The Impact of Risk-Based Pricing in the Student Loan Market:
Evidence from Borrower Repayment Decisions

Natalie Cox [*]

August 31, 2017

Abstract

Advances in data-driven underwriting have both efficiency and equity implications. In the $1.3 trillion student loan market, private lenders offer a growing distribution of risk-based interest rates, while the federal loan program sets a uniform price. I measure changes in consumer surplus that occur as low-risk types refinance out of the government pool into the private market. I use a dataset from an online refinancer to estimate a structural model that relates borrowers' repayment choices to interest rates. I estimate refinancing increases low-risk surplus by $1,302, and show substantial distortionary costs (32% of the average transfer) under a pooled, uniform interest rate. To maintain access to the current uniform rate, the government must subsidize high-risk borrowers $1,507 on average.

1 Introduction:

Risk-based pricing has increased in use and sophistication, especially as more credit and insurance firms transition to an online setting. It has been shown in theory and practice that the ability to identify and accurately price consumer risk can generate large efficiency gains in markets with selection.[1] However, if advances in risk-based pricing create clear winners and losers, this will also change the government's role in ensuring equity and redistribution of surplus.

This is especially true in the student loan refinancing market, where private online firms use rich financial and educational data to underwrite student borrowers who have finished schooling. These individualized prices contrast with those of the Federal Direct Loan program, which offers borrowers a uniform interest rate despite observable variation in the expected costs of lending. Moving from average to marginal cost pricing could correct allocative inefficiencies - low risk types might extend or increase borrowing when faced with an undistorted price in the private sector. But pricing innovations could also have complex implications for how private and public lending options coexist. As low risk types refinance into the private sector, the average risk of the remaining federal borrowers will rise, forcing the government to either raise its uniform rate or subsidize the remaining pool.

This paper studies this efficiency-equity tradeoff empirically, using a dataset of applicants from an online student loan refinancing firm that employs comprehensive risk-based pricing. Using a series of firm-conducted price changes, I show that observationally similar borrowers are interest rate sensitive: they increase monthly payments and shorten maturity when interest rates increase. This suggests that there will be a distortionary "cost" to the governments' one-size-fits-all pricing policy as it transfers from low to high risk borrowers. I combine this elasticity with the observed distribution of risk based prices to quantify the deadweight loss and

---

[1](Einav, Jenkins, Levin, 2012, 2013; Einav, Finkelstein, Levin, 2010; Edelberg 2006; Paravisini, Schoar 2013, and Phillipon 2016 show that credit-scoring can generate both efficiencies in consumer lending, and impact market structure . There is also a literature that studies uniform and average cost pricing schemes in the presence of heterogeneous risk (Bundorf, Levin, & Mahoney 2012; Einav, Finkelstein, & Cullen 2008; Hurst, Keys, Seru, & Vavra 2015) which has shown in several markets (health insurance, mortgages) that while uniform pricing policies achieve cross-sectional redistribution, they can also distort consumer choices and generate welfare loss.

redistribution that occurs under a uniform interest rate, as well as the budgetary impact of low risk borrowers refinancing into the private sector.

The dataset I use provides precise, borrower-level information on risk-scores, interest rates, repayment decisions, demographics and household balance sheets. It also contains 10 interest rate changes that were conducted at a firm-wide level to gather quasi-experimental evidence on maturity and application elasticities. I leverage this variation in price schedules to measure how interest rates impact the loan maturity decisions of observationally similar borrowers.[2] Borrowers reveal their intertemporal preferences when making a maturity decision: by extending maturity they reduce their monthly payment, but increase their total interest paid.

Reduced form evidence reveals that the borrowers in my sample decrease maturity when interest rates increase, and that this sensitivity increases with borrower quality. This suggests that under a uniform interest rate, low risk borrowers will make inefficiently high monthly payments, and high risk borrowers will make inefficiently low monthly payments. I next write down a model where borrowers choose a maturity to maximize expected utility; this model maps the reduced form maturity elasticity to the underlying parameter of interest, the intertemporal elasticity of substitution (IES), and allows for a rich counterfactual analysis. I use the first order condition that captures the maturity tradeoff to estimate the model using non-linear least squares and find a moderately high average IES of .6.

Using the estimated repayment model, I compare the size and distribution of borrower surplus under several pricing regimes: full pooling under a break-even uniform rate, pooling with a refinancing option using FICO-based pricing, and pooling with a comprehensive risk-based refinancing option. The first average-cost interest rate policy redistributes roughly $1,200 from low to high risk borrowers, but generates an average distortionary cost of $450, or 32% of the average transfer. The policy achieves more modest redistribution over income quantiles given that borrower risk type is an imperfect proxy for borrower income. When low-risk types have the option to refinance into the private sector they gain on average $1,500 in surplus

---

[2]Loan maturity is perhaps the most fundamental decision made by the borrower during repayment, given that the choice of how much to borrow is already fixed. It is the only means a borrower has during repayment to lower monthly payments and increase immediate liquidity.

– they face a lower absolute level of interest rates, and respond by making smaller monthly payments. Both the distribution of risk based interest rates *and* the gains to low risk borrowers are much larger when firms price on borrower characteristics, like savings, income, and education, in addition to FICO score. This suggests that non-traditional scoring methods can benefit individuals who are low risk, but have underdeveloped credit histories (i.e. the student borrower population).

These findings highlight how developments in the private sector's ability to price borrower risk will simultaneously i) improve welfare for low risk borrowers and ii) increase sorting of low risk borrowers out the public repayment pool. I analyze one government policy response that would prevent unraveling and maintain equity in the federal pool: providing a net interest rate subsidy. My model highlights how the effective costs of an interest rate subsidy can deviate from the mechanical costs once refinancing and maturity responses are accounted for. For example, lowering the uniform interest rate will reduce refinancing into the private sector, improve the average risk of the remaining federal borrowers, and therefore decrease the subsidy's average cost.

This paper relates to several literatures: first, it contributes research on how borrowers finance their higher education with student loans. While work has primarily focused on optimal borrowing limits and repayment structures (Lochner and Monge-Naranjo 2015; Rothstein & Rouse 2011; Avery & Turner 2012; Yannelis 2015; Lucas & Moore 2010, Beyer, Hastings, Neilson, & Zimmerman (2015)), I consider how interest rates can be used as a policy instrument. The question of how borrowing and consumption decisions respond to interest rates and are mediated by credit constraints is central to the household finance literature (Adams, Einav & Levin; Gross & Souleles 2002; Martins & Villanueva 2006; Karlan & Zinman 2005; Summers 1981). Several papers look specifically at maturity choices (Attanasio, Goldberg, & Kyriazidou 2008; Hertzberg, Liberman, & Paravisini 2016), and also find that long maturities are preferred by riskier, liquidity constrained borrowers. Finally, my paper contributes to a literature that structurally estimates parameters relating to risk aversion and consumption smoothing using micro-data on consumer choices and quasi-experimental variation in prices (Einav & Cohen 2007; Handel

2013; Best et. al. 2015; Gruber 2008).

The rest of the paper proceeds as follows. Section II describes the setting and data with an emphasis on the variables impacting borrowers' maturity choices, and the use of exogenous interest rate variation to identify maturity elasticities. Section III describes the theoretical framework, estimation of the maturity demand model, and discusses the results. Section IV outlines a welfare framework, a simple model of loan costs, and analyzes several counterfactuals. Section V concludes.

2 Setting and Data

2.1 Institutional Background

2.1.1 Student Loan Origination

While this paper focuses on the *repayment* of student loans, it is necessary to understand their origin. Loans that come from the federal government are by far the most popular option to finance post-secondary education – over 90 % of the student loan market consist of Federal Direct Loans. This paper concentrates on federally originated loans that may subsequently be refinanced in the private sector.

There are two key loan-related facts that motivate my paper: growth in the originated volume of Federal loans and heterogeneity in borrower risk. Origination rates in the Direct Loan program, where there is no risk-based underwriting and generous lending limits, have skyrocketed over the past decade. The outstanding volume of student debt has quadrupled in the last 12 years, and the median borrower's holding has grown from $14,000 to $19,500 (Looney, Yannelis 2015), leavingmany borrowers with sizable monthly payments and large amounts of accumulated interest.

Growing in line with origination rates have been average delinquency rates – the average three-year cohort default rate (CDR)[3], peaked at 14.7% for the 2010 cohort, compared to a rate of 5.2% in 2002. However, these average trends mask important heterogeneity: default rates amongst graduate students and individuals at 4 year institutions have remained consistently low. These relatively low-risk borrowers

---

[3]The percentage of loans in delinquency 3 years after entering repayment

4

make up the majority of the dollars lent by the Direct loan program – graduate students are some of the biggest borrowers, holding 33% of dollars outstanding. This means a large portion of the Direct Loan Portfolio will be "overpriced" by a break-even uniform interest rate.

