, consolidation costs will be lower than reported for the benchmark model. In column 4 of Table 4 we report the effect of lowering the default rate and the risk premium for consolidated or extended student loans by 50 basis points each. A lower discount rate is consistent with less risky loans having less market risk.

Table 4: Prospective Cost Sensitivity to Behavioral Variations in Consolidation Rates
($ per $100 of Stafford Originations)

Origination Year	Benchmark Model (1)	No Sensitivity of Consolidation to Interest Rates (2)	More Frequent Extensions in the Counterfactual Case (3)	Consolidation Reduces Default (4)	Near-Optimal Consolidation (5)
1998	1.69	2.21	1.39	0.90	4.3
1999	1.58	2.33	1.45	0.93	5.3
2000	1.23	2.00	1.18	0.62	2.8
2001	2.19	2.60	1.87	1.52	7.8
2002	2.98	3.56	2.70	2.46	11.9
2003	4.74	3.93	3.77	3.73	13.9
2004	3.48	3.61	2.89	2.64	13.0
2005	2.99	2.87	2.49	2.13	8.5

[25] The probability of loan extension is equal to the estimated base rate of consolidation in the Probit model when the borrower interest rate is at the 8.25 percent cap, and so is a proxy for the proportion of students who are anxious to extend maturity even when the new interest rate will be less favorable.

Borrower behavior deviates substantially from what would be predicted by a pricing model where interest rates were the only consideration and there are no transactions costs. For all of the years in our sample, immediate consolidation (i.e. consolidation as soon as permitted under what was then current law) is very close to the optimal strategy.[26] The cost of consolidation under this rule is two to three times that cost under the empirical consolidation rule (Table 4, column 4).[27]

4.2.3 Effects of Related Policies on Consolidation

Allowing borrowers to consolidate during their grace period at the lower grace period interest rate instead of the standard Stafford interest rate significantly increases costs. The increase averaged $0.78 over the 1998 to 2005 period, and would have been significantly higher but for the frequent failure of borrowers to take advantage of the provision. [28]

In May 2005, the Department of Education appeared to change its position on in-school consolidations in the guaranteed program, writing that, contrary to popular opinion, in-school consolidation is permissible under current law.[29] Together with favorable interest rate conditions that were expected to be transitory, the decision triggered a wave of consolidation activity. Had in-school consolidation been a common practice over the entire period of 1998 to 2005, we estimate that it would have increased the prospective cost of the consolidation option by an

[26] The naïve rule of consolidating immediately usually maximizes the value of the option because deferring is only valuable when the yield curve is downward sloping. Even with a downward sloping yield curve, the availability of term extension can make waiting more costly than consolidating immediately.

[27] The difference between the base case and optimal cost cannot be interpreted as the money left on the table by borrowers, since part of the difference is due to the effect on the amount paid to guaranteed lenders.

[28] The grace period can be reactivated by returning to school for additional education, increasing the grace period's average value.

[29] The ruling was not a response to any legislative change, but was rather a response to inquiries by guaranteed lenders. In-school consolidation was already available to direct program borrowers at that time.

average of $0.55 per $100 of Stafford loans originated.[30] Once again, these costs would have been significantly larger had borrowers more fully exploited those advantageous program terms.

Consolidations of troubled loans are much more likely to occur under the direct loan program. This may be because guaranteed lenders can avoid dealing with some troubled loans by declining to consolidate them, causing those students to turn to the direct program for consolidation. We do not try to quantify the effect on cost, the sign of which is uncertain. For a typical loan the costs of administering a loan in the direct program are lower than the statutory payments to lenders in the guaranteed program, but if the direct loan program is adversely selected against on troubled loans, the costs of administering those loans could be higher.

4.3 Realized costs

In this section we report the realized annual cost to the government from 1998 to 2005. The calculation is based on the intrinsic value of realized consolidation volume in a given year. The intrinsic value is the difference between the present value of cash flows for a consolidated loan and the present value of cash flows for an unconsolidated loan from the date that the consolidation occurs.[31] Table 5 shows that the costs vary considerably from year to year, ranging from a gain to the government of $5.44 per $100 of loans consolidated in 2000 to a loss of $18.73 per $100 of loans consolidated in 2003. In total dollar terms, the costs peaked in 2004, at $8.8 billion. The wide variation in cost over time reflects changes in interest rate conditions and borrower behavior. The cost varies more over time than in the prospective cost calculations (Table 2) because for the prospective costs the delay between origination and consolidation provides time for interest rates to revert toward their long-run average, making initial interest rate

30 We assume that in-school consolidation also is sensitive to current interest rate conditions, but that, for any level of interest rates, in-school consolidation is less likely than after graduation.