2.1.2   Student Loan Repayment Options:

Repayment of federal debt, and importantly choice of repayment plan, does not occur until a student has finished schooling (either undergraduate or graduate school).[4] Students also have the option of changing repayment plans as time progresses. Federal repayment plans fall into two general categories: fixed payment plans, which adjust the monthly payment level to ensure that the full amount of the original loan will be paid off in a specified number of years, and income-based plans, which scale the monthly payment in proportion to the borrower's income[5]. A CFPB analysis (Gibbs 2017) of the Consumer Credit Panel found 40-50% of borrowers had fully repaid their loans within 5 years, while 25-30% of borrowers took longer than 10 years.[6]

A new repayment option that is growing in popularity consists of refinancing Federal debt in the private sector. Refinancing can take place at the beginning or in the midst of a repayment schedule. Federal loans, which do not carry a prepayment penalty, are paid off by the private firm which takes over the servicing and liabilities associated with the loan. Student loans that are refinanced in the private market are still not dischargeable in the case of bankruptcy. The majority of refinancing firms are online lenders, who digitally link to applicants' financial accounts and credit reports, and use extensive amounts of data to assess their risk. Online applications reduce the frictions of refinancing both for the lender, who face

---

[4]In most cases a student can postpone repaying their debt for more than a 6-month grace period after graduation without incurring interest.

[5]A description of the federal repayment plan options currently available is provided in the Appendix.

[6]During my analysis I assume that term choice in the Direct Loan program is "flexible" – this somewhat understates the gains borrowers receive in the private market where there is flexible term choice, and is a generous assumption since over 50% of borrowers in the Direct Loan Portfolio remain in the 10 year fixed maturity plan.

lower underwriting costs, and for the borrower, who can apply and review competing interest rate quotes in a matter of minutes.

2.1.3 Interest Rates and Loan Costs:

When choosing a loan maturity, borrowers trade off between the payment made each month towards principal and interest, and the total interest cost paid over the life of that loan. A longer maturity loan will have a lower monthly payment, but a higher overall interest cost (as more interest accumulates, at potentially a higher rate). This means that different loan maturities will appeal to different types of borrowers. For example, individuals with lower incomes who are more liquidity constrained may prefer a long maturity with lower monthly payments.

The federal government charges a single interest rate for all loan terms, whereas the private sector charges a higher rate for longer maturities. Holding maturity constant, the private sector can also offer either a lower or higher interest rate *depending* on the borrower's expected risk. Low risk borrowers can therefore decrease both the monthly and total cost of their loan by moving to the private sector and refinancing at a lower interest rate. For higher risk individuals who would actually face a higher interest rate under risk-based pricing, the uniform federal interest rate is preferable and refinancing will likely not occur.

In the private refinancing sector, the variables that determine an individual's risk-based interest rate are proprietary and company-specific, but are primarily driven by free cash flow, degree, income, savings, FICO, and debt.[7]. Federal interest rates follow a formula specified under the Higher Education Act. Each year they are determined by an index rate plus an add-on margin that varies by loan type (see Table 1). Graduate students face higher federal interest rates than undergraduates; the fact that they are some of the largest, lowest risk borrowers makes them a population especially prone to refinancing in the private sector.

---

[7]Legally prohibited risk-based pricing factors under the Equal Credit Opportunity Act are: race, color, religion, national origin, sex, marital status, age, and receipt of income from any public assistance program

Table 1: 2011-2015 Interest Rates on Federal Direct Student Loans

Loan Type	Borrower Type	Index	Add-on	Fixed Interest Rate
Direct Un/Subsidized Loans	Undergraduate	10 Yr Tr	+ 2.05%	3.4-4.66%
Direct Unsubsidized Loans	Graduate/Professional	10 Yr Tr	+ 3.60%	5.41-6.8%
Direct PLUS Loans	Parents& Graduate	10 Yr Tr	+ 4.60%	6.4-7.9%

Undergraduates are able to borrow at lower interest rates up to a certain limit ($5,500 to $7,500, depending on their year in school), and then must borrow at higher interest rates.

2.2   Dataset:

I use a proprietary dataset from a student loan refinancing firm that contains extensive information on interest rates, risk score inputs and outputs, and maturity and refinancing decisions. The dataset describes individuals who both decide to refinance in the private sector, *and* on the wider population of those who apply and view an initial interest rate quote. While the first group is very low risk, the second sample allows me to measure the distribution of market-priced risk, and the refinancing propensity, of a more representative sample of Federal borrowers.

The main dataset that I use when measuring maturity elasticities is a repeated cross-section of all new borrowers refinancing with the firm over the period of a year; it links background financial information (debt amount, income, assets, credit score) about borrowers with the menu of interest rates they faced and the ultimate maturity choices they made when refinancing with the firm. When estimating refinancing propensity, I use a similarly structured dataset which includes all applicants to the firm, including those who do not necessarily choose not to refinance. The sections below describe the data and price variation used during estimation.

2.2.1   Descriptive Statistics

Refinancer Population: The population of borrowers who ultimately refinance are high income, high debt, and highly educated. The median borrower holds over $50,000 in student loans, and the median monthly payment on refinanced debt is $600 per month. The median borrower also holds $38,000 in assets, $0 in investments (the 75th percentile has $15,000 in investments), owes $89,000 in additional liabilities, and has a median monthly free cash flow (post tax income minus fixed monthly payments like housing) of $3,100. Borrowers hold a host of degrees and

occupations; JDs (lawyers) make up 13% of the sample, MBAs are 17%, MDs (doctors) 5%, pharmacists 6%, and dentists 4%. The majority of borrowers finished school in the last 4 years, with 25% graduating in 2016 and 50% since 2012; however, some are refinancing older loans, with 25% of borrowers having graduated before 2010.

Table 2: Borrower and Loan Summary Statistics

Borrower Summary Statistics				Loan Summary Statistics			
	Mean	Median	SD		Mean	Median	SD
Income	75,879	68,304	39,799	Maturity (Months)	106.8	93.0	49.7
Loan Amt	67,078	50,656	52,890	Monthly Payment	799.9	600.0	624.8
FICO	782	787	36	APR	4.494%	4.64%	1.042%
Mortgages	0.40	0.00	0.60	Variable Rate	0.40	0.00	0.50
FCF	3,636	3,100	2,574				
Graduate	0.70	1.00	0.50				
Age	32.60	31.00	6.80				
Dependents	0.50	0.00	0.90				
N	11663			N	11663		

This table contains summary statistics describing the population of borrowers who ultimately refinance their loans. The upper rows describe the income and characteristics of the borrowers themselves. Income refers to yearly income. FCF refers to the monthly post tax income minus all fixed expenses. Graduate refers to the portion of the population who has a graduate degree. The lower rows describe the terms of the refinanced loans. Variable rate is the proportion of loans that have a variable interest rate.

The impressive background of these candidates translates into them obtaining considerably lower interest rates when refinancing. The average previous federal interest rate on the loans (before refinancing) was 6.7%, and the average refinancer saved 2.21 percentage points when refinancing.

Applicant Population: In addition to refinancers, I also observe a larger, less selected sample of website visitors who see an interest rate quote but do not necessarily proceed with the refinancing process. This sample is more representative of the population of graduate student borrowers who have federal loans. In the Appendix (Figure 11), I compare the debt and income quantiles of my applicant sample to a nationally representative sample of graduate student borrowers – they look very similar. The average risk-based APR for this sample is 6.5%, which is very close to the uniform graduate rate charged by the Direct loan program – this again suggests that the applicant sample is representative of the distribution of risk underlying the federal portfolio.

Maturity Choices: Borrowers in my setting are asked to choose from a continuum of maturities from 5 to 20 years which allows them to customize their monthly payment. As borrowers extend their repayment maturity, they reduce their monthly

payment but increase their total costs. There is a steep and sizable yearly payment gradient for individuals who are starting at a low maturity: individuals at a 5 year maturity pay $1300 less per year on average when they increase maturity to 6 years. However, individuals moving from a 19 to 20 year maturity pay only $200 less per year.

Given the novelty of the choice set and complexity of interest rate/monthly payment tradeoff, one might wonder if borrowers are making completely informed decisions. There are several aspects of the user interface that borrowers interact with that make this unlikely: for one, borrowers are provided with the monthly payment, APR, and total paid for the maturity the choose at many points during the refinancing process. Borrowers use a "slider" to adjust their monthly payment, and are shown how maturity, APR, and total payments change simultaneously – Figure 1 shows an example of this user interface. This means they are aware not only of the tradeoffs inherent when choosing any given maturity, but also the rate at which these tradeoffs change when they adjust maturity.[8]

This being said, the distribution of maturity choices suggests that some borrowers may be using heuristics, like rule-of-thumb accounting, when making maturity choices. For example, the distribution of chosen monthly payments has distinct spikes at "round" monthly payments, like $500, $1000, or $1500. While these behavioral borrowers are in the minority, it is important to acknowledge that a rational model of financial decision making does not apply to all households.[9]

2.2.2 Price Variation:

The maturity choices described above show how borrowers select into various repayment contracts *given* a set of interest rates. However, they reveal little about how borrowers' repayment decisions would change if faced with *new* set of interest rates. Measuring this elasticity, a central goal of this paper, requires variation in

---

[8]In this way the information provided to the borrower is very similar to the information necessary for the first order condition calculations in our model: they see both the change in monthly payment associated with a maturity increase and the change in total interest paid.

[9]Those with specific maturity or monthly payment targets will be less sensitive to interest rate changes, and this would bias elasticity estimates downward.
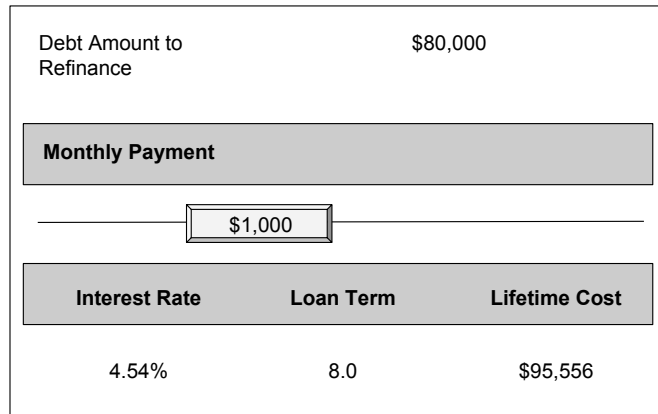
Figure 1: User Interface for Maturity/Monthly Payment Selection
This figure shows the interface individuals use when selecting a loan term. The interface shows them the customized monthly payment, APR, and total paid associated with every possible term choice. This means even less financially-savvy borrowers are well-informed of the implications of their term choices.

interest rates.

There are two main types of interest rate variation in my dataset: risk-based and within-risk. Using risk-based variation to identify the elasticity of maturity with respect to interest rates is potentially misleading. Individuals with different risk scores may also differ on unobservable dimensions (like expectations about future income growth or volatility) that will impact their maturity choices. Ideally one would instead use price variation that is orthogonal to all borrower characteristics, including risk score. I refer to this exogenous variation as "within-risk" variation.[10]

I use 10 small within-risk score price changes that were conducted at a firm-wide level, and were unrelated to the characteristics of any given borrower. The exogenous price changes were conducted primarily to gather quasi-experimental evidence for the firm on maturity choice and application volume elasticities with respect to interest rates. The price changes occurred over time, not simultaneously for different groups of borrowers, and at a frequency of once to twice a month.[11] While not all price regimes lasted the same amount of time or effected the same number of borrowers, on average they each impacted 1,100 borrowers. Borrowers were not aware of the timing of these price changes, and therefore could not respond

---

[10]Figure 12 in the Appendix provides a graphical explanation of why within-risk score price variation is necessary for identification, and how it can be used.

[11]This frequency helps alleviate concerns about significant changes in the composition of customers over time (a period of rapid growth), but there are still changes in the observable characteristics of the population over the full set of price changes, addressed below.

by adjusting when they refinanced.

2.3   Reduced Form Maturity Elasticity wrt. Interest Rates

To estimate a reduced form elasticity, I run a series of regressions of maturity choice on observable characteristics and offered interest rates.[12] To quantify the interest rate an individual is offered, I calculate the average fixed rate APR (P_i) over all maturities that individual i with risk type p_i faces. I first regress maturity choice in months (T_i) on APR (P_i), and observables(X_i), pooling both sources of price variation[13]:

$$T_i = \alpha + \beta P_{it} + X_i'\mu + \epsilon_{it}$$

The results (see column 1 of Table 9) show individuals who face higher interest rates as measured by P_i are actually *more* likely to choose a longer term, even conditional on income and loan amount. This seems counterintuitive, since it implies that as borrowing becomes more expensive, individuals want to borrow for a longer amount of time. However, if higher risk types have a higher demand for maturity (due perhaps to unobserved liquidity constraints or income variability), then our price coefficient will suffer from omitted variable bias.

The next specification therefore includes risk score, controlling for risk-based price differences and using *only* the remaining within-risk price variation to identify β. The coefficient on the price variable now has a significant, negative sign – this means that when faced with higher interest rates, similar individuals choose shorter loans. The elasticity that corresponds with this coefficient (see Table 3) says that a one percent increase in the average offered APR causes a .8% decrease in the average maturity chosen. In unit terms, this means that increasing the average APR from 5.5% to 5.6% would decrease the average term by 1.7 months (from a mean of 108 months). For the average $70,000 loan, this would increase monthly payments by 1.6%, but keep total interest payments relatively constant, increasing

---

[12]The results of the regression are in the Appendix in Table 9

[13]I use a tobit specification to account for the truncation of the choice set at 60 and 240 months.

Table 3: Maturity Elasticities

	(1)	(2)
Overall	-0.819***	
	(0.307)	
Highest Risk		-0.226
		(0.352)
Mid Risk		-0.708**
		(0.313)
Lowest Risk		-1.524**
		(0.626)
N	11663	11663

Standard errors in parentheses

* $p < 0.10$, ** $p < 0.05$, *** $p < 0.01$

This table converts the coefficients in specifications (2) and (3) from Table ?? into elasticities calculated at the mean values of the independent and dependent variables. These values can therefore be interpreted as the percentage change in term in response to a 1% increasing in average APR. The "Overall" elasticity calculates the elasticity for the entire sample, whereas the second column separately calculates the elasticity for the upper, middle, and lowest thirds of the risk distribution.

by only .005%. This shows that the majority of borrowers place more weight on minimizing total interest payments than on minimizing monthly payments.

In a final specification, I allow risk type and price to interact – this allows for different price elasticities across risk types. Table 3 shows that the highest risk individuals are essentially inelastic to price changes, whereas the lowest risk individuals exhibit a much higher elasticity and reduce term when interest rates increase. This is interesting, since it suggests that while the lowest risk individuals are interest rate sensitive, the highest risk individuals are primarily driven by the level of monthly payment. This difference in price sensitivity may at least partially explain the difference in the levels of term choices across risk types.

2.3.1 Sample Selection Over Price Regimes

Using temporal price variation presents a selection concern: while some individuals may respond to interest rate changes on the intensive margin by adjusting maturity, others may respond on the extensive margin by no longer refinancing or refinancing with a different company. If those who join or leave the population after a price change have systematically higher or lower maturity preferences, then this extensive margin response will bias our intensive margin estimates.

One empirical way to gauge the extent of extensive margin responses is to test whether changes in observable borrower characteristics over price regimes are correlated with the exogenous variation in APR. If the composition of observable char-

acteristics is predicted by the price changes, then we would be worried that there may also be selection on unobservables. Table 10 tests whether four important observable characteristics, income, debt, FICO, and savings, are predicted by the price regime shifts. These insignificant results show that price changes did not cause any differential attrition across observable characteristics: while characteristics like income and FICO did vary over price regimes, this variation was not correlated with the price level. Figure 15 graphically shows the lack of correlation between observable characteristics like debt and income and the price shifts.I also predict individuals' maturity choices, $\hat{T}_i$, using all observable characteristics other than APR, and test whether this variable is predicted by the price regime shifts; again, these results are insignificant. This makes sense since the highly competitive nature of the refinancing market, and growth of risk based pricing, means that even "high" risk borrowers have outside options that are close to their quoted price. It is helpful to note that the price changes were not monotonic: interest rates both increased and decreased over time, and therefore will not be confounded by other monotonic trends occurring over time like growth of the company.

3   Loan Repayment Model

In this section I outline a simple model which describes how borrowers make maturity and refinancing choices as they begin to repay their loans. One goal of the borrower model is to be able to interpret the reduced form maturity elasticity in terms of the intertemporal elasticity of substitution.

3.1   Basic Set-up

Borrowers entering repayment have incurred a fixed amount of student debt while attending school, D_i, are now finished with school, and are beginning repayment.[14]

---

[14]I model all borrower decisions conditional on debt, schooling, and educational choices, which are made at an earlier period before repayment begins. This equates to the assumption that these decisions are fixed and not impacted by the level of interest rates or ability to refinance debt. This assumption is valid for the population of student borrowers who have already made their loan principal decisions and are yet to make repayment choices (i.e. those currently in school or beginning repayment) – however, it does not apply to individuals who have yet to make borrowing

In this two-stage model, borrowers choose a repayment maturity, T_i, to maximize their present discounted stream of expected future utility, and whether to refinance their federal loan into the private sector. Two main things distinguish the public and private repayment options: risk-based pricing and maturity-based pricing. The private sector offers interest rates $(r(T_i, p_i))$ that are increasing in a borrower's observed risk p_i and chosen maturity T_i. The government offers a single price for all risk types and maturities, g. Monthly and total payments could be lower or higher for a given individual in the private vs. public sector - this depends on their risk type and maturity preference.