31 In each case, we assume that borrowers take the maximum term extension available to them. For each program year, we take as a starting point the interest rate conditions in that year, and simulate the loan cash flows and discount rates over the life of the representative set of loans. The behavioral parameters are not relevant, however, because consolidation is assumed to have occurred in that year.

conditions less important.[32] Also, the prospective cost is lower on average because it includes loans that never consolidate or that have a long delay until consolidation. The scatter plot in Figure 3 shows that borrowers consolidate at higher rates when the subsidies they receive from the government are greater, although a considerable number consolidate even when it results in cost savings to the government.

4.3.1 Were High Costs Predictable?

An interesting question is to what extent the high costs of this policy in some years could have been predicted? In other words, was there just a particularly unlucky draw of interest rates, or could the magnitude of realized costs have been anticipated by lawmakers?

A simple back-of-the-envelope calculation suggests the possibility that large costs should have been easy to predict. Historically a 2 percent yield spread between 3-month Treasury and 10 year bond rates has prevailed at numerous points in history. Locking in a two-percent annual interest rate advantage on a 20-year amortizing loan produces a cost estimate of $12 per $100 consolidated, which is in line with the predictions of our more complicated option-pricing model.

To establish more rigorously that the high realized costs were to be expected, and to demonstrate the plausible range of variation of future costs, the CIR model, calibrated to interest rate conditions in 1998 and to historic interest rate volatility, is used to generate the distribution of the term structure of interest rates and corresponding consolidation costs for subsequent years. Figure 4 shows the distribution of cost per $100 consolidated, at the time of consolidation, based on 2000 draws of the forward distributions of term structures. Although the very high costs realized in 2003 and 2004 were unlikely ex ante, the calculations suggest that there was a 23 percent chance of costs in excess of $8 per $100 of loans consolidated at the end of four years.

[32] The interest model implies that in the steady state the 3-month yield is 3.5 percent, the 10-year yield is 4 percent, the annual volatility of the 3-month yield is .8 percent and the annual volatility of the 10-year yield is .5 percent. In 2003, when the 3-month yield was .92 percent, the half-life to the steady state is 15 months.

Table 6 contains descriptive statistics for the distribution of intrinsic values at other terminal horizons, in each case starting from initial interest rate conditions in 1998. One way that this risk exposure could have been avoided, while still allowing conversion to a fixed rate, is if the consolidation rate had been indexed to current long-term fixed rates rather than to current short-term rates.

Table 5: Consolidation Volume and Estimated Cost (1998 – 2005)

Consolidation Year	Consolidation Volume (millions of $)	Consolidation Cost (millions of $)	Consolidation Cost ($ per $100)
1998	12,312	21	0.17
1999	10,217	409	4.01
2000	15,466	-841	-5.44
2001	26,411	1,099	4.16
2002	39,283	3,480	8.86
2003	43,770	8,198	18.73
2004	55,272	8,800	15.92
2005	29,203	5,832	7.60

Figure 3: Percent of Eligible Loans Consolidating versus Government Consolidation Cost, 1998 – 2005.

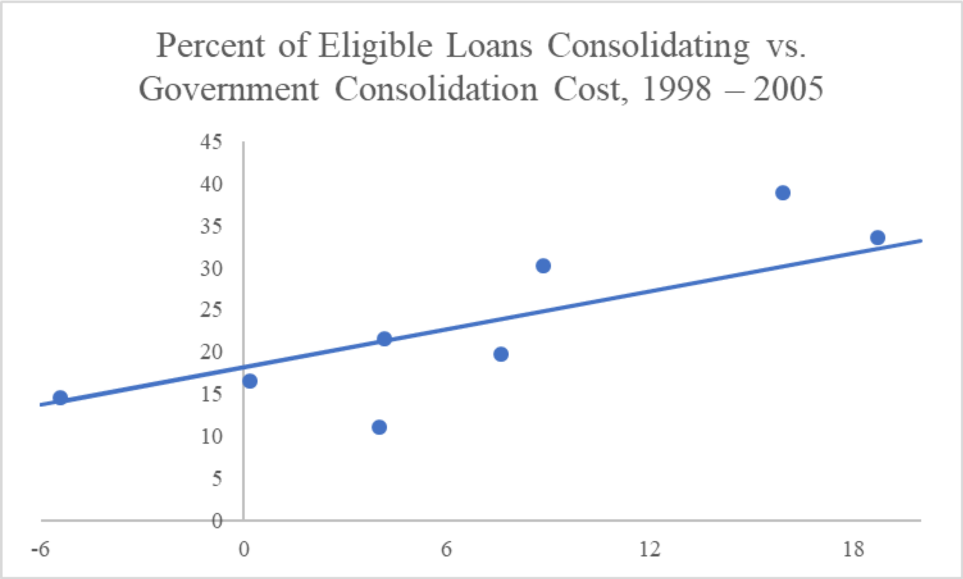

Figure 4: Simulated Distribution of Intrinsic Value, Simulated Four Years Forward From Interest Rate Conditions in 1998