All borrowers have the same per-period CRRA utility function $u(c) = \frac{c^{(1-\gamma)}}{(1-\gamma)}$, and discount factor, β. While borrowers can control their yearly payment level, they cannot control their variable, growing income stream, which I parametrize using a unit root process:

$$ln(w_{it}) = ln(w_{it-1}) + u_{it}$$

$$u_{it} \sim N(g_i, \sigma_i^2)$$

where g_i is an individual-specific yearly growth rate and σ_i^2 is individual-specific income variance.

3.2   Step 1: Maturity Demand

Borrowers of risk type p and debt amount D choose a maturity T to solve:[15]

$$\max_T E[\sum_1^T \beta^t u(w_{it} - d_i) + \sum_{T+1}^Q \beta^t u(w_{it})]$$

where Q is the individuals' maximum age (i.e. finite), w_{it} is post-tax income, and d_i is the yearly payment associated with maturity T. Here consumption is defined as post-tax income minus the student debt payment $d(T)$ – this implies that

decisions (prospective borrowers who have yet to start school). In a full equilibrium analysis, the level of interest rates and refinancing options could also impact choices like loan principal.

[15]In the public sector the interest rate $r(T, p_i)$ is replaced with g

individuals are "hand-to-mouth" and not smoothing consumption through other debt or savings.[16] It also means that changes in maturity translate directly into changes in consumption by impacting the size of the yearly payment, $d(T)$.

The debt payment $d(T)$, that individuals make each period is a function of their total debt amount D, their chosen maturity, T, and the (potentially risk, maturity-specific) interest rate schedule that they are offered, $r(T, p)$:

$$d(T) = T * D * \frac{r(T, p)}{(1 - (1 + r(T, p))^{-T})}$$

As borrowers extend maturity, each periods' payments become lower ($\frac{d(T)}{(dT)} < 0$), but they pay more over the life of the loan.

Solving the maximization problem results in the first order condition:

$$0 = -E[\sum_0^T \beta^t \frac{\partial d}{\partial T} u'(w_{it} - d_i) + \beta^{T+1} u(w_{i(T+1)} - \frac{d_i}{2}) - \beta^{T+1} u(w_{i(T+1)})]$$

which can be rewritten as:

$$E[\sum_0^T \beta^t \frac{\partial d}{\partial T} u'(w_{it} - d_i)] = E[\beta^{T+1}(-d_i) u'(w_{i(T+1)} - \frac{d_i}{2})]$$

This condition says that at the optimal loan maturity, the sum of marginal utility gained from a slightly lower monthly payment (from a slightly longer term) is equal to the marginal utility lost from paying additional interest for an extra year.[17]

3.2.1 The Influence of Interest Rates on Maturity Choice:

The first order condition captures how maturity choices, and therefore utility levels, change under various price regimes. All else constant, as the level of interest rates increases, individuals must decrease maturity to maintain the optimality

---

[16]I address this assumption empirically in the robustness section.

[17]To make this condition empirically tractable, I approximate the second term, $\beta^{T+1} u(w_{i(T+1)} - d_{jit}) - \beta^{T+1} u(w_{i(T+1)})$, with the expression:

$$\beta^{T+1} u(w_{i(T+1)} - d_i) - \beta^{T+1} u(w_{i(T+1)}) \approx \beta^{T+1}(-d_i) u'(w_{i(T+1)} - \frac{d_i}{2})$$

condition. The exact formula for the response $\frac{dT}{dr}$ does not have an analytical solution, but Figure 16 in the Appendix shows how optimal term choices, simulated according to this model, vary with the level of interest rates – as interest rates increase, the optimal maturity choice decreases. It also plots this relationship for two calibrated values of the intertemporal elasticity of substitution (IES), 1.25 and .3. As the IES increases, individuals become more interest rate sensitive. This leads to a lower optimal maturity choice at any given interest rate, and also to a steeper relationship between T^* and r.

3.2.2 The Influence of Non-Interest Rate Factors on Maturity Choice:

The other non-interest rate factors in our model that influence maturity demand and interact with the interest rate elasticity are: income level, debt level, income growth and volatility, and the intertemporal elasticity of substitution $\frac{1}{\gamma}$. Due to concave utility, individuals who are low income or high debt gain more marginal utility from decreasing yearly payments, and thus have a higher willingness to pay for long maturities. Individuals who expect income to grow in the future will also prefer a long maturity, since it acts as a means to transfer consumption from the future to the present. Individuals with higher income variability have both higher and less elastic maturity demand due to the fact that longer loans help to smooth consumption across a more variable income profile.

Note that the income-related factors that drive demand (income levels, growth, and volatility) are potentially correlated with risk score p_i. Therefore, the same variables that increase demand for maturity on the borrower's side will also increase interest rates on the supply side. This means that even when faced with higher risk-based prices, high risk borrowers may choose longer loans. This is in line with our reduced form evidence, which showed that, all else constant, riskier borrowers had higher demand for long maturities.

The first order condition also helps us understand how the intertemporal elasticity of substitution ($\frac{1}{\gamma}$) influences demand. As γ increases, an individual will prefer a longer maturity holding all else constant (see figure 16 for a comparison of maturity demand with a high and low IES). Intuitively, this is because an individual with

concave utility will prefer to smooth consumption by lowering yearly payments, even if it means paying more interest overall. A high level of γ (ie a low intertemporal elasticity of substitution) also means that the term choices of individuals will be less responsive to price changes. Thus γ is essential for understanding how a borrower's decisions, and utility, would respond to *changes* in price, the central goal of this study.

3.3 Step 2: Refinancing Choice

Borrowers also decide whether or not to refinance by comparing the level of utility at the optimal maturity across sectors. For individuals on the high or low end of the risk distribution, the risk based price differential determines whether they should refinance. For very low risk individuals there will be a clear incentive to refinance in the private sector, and for very high risk individuals, there will be no incentive to refinance because government interest rates will be significantly lower than private rates. For marginal individuals who face a similar level of prices in the private and public sectors, the term based price differential (and whether they prefer shorter or longer loans) could also determine whether they will sort into the private sector. While the model describes the decision to refinance as a discrete choice problem, in reality borrowers might face frictions (inertia, search costs), value non-pecuniary repayment benefits, or have idiosyncratic preferences that prevent them from refinancing even when they would receive lower interest rates.

3.4 Modeling Borrower Delinquency

This model ignores the impact of delinquency or default on maturity choice, focusing only on the intertemporal consumption tradeoff that comes from a higher or lower monthly payment.[18] This equates to the assumption that adjusting loan maturity does not impact the probability of delinquency, despite the fact that maturity impacts the size and duration of monthly payments.

---

[18]Currently, borrower income levels and risk impact maturity decisions due to the fact that a low income draw minus a large debt payment will generate a very high marginal utility.

One reason for this assumption is that the dataset used to estimate the model does not observe instances of default or delinquency.[19] The main model focuses solely on relationships that can estimated empirically – i.e. the response of maturity to changes in interest rates. While the dataset captures how maturity choices relate to interest rate variation, absent data on delinquency or default it is impossible to *empirically* link maturity choices to repayment outcomes. In the Appendix, I instead write down and analyze a model that *theoretically* links delinquency and maturity choice. In this model, extending maturity will decrease the probability of delinquency, $(P(T)$, in any single period by lowering the monthly payment, but may also increase the probability of delinquency over the life of the loan by extending the repayment period. Misspecification will be more or less of a problem if changes in maturity have a large impact on utility via these delinquency channels. This exercise also provides a robustness test, for various calibrated values, of the central, simplified model which I use in the estimation section. The results suggests that the estimates are not very sensitive to the exclusion of the delinquency channel – for this low risk group, the first order effect of a change in maturity seems to operate instead through the monthly payment consumption smoothing channel.

4   Estimation of Borrower Model

In this section I use the same maturity choices and within-risk price variation explored in the reduced form section to estimate the structural borrower model. This structural exercise will allow us to estimate borrower utility using the same price variation and maturity response as in the reduced form exercise. While it imposes stronger assumptions on the borrower's problem (i.e. calibrating the income process and assuming CRRA utility), it allows us to map the reduced form maturity elasticity to a parameter of economic interest, γ, and to ultimately measure changes in consumer surplus.

In this section I first discuss estimation of the first order condition using nonlinear

---

[19]Note that the dataset captures only the first 6-36 months of repayment of potentially very lengthy repayment maturities (5-20 years). The absence of delinquency in these initial months does not necessarily mean that delinquency will never occur.

least squares, then results, identification, and robustness.

4.1 Estimation of FOC:

Recall that individuals choose T to maximize a discounted stream of yearly utility; the resulting first order condition provides our main estimating moment. Most elements of this equation are observed, including: T_i, the optimal term choice, d_i which represents the yearly payment for individual i at term T_i, $r(T_i, p_i)$ which is the risk, term specific interest rate faced by individual i at term T_i, $\frac{dd}{dT_i}$, and w_{i0}, which is defined as after tax income minus fixed expenses.

Future income, w_{it}, is not observed, but I assume log income follows a unit root process and calibrate both the growth rate and volatility[20] The calibration details are provided in the appendix – I predict $\hat{g}_i$ using a separate cross-sectional dataset of personal loan applicants, and calibrate income volatility, $\hat{\sigma}_i^2$, using the repayment probabilities implied by a simple lending model.

The remaining parameter left to estimate is the intertemporal elasticity of substitution, which I model as a function of observable characteristics (X_i) including degree type, risk score, current disposable income, student loan amount, age, age^2, FICO score, home ownership, and number of dependents:

$$\gamma_i = \gamma + X_i'\mu$$

Allowing γ_i to vary with X_i will control for changes in observable characteristics across price regimes. Following the logic outlined in the reduced form section, it is also important to include observed risk type (p_i) directly in the model. If p_i was not included in the estimation, the model would wrongly attribute differences in maturity choices across risk type to differences in offered APR, and our estimate of γ would suffer from omitted variable bias.

In my main specification, I use a certainty equivalence approach to write the first order condition as a closed form analytical expression. Specifically I rewrite the expected marginal utility as the marginal utility of a certainty equivalent given

---

[20]In the robustness checks I relax this assumption and try other specifications.

by:

$$E[(w_{i0} * e^{t*(\hat{g}_i)} * e^{\sum_1^t u_{it}} - d_i)^{-\gamma_i}] = (w_{i0} * e^{t*(\hat{g}_i)} * e^{\pi_{it}} - d_i)^{-\gamma_i}$$

where π_{it} is the certain amount an individual would have to be given in that period to make the certainty equivalent equal to the expected marginal utility. Specifically[21]:

$$\pi_{it} = \frac{1}{2} * t * \hat{\sigma}_i^2 [1 - (1 + \gamma_i)\frac{w_{i0} * e^{t*(\hat{g}_i)}}{w_{i0} * e^{t*(\hat{g}_i)} - d_i}] \qquad \text{for t<T+1}$$

$$\pi_{it} = \frac{1}{2} * t * \hat{\sigma}_i^2(-\gamma_i) \qquad \text{for t$\geq$T+1}$$

One can see that as income volatility, risk aversion, and the debt to income ratio $(w_{i0} * e^{t*(\hat{g})} - d_i)$ ratio increases, the certainty equivalent becomes more negative.

Using this expression, our analytical estimating moment becomes:

$$h_i(\theta) = \sum_1^T \beta^t \frac{\partial d}{\partial T}(w_{i0} * e^{t*(\hat{g}_i)} * e^{\pi_{it}} - d_i)^{-\gamma} - \beta^{T+1}(-d_i)(w_{i0} * e^{(T+1)*(\hat{g}_i)} * e^{\pi_{i(T+1)}})^{-\gamma_i}$$

This makes the first order condition a nonlinear function of observable variables, $(r(T_i, p_i), T_i, w_{i0}, D_i, X_i, \hat{g}_i, \hat{\sigma}_i^2)$, and unobservable parameters, $\theta = \{\gamma, \mu\}$, that we need to estimate[22]. To estimate the model, I use nonlinear least squares, choosing the parameters that minimize the quadratic form:

$$b = \arg\min_\theta h(\theta)'h(\theta)$$

4.1.1 Empirical Identification:

The identification intuition for our structural model remains very similar to that of the reduced form section: there, maturity choice was expressed as a linear

---

[21]for derivation of π_{it} see Appendix

[22]I calibrate $\beta = .98$

function of observables, risk type, and interest rate (X_i, p_i, and $r(T,p)_{it}$), with the interest rate coefficient (β) identified using the series of price changes that were orthogonal to p_i. In the structural estimation, T_i is instead expressed as an implicit non-linear function of X_i, p_i, and $r(T,p)_i$ derived from the borrower's first order condition, and the price changes now serve to identify γ.

In both models, the price changes provide moment conditions in which observationally identical individuals (in terms of both X_i and p_i) face a different set of interest rates (r vs r') and make potentially different maturity decisions. The parameters, β and γ, must rationalize how *changes* in maturity decisions respond to *changes* in interest rates. They are therefore estimated using *shifts* in the maturity distribution over price regimes, and not the maturity distribution itself. Using this price variation requires the assumption that conditional on X_i and p_i, any unobserved characteristics of borrowers across price regimes are uncorrelated with their maturity choices.

One key difference between the two specifications is that income growth and volatility are directly modeled in the structural section. In the reduced form analysis, income growth and volatility are one of many sources of unobservable heterogeneity that might impact maturity choice – they are left in the error term, and assumed to be orthogonal to the price changes. In the structural analysis, income growth and volatility become a source of observable heterogeneity. By modeling the income path over time and over states of nature, we can translate changes in maturity into changes in consumption. While this puts more structure on the borrowers problem, it also allows us to interpret the reduced form maturity elasticity as a more economically relevant parameter, the IES.

In the model, γ plays the role of both the IES and the risk aversion parameter. Empirically, the parameter is identified from a consumption smoothing decision – how changes in monthly payment respond to changes in interest rates – and thus should be interpreted as the IES. In the counterfactual, changes in welfare also come from intertemporal changes in consumption, not changes in income risk.

A final concern is to correctly attribute what changes in term choice across price regimes comes from actual changes in interest rate, and what portion comes from

changes in sample composition. Again, modeling γ as a function of X_i and p_i will control for the impact of *observable* heterogeneity across price regimes on maturity choice.

4.2 Results

The results from the structural estimation are shown in Table 4. The first column estimates come from our preferred specification, which models γ as a function observable characteristics, and the remaining columns report results from specifications with alternative assumptions, discussed in more detail below.

The average estimate of γ in the primary, and in all specifications, falls on the moderate to low end of estimates in the existing literature. It translates into an IES of .55, and most recent micro estimates have found a IES from .2-.6. This value implies that on average there is a sizable consumption response to changes in interest rates. The small estimate of γ is not surprising in light of our sample, setting, and earlier results. It is reflected in both the distribution of term choices, in which one quarter of borrowers choose a maturity below 6 years, and the reduced form results, which found a relatively large maturity elasticity. In addition, my dataset is unique in that I can fully observe household's balance sheets and explicitly control not only for borrower income, but also for other monthly fixed expenses. Using pre-tax income, rather than this more accurate measure of monthly free cash flow, would overstate borrower liquidity and bias estimates of γ upwards.

Individuals actively refinancing their loans are also likely more cognizant of the interest rate tradeoffs they are making then individuals in studies that examine credit card use or saving rates. The online interface also explains how interest rates and debt maturity interact, potentially making my sample more informed than those studied in a traditional loan setting. These are complex calculations that the borrower may not make independently, and make the total interest/monthly payment tradeoff extremely salient.

It is also possible that the *type* of debt studied here could also have a unique psychological impact on estimates of γ. The amount of student debt a borrower has is more often determined by "necessity" (due to the level of tuition or financial

aid available at the individual's school), rather than choice and thus be perceived as more burdensome and unwanted. Therefore borrowers might treat their student loans differently than other forms of debt or savings, and want to pay it off more quickly. This results underscore the importance of considering several models of consumer behavior when analyzing saving and borrowing decisions - while our life-cycle model of repayment rationalizes these choices under the assumption of full information and rational expectations, there may in fact be behavioral tendencies, for example debt aversion or rule of thumb accounting, that are driving some portion of individuals' behavior.

The main specification estimates γ as a function of observables, returning a distribution of predicted values from 1.2 to 2.6. As explained in the identification section, this helps isolate the exogenous interest changes as the main source of identifying variation for the intertemporal elasticity of substitution, but it also allows us to compare the IES across different characteristics. There is a lower elasticity amongst older individuals, those with larger families, and those with lower credit scores. Perhaps more surprising is that individuals with higher income, lower amounts of student debt, and lower risk scores have a lower IES, which seems to contradict the reduced form results. However, heterogeneity in the IES estimates should not necessarily map 1-to-1 to the reduced form findings: the structural model allows for much more complexity than our OLS analysis, as it captures, for example, the curvature of the monthly payment tradeoff at different points in the maturity distribution, the full interest rate schedule, or the role of expected income growth. A high income individual may be making a much smaller consumption tradeoff when they reduce maturity by one month than a low income individual making the same decision. Therefore, even though low risk or high income individuals may have a higher reduced form maturity elasticity, when these additional non-linear factors are considered, they actually have a lower IES.