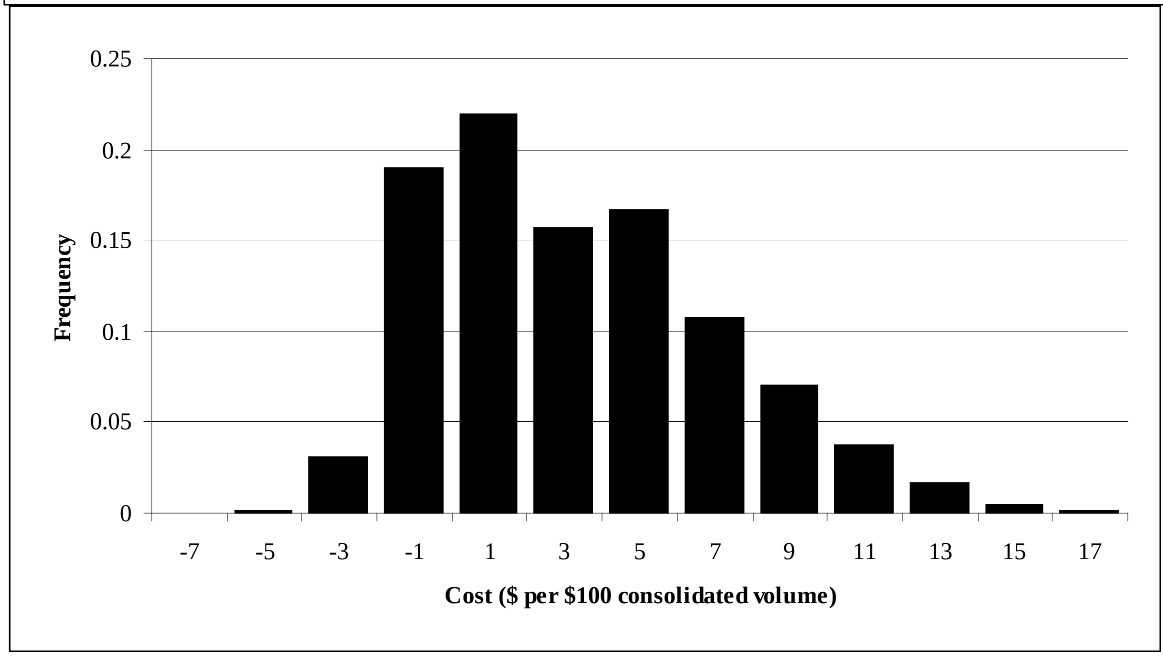

Table 6: Intrinsic Cost of Consolidation ($ per $100 loans consolidated)

	Number of Years				
	2	4	6	8	10
Mean	3.33	3.51	3.58	3.56	3.66
Median	3.01	2.95	2.90	3.01	3.10
Standard Deviation	3.40	3.77	3.97	3.92	3.96
Relative Skewness	0.61	0.70	0.71	0.68	0.65

Notes: Initial interest rate conditions are set at 1998 levels. Interest rates are simulated the set number of years forward indicated.

5. Discussion and Conclusions

In this study, we develop a model to value the student loan consolidation option, and use it to estimate the cost to the government of this provision under a variety of assumptions about economic conditions and borrower behavior. The analysis implies that, between 1998 and 2005, the option had an ex post cumulative cost to the government of about $27 billion. Although this estimate is sensitive to model assumptions, the estimate remains strikingly large for a wide range of alternative assumptions. Conferring a subsidy that is opaque and takes many years to be realized appears to be a less effective means of supporting investment in higher education than more transparent alternatives. The option confers the greatest benefit on those cohorts who happen to graduate when the yield curve is steep and rates are low, to professional students with the largest loan balances, and to borrowers with the sophistication to manage their loans efficiently.

An important behavioral finding that emerges from the analysis is the high take-up rate by borrowers for whom the option is most valuable. Comparing observed consolidation behavior with an optimizing strategy suggests that many borrowers had sufficient understanding of the option’s value to act upon it. The extent to which this was due to a fundamental understanding of the value of the swaption, or rather to the intense marketing efforts of consolidators or to the side-effect of lower payments through maturity extension, however, remains an open question.

Appendix A: Data and Estimation

The National Student Loan Database System (NSLDS) is maintained by the Department of Education. For each borrower, the database contains a record of all current and previous loans, which avoids any potential survival bias. In our sample, each loan record contains the loan amount disbursed, the amount outstanding at January 2006, the type of loan, and current status of the loan. In addition, each loan record contains a full history of loan status changes and associated dates. For each borrower, we recorded an annual status for each year the borrower had loans outstanding. The status is mapped into seven categories, four of them open and three closed. The open status categories, in which the loan is still open and collectible, are:

- G: In school or in the six-month grace-period. Usually the borrower does not make payments on a loan in this status.
- F: In forbearance or deferment. The borrower does not make payments in this status.
- R: In repayment.
- d: In default. A borrower who fails to make the prescribed minimum repayments for more than 270 days is in default.