4.3 Robustness Analyses:

In this section I test the robustness of my modeling assumptions in several ways: first, I test how sensitive estimates of γ are to the inclusion or exclusion of delin-

Table 4: Maturity Choice Model Estimates

Coefficients	(1) Main Specification - $\gamma_i = f(X_i)$	(2) Homogenous γ	(3) No Income Risk	(4) No Income Growth or Risk
γ - Constant	1.704 (0.269)	1.416 (0.01)	5.900 (1.771)	8.369 (3.122)
γ - log(Income)	0.203 (0.023)		2.356 (0.130)	3.269 (0.247)
γ - log(Debt)	-0.234 (0.010)		-2.361 (0.064)	-3.434 (0.147)
γ - Home Owner	-0.0532 (0.0126)		0.1689 (0.0800)	0.1484 (0.1328)
γ - # Dependents	3.76E-02 (8.60E-03)		2.56E-01 (4.83E-02)	3.50E-01 (8.44E-02)
γ - Age	6.10E-03 (1.10E-03)		2.40E-02 (7.15E-03)	3.08E-02 (1.31E-02)
γ - Risk Score	0.124 (0.007)		-0.455 (0.047)	-0.464 (0.081)
γ - FICO	-3.00E-04 (2.00E-04)		-1.98E-03 (1.14E-03)	-2.15E-03 (1.97E-03)
α_{60}	2.477 (0.315)	-0.578 (0.38)	4.219 (0.349)	4.691 (0.424)
α_{240}	-71.695 (1.992)	-72.00 (2.06)	-75.043 (2.149)	-84.091 (2.400)
$\bar{\gamma}$	1.801	1.4161	4.688	5.892
σ_γ	0.242	−	1.673	2.407
N	11585	11585	11585	11585

Standard errors in parentheses. This table presents results from the non-linear least squares estimates of the borrower's maturity choice model. Income is defined as yearly post-tax income. Debt is the amount of student loan debt the individual is refinancing. Home Owner is a dummy indicating whether an individual owns a home. Risk Score is the firm specific score that is used as the basis for risk based prices; a higher score indicates lower risk. $\alpha_{(x)}$ are dummy variables that indicate whether an individual chose the minimum or maximum maturity to account for the truncation of the maturity choice set a 5 and 20 years.

Column (1) presents the main specification, in which the intertemporal elasticity of substitution parameter (γ) is modeled as a function of observables; Column (2) does not allow γ to vary with observables. In both (1) and (2) the future income growth and volatility of the borrowers are calibrated using an external data source (see the appendix for details). In (3) I only calibrate income growth, and make income deterministic (not risky). This increases the estimate of γ, since without income volatility individuals must have a lower IES to rationalize the same maturity choices in the data. In (4) I remove income growth (income stays constant) and income volatility. Again, this increases the estimate of γ.

quency in the borrower's problem. Second, I examine whether borrowers are adjusting on other financial margins, to see if changes in maturity truly represent changes in consumption. Finally, I test the modeling assumption that borrowers are making a single, not recurrent, maturity choice by looking at prepayment rates in the data.

4.3.1 Impact of Delinquency on Estimates

While the estimated model abstracts away from delinquency, in the Appendix I write down a model that includes delinquency and adds additional terms to the first order condition. Here I calculate what happens to $\hat{\gamma}$ as we increase or decrease the magnitude of the additional terms.

The first order condition with delinquency included an unobservable utility "penalty", G, that a borrower would face if they ever missed a payment on their loan.[23] The value of G that would provide the same estimate of γ using either the delinquency FOC or the simplified FOC is $G = -.1269E - 04$[24] – this penalty is equivalent to 1.4 years of utility at the starting income level, or cutting utility to 97% of the starting level for the next 50 years. Figure 2 panel 2 tests how γ changes across a broad range of values of G, from a penalty that would reduce utility by 0% for the next 50 years, all the way to 70%. The value of G implied by our estimate of $\hat{\gamma}$ is relatively low, suggesting our estimate may be biased downwards. But even across a wider domain of penalty values, the implied range of γ is relatively narrow.

I also test how a range of delinquency probabilities, $P(T) = .001$ to $.3$, would impact an estimate of γ. Again, the value of γ does not change considerably, moving from 1.6 to 2. The insensitivity of γ to these additional terms makes some intuitive sense. For borrowers in our dataset who are relatively low risk, the derivative of the probability of delinquency with respect to maturity is small. This means that the first order impact of a change in maturity on utility comes from the change in monthly payment, which we capture in the main specification, not from the change

---

[23]This was one obstacle to estimating the delinquency model: while the lender's problem provides observable delinquency probabilities, there is no clear empirical proxy for G.

[24]For these calculations I use the empirical proxies for $P(T)$ that I derive from the lender model in the appendix. For the average borrower in the dataset, and for the value of $\hat{\gamma}$ estimated above using the simple FOC, I calculate the delinquency first order condition. I then find the value of G that sets this condition equal to zero.
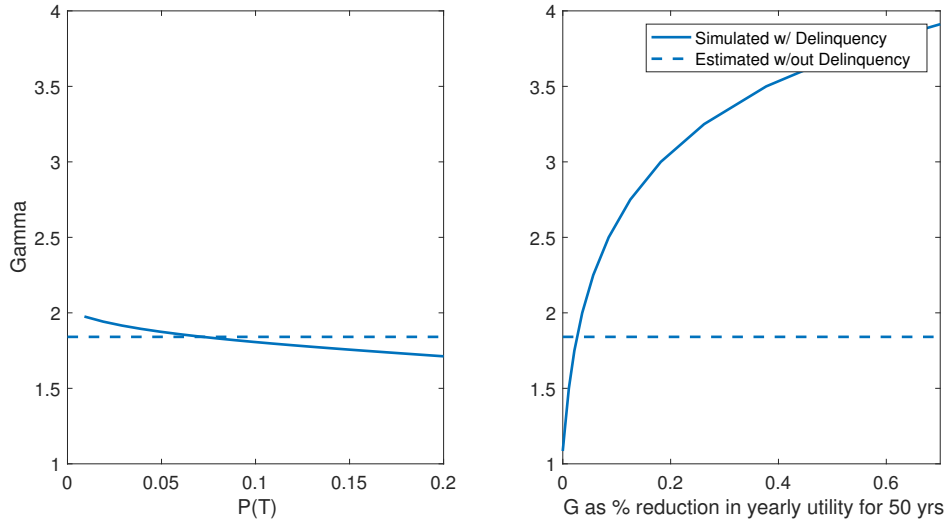
Figure 2: Comparison of γ under delinquency and non-delinquency model.

in the probability of delinquency.

4.3.2   Contemporaneous Financial Decisions

The above model defines yearly consumption as post-tax income minus the student debt payment; in reality individuals may be saving or paying down other debt, rather than consuming, this residual.

One way I address this empirically is by using a measure of borrowers' "free-cash-flow" (FCF) rather than monthly income when estimating the model. Free cash flow is defined as the remaining income an individual has after paying taxes and other fixed monthly expenses, like housing payments and/or other debt payments. This is an important empirical adjustment: the median monthly free cash flow ($3,100) is less than half the median monthly income in my dataset. Over 40% of borrowers have a mortgage (which on average translates into a $1,900 payment), and the median monthly fixed expenses for borrowers is $2,400. All of the borrowers have some sort of fixed monthly payment on their credit reports: 40% of borrowers have monthly auto payments which are on average $450, 75% have credit card payments, and 90% have uncategorized installment debt.

I can also directly observe the savings and investment behavior of borrowers in my sample: because individuals in my sample are young, they have relatively low

levels of savings to begin with. Slightly under 40% have a formal retirement savings account – for example 25% have a 401k, with a median balance of $24,000. The number of individuals with investment holdings increases with age. Figure 20 shows that while the median borrower continues to not have substantial savings through age 60, the 75th percentile has accumulated over $80,000 by age 50. However, 90% of my borrowers are under 40 years old, and therefore even the most active savers have investment holdings that are much smaller than their student debt amount.

While FCF is a more accurate depiction of monthly borrower liquidity, the model also assumes that borrowers are not readjusting on other financial margins when refinancing. In other words, contemporaneous savings and debt decisions are assumed to be exogenous, predetermined, and unaffected by maturity and refinancing decisions. I can test this assumption by looking at borrowers' other monthly payments before and after refinancing, and measuring whether they adjust immediately during refinancing. Table 19 in the appendix describes changes in other monthly payments (mortgages, auto loans, credit cards, etc) before vs. after refinancing for individuals who had positive monthly payments to begin with, and shows that for the vast majority of borrowers these stayed constant. This makes sense, since many of these payments are fixed installments, and it would take active work on the borrower's part to readjust them.

4.3.3   Evidence on Permanence of Term Choice:

Our model assumes that borrowers make a maturity choice in year 1 to maximize expected utility over the life of the loan. One might question whether borrowers are actually optimizing over such a long time horizon, or if they are in fact choosing a monthly payment to fit their *current* income level, with the intent to refinance and change term yet again in the future when their income level changes.

To address this, I look at payment patterns over time within my sample of refinancers – in other words, do any individuals keep their payment level over time constant, or do they systematically make higher or lower payments on their debt. I find that there are some extra payments in the data, but they are small and do not vary systematically over time. Figure 17 in the Appendix shows that each month

27

borrowers pay on average 1.5% more than their regular payment, and this is driven by on average only 1% of borrowers making a extra payment each month. There is also no systematic trend in the extra payments. One might expect payments to increase with time as income increases, but here the level of extra payments stays constant over the two year period.