The closed status categories are:

- D: Collected, negotiated, or written off.
- P: Paid in full.
- C: Consolidated. Some loans that were coded as paid in full in the data were actually consolidated. We corrected this error where possible by linking paid-in-full records to the corresponding origination of a consolidated loan for that borrower.

We analyzed the transition of loans from the repayment status (R) to consolidation status (C) at the borrower level by constructing an annual time series of consolidation events and associated controls for each borrower. This gave us approximately 3 million pooled observations on the following variables:

1. CONSOL: Consolidation indicator variable taking the value 1 if the borrower consolidates currently outstanding loans in a given year, 0 otherwise.
2. BALCATx: Dummy variable indicating the estimated total value of the borrower's outstanding original loans. The categories are as follows: x = 0, balances of $0 to $20,000; x = 1, balances of $20,001 to $60,000; and x = 2, balances exceeding $60,000. Because we do not observe the principal outstanding for every loan in every year, we estimated it from the disbursed amount of each loan under assumed interest rates and an assumed amortization schedule for loans in repayment.
3. RATE: The weighted average interest rate on the borrower's outstanding original loans. This is the rate a borrower would lock in if he or she consolidated (ignoring any rate incentives he or she may have received from a guaranteed lender).
4. YRS IN REPAY: Weighted average time the loans have been in repayment.
5. Dummy variables for the program the loans originated in.
6. Dummy variables for the year of the decision, proxying for the changes over time in the information borrowers receive about consolidation.

The Probit model described earlier was used to estimate the coefficients of each borrower's annually observed decision to consolidate their loans given the covariates 2 through 6 (and some interactions between them). The covariates are intended to pick up the economic incentives to consolidate and time variation in decision making. In particular, consolidation is more likely for borrower with larger balances and when the borrower's lock-in rate is low relative to market rates (market rates were relatively stable over this period so we did not include them).

An alternative for estimating the probability of consolidation would be to use a duration model, such as the Cox proportional hazards framework used in Deng, Quigley and Van Order's (2000) study of mortgage prepayment and default. In fact, our Probit regression can be interpreted as a discrete time approximation of a duration model. [33] Estimating coefficients for time varying covariates (such as interest rates) is more straightforward in discrete than continuous time.

Finally, we have not attempted to account for individual specific random effects because we believe the impact on our estimates of the cost of consolidation would be minor.[34] Deng et al use a competing hazards specification of the Cox model in order to jointly estimate the intensity of default and prepayment, controlling for unobserved individual heterogeneity. A more complete specification of our model would capture the transition between various states including consolidation, prepayment, default and forbearance as well as individual heterogeneity. However, as the Deng et al results show, parameter estimates are relatively robust to these omissions.[35]

---

[33] The Cox proportional hazard model is the limiting case of a binary logit model on panel data with time varying dummy variables (see Amemiya, 1985, pp 433 - 435). Han and Hausman (1990) use discrete time probit and logit approximations interchangeably in their exposition of competing hazards estimation and the two models generally produce similar estimates.

[34] An example of a controllable temporal dependence that has been excluded from the specification is the previous year's incentive to consolidate.

[35] In the absence of individual level heterogeneity, parameters separately estimated from each hazard are identical to parameter estimates from a jointly estimated competing hazard model.

Appendix B: Model and Methodology

To estimate the cost of the consolidation option we:

1. Use the Cox-Ingersoll-Ross model to simulate a risk-neutral path of interest rates for various maturities. It is the basis for discounting credit-risk-free nominal cash flows in the model;
2. Compute cash flows: In the prospective cost estimates, the government's stream of cash flows is computed across each interest rates path, the default state of the loan, and stochastic consolidation and extension behavior. In the intrinsic value calculations, cash flows across each interest rate path and the default state are computed as in the prospective case. However, the date of consolidation, and hence current interest rate conditions, are taken as known and given. Consolidation occurs at the beginning of repayment, and the maximum allowable extension is assumed.
3. Discount the government's default contingent cash flows using a state price deflator to account for default risk and the simulated risk free interest rate to account for time value of money;
4. Compute the cost of consolidation to the government for each loan type (direct or guaranteed) by taking the difference between the present value of cash flows with and without the option to consolidate, averaging across the simulated risk-neutral interest rates and stochastic consolidation and extension behavior; and
5. Aggregate cash flows for a representative set of loans made under the two lending programs.