5 Welfare Analysis:

In this section, I use the estimated demand model to analyze how advances in private sector risk-based pricing impact the size and distribution of consumer surplus. My benchmark for these comparisons pools all borrowers (high and low risk) under a uniform interest rate, which represents how the federal loan program would operate with no private refinancing option. I measure the extent of cross-subsidization generated under uniform pricing across risk types and income levels, as well as the deadweight loss. I next introduce a private refinancing option with varying degrees of pricing sophistication, from very coarse (FICO score) to fine grained (our risk score p_i), and measure two main effects: the net increase in consumer surplus, as low risk types refinance into lower, more efficient risk-based prices, and the increase in average cost for the federal program as low risk types select into the private market.

For these exercises, I use the sample of all refinancing applicants – individuals who received a refinancing price quote, but who did not necessarily complete the entire refinancing process. This is different from my estimation sample, which included only approved, agreed refinancers. The applicant sample is more representative of the federal loan portfolio, but the exercise requires that I extrapolate my estimates to a group with a much wider distribution of income, FICO score, and debt amount. To limit the extent of the extrapolation, I restrict the applicant sample to individuals who have a debt-to-income ratio that overlaps with the support of the refinancing sample.

5.1 Supply Side Assumptions:

In the private refinancing sector, I assume that the risk-based interest rates that firms offer $(r(T, p_i))$ are equal to the expected costs of lending to individual i over maturity T. This equates to the assumption of a perfectly competitive refinancing market - i.e. if a firm charged a mark-up, another firm could enter the market and offer a slightly lower price to the same individual while still breaking even. The refinancing market displays most features of perfect competition, including rapid entry into the industry by many firms, and little product differentiation. It is very easy for consumers to price shop and compare price quotes online across refinancing firms. I also estimate very high elasticities (larger than 5) to refinance with respect to offered interest rates in the data, which suggests that price competition across refinancing firms is very high.[25] Importantly, this assumption means that there will be no changes in producer surplus $(PS = 0)$ during the counterfactual, only consumer surplus.

For the federal loan program, I use risk-based discount rates to estimate the size of each per-borrower subsidy under a uniform interest rate regime – specifically, I discount future cash flows under the uniform price regime with the risk, maturity-specific interest rates that would be assigned to that loan in the private sector.[26] The risk-adjusted stream of cash flows, where monthly payments under uniform pricing are given by $d_i(g)$ and term choice is T, is:

$$PRDV_i(g) = \frac{d_i(g)}{r(T, p_i)}[1 - \frac{1}{(1 + r(T, p_i))^T}]$$

The value of the subsidy is given by the difference between the risk-adjusted present value of the loan, and the loan principal (which is equivalent to the present value of the loan without risk adjustment). The subsidy is positive high risk borrowers and negative for low risk borrowers. If the government were to conduct revenue-neutral

---

[25]See Table 7 in the Appendix for the analysis regression underlying this elasticity.

[26]By using these market prices, I focus only on risk that is observable and priced in the private sector but unpriced in the public sector. I also assume that expected losses given risk type and term are the same in the private and public sector. This assumption seems reasonable given that both sector treat default and delinquency similarly.

pricing, the breakeven interest rate g would be defined by:

$$\tilde{g} = \{g : \sum_{i=1}^{N} D_i - \sum_{i=1}^{N} PRDV_i(g) = 0\}$$

5.2 Baseline: Fully Uniform Pricing

As a benchmark for my all my analyses, I assume individuals in my sample are
forced into a uniform pricing scheme at the rate that is revenue neutral . Uniform
pricing will generate deadweight loss as some individuals choose maturities at a
price that is above or below the expected cost of providing to them - the low risk
types end up choosing shorter loans than they would in a setting where they are
charged the cost of providing the loan, and this distortion is large because of their
elastic demand. Figure 3 illustrates how the equilibrium maturity choices of low and
high risk types under uniform pricing ($T_{L,g}$ and $T_{H,g}$) are pushed further apart then
in the efficient setting (T_L^* and T_H^*).The graph also illustrates how uniform pricing
will also increase equity, effectively "taxing" low risk types in order to "subsidize"
high risk types.
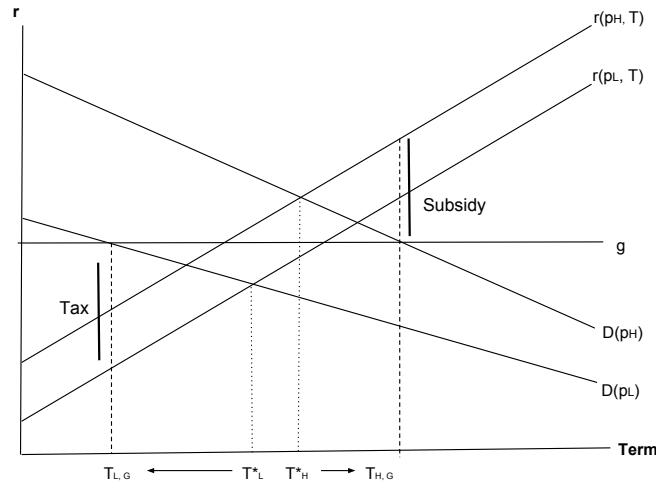
Figure 3: Equity and Efficiency Impact of Uniform vs. Risk-based Pricing

This figure depicts the maturity choices a high and low risk borrower would make in the private sector and the
public sector. There are two maturity demand curves, the lower one for the low risk borrower, and the higher one
for ahigh risk borrower. In the private sector they face the two risk specific price schedules, $r(p_H)$ and $r(p_L)$, and
choose terms T_H^* and T_L^*, which are efficient. In the public sector, they instead both face the uniform price g. The
low risk type chooses a much shorter term $T_{G,L}$ and the high risk type chooses a much longer term $T_{G,H}$ than in
the private sector.

Using our model, we can quantify both the deadweight loss and the redistribu-

tion represented here graphically. The breakeven interest rate for this sample is $g = 6.4\%$, which is in the range of existing Federal Interest Rates for graduate students. To calculate the deadweight loss associated with the uniform interest rate, I calculate the transfer, in addition to the revenue collected, that would make each individual indifferent between a uniform and risk-based pricing regime.[27] Simply returning the revenue to the borrower is insufficient compensation, since the maturity elasticity (i.e. the substitution effect) is nonzero. On average, the per borrower DWL due to the maturity distortion is \$448, or 32% of the average tax/transfer. These calculations suggest that a uniform interest rate is a relatively inefficient means of redistribution, which is unsurprising in light of the high estimated IES and maturity elasticity

Redistribution occurs primarily over risk type, given that risk type directly determines an individuals true "price" and therefore the size of the implicit tax or subsidy they face. On average, individuals who are low risk are taxed \$1,082 under uniform pricing (relative to risk-based pricing), whereas individuals who are high risk (and thus face lower interest rates under uniform pricing) gain an average of \$1,507. The redistribution achieved over a more equity-relevant variable, income, is modest. Figure 4 plots the average subsidy given to each borrower under uniform pricing over both borrower risk type and borrower income. The lowest income borrowers get a subsidy of slightly more than \$1000, while the riskiest borrowers get an average subsidy of almost \$3,000. This is because income is not perfectly correlated with risk type or maturity preferences (the two dimensions that differentiate costs and thus directly generate redistribution), and therefore the uniform rate is an imperfect instrument for achieving redistribution over income.

---

[27]This requires predicting individual's optimal maturity choice under each pricing scenario using the calibrated demand model, and calculating the size of the subsidy or tax they face at that maturity under the uniform rate.
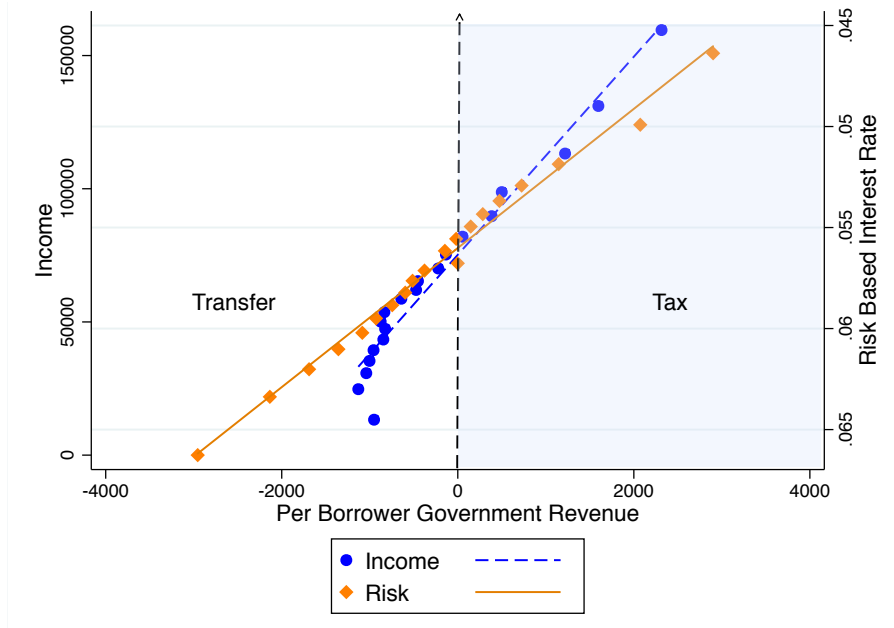
*Figure 4: Redistribution over Risk and Income under a Uniform Interest Rate*

This figure plots the average per borrower tax or transfer under a uniform interest rate policy for 20 income quantiles and 20 risk type quantiles. While the redistribution from lowest to highest risk quantile is large (over $3,000 on average per borrower), less redistribution occurs from lowest to highest income quantile.