B.1 Interest Rates

We adopt the Cox-Ingersoll-Ross (CIR) model to simulate future paths of Treasury rates. For discounting risky loan cash flows, we adjust rates upward for the systematic component of default risk, as described in section B.4. In the CIR model, the instantaneous interest rate, $R(t)$, is the sum of a constant and n factors, $z_i(t)$, for $i = 1,\ldots,n$, the state variables in the model:

$$R(t) = \overline{R} + \sum_{i=1}^{n} z_i(t) \tag{1}$$

The constant $\overline{R}$ takes care of the well-known difficulty of the standard CIR model (without constant) to fit all term structure shapes with strictly positive factors. Each factor obeys a mean reverting square root process:

$$dz_i(t) = \kappa_i \left[\theta_i - z_i(t)\right] dt + \sigma_i \sqrt{z_i(t)}\, dZ_i(t) \tag{2}$$

where θ_i is the mean reverting rate, κ_i is the speed of mean reversion, σ_i is the volatility parameter, and $dZ_i(t)$ is a standard Weiner process independent across factors.

Under the risk neutral (or equivalent martingale) measure:

$$dz_i(t) = \overline{\kappa}_i \left[\overline{\theta}_i - z_i(t)\right] dt + \sigma_i \sqrt{z_i(t)}\, dZ_i(t) \tag{3}$$

where

$$\overline{\kappa}_i = \kappa_i + \lambda_i \tag{4}$$

and

$$\overline{\theta}_i = \frac{\kappa_i \theta}{\kappa_i + \lambda_i} \tag{5}$$

λ_i is the constant market price of risk for factor, i. The price at time t of a zero coupon bond with unit coupon and expiry at time T is:

$$p(t,T) = e^{-\overline{R}(T-t)} \prod_{i=1}^{n} A_i(t,T)\, e^{-B_i(t,T) y_i(t)} \tag{6}$$

where

$$A_i(t,T) = \left[\frac{2\gamma_i \exp\left[(\gamma_i + \overline{\kappa}_i)(T-t)/2\right]}{(\gamma_i + \overline{\kappa}_i)\left[\exp\left[\gamma_i (T-t)\right] - 1\right] + 2\gamma_i} \right]^{2\overline{\kappa}_i \overline{\theta}_i / \sigma_i^2} \tag{7}$$

$$B_i(t,T) = \frac{2\left[\exp\left[\gamma_i(T-t)\right]-1\right]}{(\gamma_i+\bar{\kappa}_i)\left[\exp\left[\gamma_i(T-t)\right]-1\right]+2\gamma_i} \tag{8}$$

and

$$\gamma_i = \sqrt{\bar{\kappa}_i^{\,2} + 2\sigma_i^{\,2}} \tag{9}$$

The yield to maturity, y, of a zero-coupon bond maturing at T is

$$y(t,T) = \frac{-\ln p(t,T)}{T-t} \tag{10}$$

Jagannathan, Kaplin and Sun (2001) estimate the factors from the two-factor model, using weekly LIBOR rates of various maturities from 1995 through 1999 as shown in Table B-1.

Table B1: Parameters for the Cox-Ingersoll-Ross Two-factor Interest Rate Model

$\bar{R}$ = -0.229

Factor	κ	θ	σ	λ
1	0.392	0.272	0.0153	-0.00038
2	0.0532	0.0162	0.0430	-0.0592

Under these parameters, factor 1, with the stronger degree of mean reversion, drives the gap between long- and short-term rates, and factor 2 determines long-term rates. We subtract 20 basis points from $\bar{R}$ to reflect the average spread between three-month Treasury and LIBOR yields.

For each Monte Carlo run, initial levels of the state variables are calibrated to fairly price an initial three-month Treasury bill and ten-year Treasury bond. For each simulation, the instantaneous rate is sampled monthly for as many months as the maximum maturity of the student loan, using a discrete Euler approximation of the risk-neutral process in equation (3).[36]

B.2 Government Cash Flows in the Direct and Guaranteed Programs

The cost of consolidation for a particular loan is the difference in the present value of government cash flows when a student has access to the consolidation option and when the student does not have access. The cash flows themselves depend on the evolution of the principal balance and interest charged on the student loans, which we outline now.

Loans originate at time 0, begin repayment at time T^R, and have a maturity of T^M, so the loan is repaid in T^R+T^M months. T^M depends on whether the consolidation option is exercised or, in the counterfactual case, the loan term is extended. $T^{C,j}$ is the month that borrower consolidates in simulation j ($T^{C,j} > T^R+T^M$ if the borrower never consolidates). The original maturity of Stafford loans is 10 years. Section B.3 describes the stochastic rules governing consolidation and extension.

Interest accrues on outstanding principal every month. The interest rate prior to consolidation is linked to the yield on the three-month Treasury bill on May 30 each year, and is fixed for a year. Thus the reference rate, $\tilde{R}^j_{12k+i}$, in month i of year k is:

$$\tilde{R}^j_{12k+i} = \exp\left[y^j \left(k, k+3/12\right)\right] - 1, \forall\, i = 1,...,12,\ k = 0,1,2,... \qquad (11)$$

The borrower pays an interest rate that depends on whether he or she is in repayment and whether he or she has consolidated. Extension of loan term without consolidation does not affect the interest rate. Before consolidation and before entering repayment the student rate, $R^j_{S,t}$, is the lower of the reference rate plus 1.7 percent and the interest rate cap of 8.25 percent,

[36] The evolution of each rate factor over a discrete interval of time obeys a noncentral chi-squared distribution. The standard Euler approximation with random draws taken from a normal distribution and truncated to prevent non-positive factor values produced only negligible bias for our interest rate parameters (truncation occurred in less than 0.1 percent of draws) and was ten times faster to calculate.

$$R_{S,t}^{j} = \min\left[\tilde{R}_{t}^{j} + 1.7\%, 8.25\%\right], \ \forall\, t < T^{R} \text{ and } t \le T^{C,j} \tag{12}$$

After entering repayment but prior to consolidation, the student rate is the lower of the reference rate plus 2.3 percent and the interest rate cap:

$$R_{S,t}^{j} = \min\left[\tilde{R}_{t}^{j} + 2.3\%, 8.25\%\right], \ \forall\, t \ge T^{R} \text{ and } t \le T^{C,j} \tag{13}$$

After consolidation, the student rate is fixed at the rate prevailing at the time of consolidation, $T^{C,j}$:

$$R_{S,T^{c}+k}^{j} = R_{S,T^{c}}^{j} \quad \forall\, k > 0 \tag{14}$$

The monthly compounding student rate, $r_{S,t}^{j}$, is:

$$r_{S,t}^{j} = \left(1 + R_{S,t}^{j}\right)^{\frac{1}{12}} - 1 \tag{15}$$

Let P_t^j denote the evolution of principal (prior to default) over month t in each simulation j. Given an initial principal of $P_0^j = P_0$, monthly principal evolves according to:

$$P_{t+1}^{j} = P_{t}^{j}\left[1 + r_{S,t}^{j}\right] - A_{t+1}^{j} \tag{16}$$

The prescribed monthly payment, A_t^j, depends on the loan status, and is based on amortizing the principal at the current interest rate over the remaining life of the loan:

$$A_{t+1}^{j} = \begin{cases} \dfrac{P_t^{j} r_{S,t}^{j}}{1 - \left(1 + r_{S,t}^{j}\right)^{T^{R} + T^{M} - t}}, & t \ge T^{R} \\ 0, & t < T^{R} \end{cases} \tag{17}$$

In the direct program, the government’s cash flows on performing loans are the student loan payments less any administrative fees:

$$A_t^{j} - fP_t^{j} \tag{18}$$

where f is the proportional administrative fee. The fee is 0.75 percent per annum in the benchmark calibration, reflecting typical servicing and other administrative costs. In default, the government recovers in proportion to the present value of remaining payments.

In the guaranteed lending program, the government cash flows are the quarterly payments to lenders – the SAP less any consolidation fee paid by lenders to the government – while the loan is in good standing, and the lump sum payment of outstanding principal and accrued interest in the event of default. We ignore administrative costs because they are largely borne by the guaranteed lender.

The quarterly SAP is the difference between the student rate and the three-month commercial paper rate plus a spread, but has a floor of zero. We assume the annualized three-month commercial paper rate, R_C, tracks the T-bill rate with a 20 basis point spread:

$$R_{C,t}^{j} = \exp\left[y^{j} \left(t, t+3/12 \right) \right] + .002 - 1, \forall\, t = 1, 2, ..., T \quad (19)$$

Absent default, the government cash flow in each month is the SAP less any consolidation fee paid from lenders to the government, or 1.05% of principal. We denote the net guarantee payment from the government while the loan is in good standing by G:

$$G_t^j = \begin{cases} -P_{3k}^j \max\left[R_{C,3k}^j + 1.74\% - R_{S,3k}^j, 0 \right] / 4, & 3k < T^R \text{ and } 3k < T^{C,j} \;\; \forall\, k = 0,1,2,... \\ -P_{3k}^j \max\left[R_{C,3k}^j + 2.34\% - R_{S,3k}^j, 0 \right] / 4, & 3k \geq T^R \text{ and } 3k < T^{C,j} \;\; \forall\, k = 0,1,2,... \\ -P_{3k}^j \left(\max\left[R_{C,3k}^j + 2.64\% - R_{S,3k}^j, 0 \right] - 1.05\% \right) / 4, & 3k \geq T^{C,j} \;\; \forall\, k = 0,1,2,... \\ 0, & otherwise \end{cases} \quad (20)$$

In default, the government pays the outstanding principal, P_t^j, to the lender, assumes the loan, and recovers in proportion to the present value of the remaining outstanding payments. The default and recovery rates used in the calibration are described in Section B.5.

Several additional factors affect the timing and magnitude of student loan cash flows but we omit them from the analysis because we expect their impact on consolidation cost to be small.

Voluntary prepayment has the effect of reducing the cost to the government by shortening effective loan maturity, but deferment and forbearance, by delaying payment, increase it.[37] Subsidized Stafford loans offer a higher subsidy than so-called unsubsidized Stafford loans, because interest does not accrue while the student is in school. In the prospective cost estimates, this affects the principal balance at the time repayment begins, making it smaller than it otherwise would be. In the analysis, all loans are unsubsidized loans.

B.3 Timing of Consolidation and Extension

We use a Probit rule for the intensity of consolidation of a given loan type, using the coefficient point estimates from Table 1 in the text. Specifically, consolidation is decreasing in the student interest rate and decreasing in the time since repayment begins. We assume borrowers consolidate loans during the grace period consistent with the rule for consolidation at other times, but cannot consolidate at all while they are in school.[38] Thus, the annualized probability of consolidation, $q_{C,t}$, at month t is

$$q_{C,t}^{j} = \begin{cases} 0, & t < T^{R} - 6 \\ \Phi\left(\beta_1 + \beta_2 R_{S,t}^{j}\right), & T^{R} - 6 \leq t < T^{R} \\ \Phi\left(\beta_1 + \beta_2 R_{S,t}^{j} + \beta_3 \max\left(\left[t - T^{R}\right]/12, 0\right)\right), & t \geq T^{R} \end{cases} \tag{21}$$

where Φ is the cumulative standard normal distribution function and β_1 through β_3 are the loan type specific parameters reported in Table B-2 below (computed from the Probit estimates in Table 1 after substituting all static variables). We also apply the rule for consolidation to capture extension behavior in the counterfactual case with no consolidation. In the benchmark simulation

---

37 Although borrowers should theoretically never exercise the option to prepay a loan when the interest on the loan is below the rate on their best alternative, in fact they frequently do. Offering prepayment thus can be a source of cost savings for the federal government.

38 For the 2006 academic year, borrowers were allowed to consolidate during their in-school period.

without consolidation, we assume a constant rate of extension equal to the rate of consolidation when the interest rate is at the 8.25 percent cap. This implies the rate of extension is invariant to the level of the student rate. If anything, standard theory would predict that the sensitivity of extension behavior to short-term interest rates should be positive because, relative to market rates, the present value of an extension increases when interest rates are closer to the cap.

Consolidation allows students to extend their original 10-year loans to as long as 30 years. To estimate the cost of the option to the government, several factors relating to the maturity change are accounted for. One such factor is that annual administrative expenses such as servicing and collection will be incurred for a longer period. Losses from defaults also must be treated consistently. The performance history of consolidated and unconsolidated loans suggests a higher cumulative default rate for consolidation loans. On a per annum basis, the default rates are similar, so we assume the same 2 percent per annum default rate for all loans.

Table B-2: Parameters Determining the Annual Frequency of Consolidation

	β_1	β_2	β_3
Consolidation: Eligible for 10-year term (balance $0-$20k)	-1.14	-0.08	-0.017
Consolidation: Eligible for 20-year term (balance $20k - $60k)	-0.37	-0.13	0.007
Consolidation: Eligible for 30-year term (balance exceeds $60k)	-0.15	-0.13	0.02
No Consolidation: Eligible for 10-year term (balance $0-$20k)	-1.44	0.00	-0.017
No Consolidation: Eligible for 20-year term (balance $20k - $60k)	-1.33	0.00	0.007
No Consolidation: Eligible for 25- or 30-year term(balance > $60k)	-1.22	0.00	0.02

B.4 Aggregate Cost of Consolidation

The forward-looking measure of consolidation cost is the sum of the cost of consolidation at origination for a representative cohort of loans issued in each year of the program. The cost of consolidation at origination of any individual loan in a cohort is:

- The present value of government cash flows under a loan originated when the borrower is eligible to consolidate; less

- The present value of cash flows under a loan originated where the borrower is not eligible to consolidate but may be eligible for term extension.

Each year's cohort of originated loans comprises direct and guaranteed loans, loans from borrowers with different total balances (determining their eligibility for various loan maturities under consolidation and term extension), and loans with different repayment start times based on how close to graduation the borrower is. In each year and roughly consistent with data, we assume that the fraction of volume originated in the guaranteed program is 75 percent and the fraction in the direct loan program is 25 percent. For guaranteed loans, the assumed distribution of loan balances is such that:

- 30% of originated volume is not eligible for any term extension. The borrower must pay off the loan over the original 10-year term.
- 30% of originated volume is eligible for 20-year term extension if consolidation is available. In the counterfactual where consolidation is not available, borrowers are not eligible for term extension.
- 40% of originated volume is eligible for 30-year term extension if consolidation is available. In the counterfactual, a term extension of 25 years is available.

For the direct program, borrowers have a symmetric opportunity to extend the maturity of their loans with and without consolidation:

- 30% of originated volume is not eligible for any term extension.
- 30% of originated volume is eligible for 20-year term extension.
- 40% of originated volume is eligible for 30-year term extension.

In a given year, repayment of each loan originated begins only after the borrower finishes school. We assume that the loans originated have repayment times uniformly distributed between one and five years. The aggregate cost of consolidation is the sum of the individual costs of consolidation across both direct and guaranteed programs, categories of maturity extension eligibility, and repayment start times.

For the intrinsic value calculations, we compute the cost of consolidation at the time the loans are consolidated. For simplicity, we assume the distribution of loan amounts and eligibility for term extension is the same as for the forward-looking estimates. On the other hand, we no longer treat term extension in the absence of consolidation as stochastic, and simply assume that borrowers take advantage of any term extension available, both in fact and in the counterfactual.

References

Amemiya, T (1985), *Advanced Econometrics*, Harvard University Press.

Calvet, L.E., J.Y. Campbell, and P. Sodini (2006), "Down or Out: Assessing the Welfare Costs of Household Investment Mistakes," manuscript, Harvard University

Collin-Dufresne, Pierre and Bruno Solnik (2001), "On the Term Structure of Default Premia in the Swap and LIBOR Markets," *Journal of Finance* 56, 1095-1115

Congressional Budget Office (2004), *Estimating the Value of Subsidies for Federal Loans and Guarantees*, Washington, D.C., August.

Congressional Budget Office (2006), *The Cost of the Consolidation Option for Student Loans*, Washington, D.C., CBO Paper, May.

Cox, J. C., J. E. Ingersoll and S. A. Ross (1985), "A Theory of the Term Structure of Interest Rates," *Econometrica*, vol. 53, pp. 385-408.

De Fraja, G. (2002), "The Design of Optimal Education Policies," *The Review of Economic Studies*, vol. 69, no. 2, pp. 437-466.

Deng, Y., J.M. Quigley and R. Van Order (2000), "Mortgage Terminations, Heterogeneity and the Exercise of Mortgage Options," *Econometrica*, vol 68 issue 2, pp. 275-307.

Dynarski, S. (2002), "The Behavioral and Distributional Implications of Aid for College," *The American Economic Review* vol. 92, no. 2, pp. 279-285.

Edlin, A.S. (1993), "Is College Financial Aid Equitable and Efficient?" *The Journal of Economic Perspectives*, vol, pp. 143-158.

Falkenheim, M. and G. Pennacchi (2003), "The Cost of Deposit Insurance for Privately Held Banks: A Market Comparable Approach," *Journal of Financial Services Research*, vol. 24, no. 213, pp. 121-148.

Han, A and J. A. Hausman (1990), "Flexible Parametric Estimation of Duration and Competing Risk Models" *Journal of Applied Econometrics*, vol. 5, pp 1-28.

Gale, W.G. (1991), "Economic Effects of Federal Credit Programs," *The American Economic Review*, vol. 81, no. 1, pp. 133-152.

Gross, D.B. and N.S. Souleles (2002), "Do Liquidity Constraints and Interest Rates Matter for Consumer Behavior? Evidence from Credit Card Data," *Quarterly Journal of Economics*, vp;/ 117, no. 1, pp.149-185.

Hanushek, E.A. (1989), "Expenditures, Efficiency, and Equity in Education: The Federal Government's Role," *The American Economic Review*, vol. 79, no. 2, pp. 46-51.

Jagannathan, R., A. Kaplin, and S.G. Sun, (2001), "An Evaluation of Multi-Factor CIR Models Using LIBOR, Swap Rates, and Cap and Swaption Prices," *Journal of Econometrics*, vol. 116(1-2), pp. 113-146.

Keane, M.P. (2002). "Financial Aid, Borrowing Constraints, and College Attendance: Evidence from Structural Estimates," *The American Economic Review* vol. 92, no. 2, pp. 293-297.

Knapp, L. G. & Seaks, T. G. (1992), "An analysis of the probability of default on federally guaranteed student loans" *The Review of Economics and Statistics*, 74(3), 404-411.

Lochner, L. and A. Monge-Naranjo, (2002), "Human Capital Formation with Endogenous Credit Constraints," NBER Working Paper 8815, National Bureau of Economic Research, Inc.

Lucas, D. (2012), "Valuation of Government Policies and Projects," *Annual Review of Financial Economics*.Vol 4, pp. 39-58.

Lucas, D. and D. Moore (2010). "Guaranteed vs. Direct Lending: The Case of Student Loans," in *Measuring and Managing Federal Financial Risk* (edited by D. Lucas), University of Chicago Press.

Maddala, G. S. (1987), "Limited Dependant Variable Models Using Panel Data" *The Journal of Human Resources*, vol. 22, no. 3, pp. 307-338.

Pennacchi, George G. and Christopher M Lewis (1994), "The Value of Pension Benefit Guaranty Corporation Insurance." *Journal of Money, Credit and Banking* , vol. 26, no. 3, pp. 735-53.