5.3 Counterfactual I: Innovations in Risk-Based Pricing and Expansion of the Private Refinancing Market

I next analyze what happens to sorting and welfare in the market as risk-based pricing technologies advance and refinancing firms are able to price on more characteristics. Panel (a) of Figure 5 shows how innovations in risk-based pricing increase the distribution of interest rates charged in the private sector relative to a more coarse measure of borrower quality like FICO score. Here I calculate and plot the 10 year fixed interest rate each borrower would face if the firm could only price on FICO score, as well as the 10 year fixed interest rate each borrower would face at the current "frontier" of risk-based pricing.[28] The graph shows that more comprehensive risk-based pricing expands the distribution of interest rates, in particular extending lower interest rates to the least risky types. The gains to considering additional characteristics are especially large for the student borrower population:

---

[28]I use the observed interest rates schedules from my dataset as empirical proxies for $r(T, p_i)$. The firm estimates borrower risk using a predictive algorithm to estimate the probability of delinquency. This algorithm produces a risk score p_i for each individual based on a vector of characteristics, X_i. This score maps to a schedule of risk, maturity specific interest rates $r(T, p_i)$.

because these borrowers are young and have less-developed FICO scores, this allows them to signal their risk type through other characteristics like degree type, savings behavior, or income.
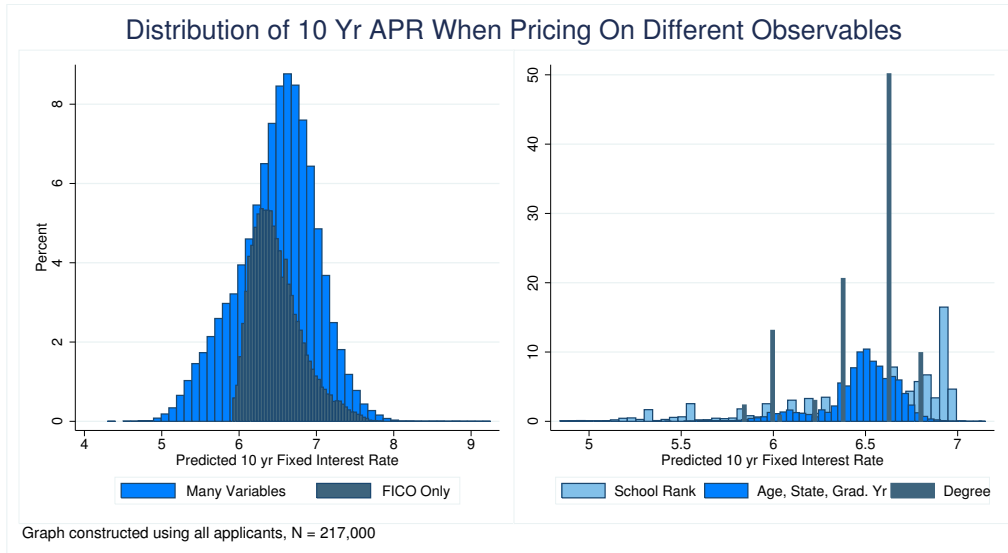
Figure 5: Distribution of Interest Rates when Pricing on Different Observables
Panel (a) shows how innovations in risk-based pricing increase the distribution of interest rates charged in the private sector relative to a more coarse measure of borrower quality like FICO score. Here I calculate and plot the 10 year fixed interest rate each borrower would face if the firm could only price on FICO score, as well as the 10 year fixed interest rate each borrower would face if the firm could price on a more comprehensive set of variables including monthly free cash flow, assets, degree type, and occupation. Panel (b) shows how further restricting the set of variables a firm could price on, for example to school rank or degree type, would substantially reduce the spread of the price distribution.

Using these prices, we can calculate changes in borrower surplus as low risk individuals refinance out of the public sector to take advantage of lower rates. In this initial analysis I assume all individuals who would benefit from refinancing do so. Table 5, Column 1, displays the average gain from risk-based pricing, defined as the change in present discounted cash flows, for individuals who refinance. By leaving the uniform regime, these consumers gain on average $ 1082.52. Using comprehensive risk-based pricing, rather than FICO-based pricing, increases these gains substantially by $341 per borrower.

From the government's perspective, when these low-risk individuals leave the break-even program to refinance, the average risk of the remaining pool increases. Under the assumption that the government maintains the same interest rate even when low risk types exit, I calculate the average per borrower subsidy for the remaining set of borrowers in column 2. The exit of low risk types under full risk-based

Table 5: Impact of Pricing and Refinancing on Surplus and Revenue

Average Changes in Surplus under Different Pricing Schemes			
	Average Tax*	Average Subsidy**	Average DWL***
Complete Pooling, Break-even g	1082.52	1507.48	448.19
Complete FICO-based pricing	580.99	745.80	13.06
Complete Risk-based pricing	0	0	0

Average Changes in Surplus as Individuals Refinance Out of Break-even Pool			
	Avg. $\Delta\ PRDV_i$ for Refinancers [†]	Avg. Δ Subsidy[††]	Avg. Δ DWL[†††]
Refinance into FICO-based Prices	741.20	1259.80	-174.29
Refinance into Full Risk-based Prices	1082.52	1507.48	-221.87

* Average Tax defined as the average *negative* change in present risk discounted value relative to the risk-based price regime. For government pooling this change is defined as, $D_i - PRDV_i(g)$, for FICO-based pricing this is defined as $D_i - PRDV_i(r(FICO_i))$. The present value is risk adjusted using the risk-based interest rates $r(p_i)$, which means that $PRDV_i(r(p_i)) = D_i$.

** Average Subsidy defined as the average *positive* change in present risk discounted value relative to the risk-based price regime.

*** DWL is calculated as the transfer, in addition to the tax or subsidy collected, that would make each individual indifferent between g or $r(FICO_i)$ and $r(p_i)$.

[†] Average $\Delta\ PRDV_i$ for refinancers is defined as $PRDV_i(r(FICO_i)) - PRDV_i(g)$ for individuals refinancing under a FICO based regime. For individuals refinancing under a fully risk based price regime, $\Delta\ PRDV = PRDV_i(r(p_i)) - PRDV_i(g) = D_i - PRDV_i(g)$ = average tax under full pooling.

[††] Avg. Δ Subsidy is defined as the new average subsidy the government will be providing under their original breakeven interest rate for individuals who do not refinance into the private sector

[†††] Δ DWL is defined as $DWL(r) - DWL(g)$. There will be no change in DWL for individuals who do not refinance – $DWL(g) - DWL(g) = 0$. There will be a decrease in DWL for individuals who refinance, but the change will be smaller for individuals refinancing into FICO based prices – $|DWL(r(FICO_i)) - DWL(g)| < |DWL(r(p_i)) - DWL(g)| = DWL(g)$.

pricing increases the subsidy from \$0 to \$1,507 . While innovations in risk-based pricing increased the gains for individuals who refinanced relative to FICO-based pricing, the two pricing schemes generate smaller changes in the new subsidy, only a \$247 value. This is because innovations in risk-based pricing do not increase the extensive *number* of refinancers, but rather extend lower interest rates to individuals who would have already benefited from FICO-based pricing and exited the public pool.

Finally, we can calculate the decrease in deadweight loss that occurs as accurate risk based prices correct the maturity distortion for the group of borrowers who refinance. While FICO-based prices are more accurate than an entirely uniform interest rate, they still generate some (very small) maturity distortions and deadweight loss.

5.4 Counterfactual II: Transitioning to a Net Subsidy

The previous counterfactual showed that as risk-based pricing advances, low risk types will exit the public sector and the originally break-even rate will become an effective subsidy for the remaining borrowers. It is unclear whether the original value of g minimizes the cost of this subsidy. In fact, our model highlights how

policy makers must consider behavioral responses on two budget relevant margins when setting g: maturity choice and refinancing decisions. These responses change the costs associated with charging any given interest rate g.

The graphs in Figure 6 analyze how subsidy costs vary for the government with g, allowing for various behavioral responses by borrowers. The horizontal axis represents the different values of a uniform interest rate g that the government could charge on federal loans. On the vertical axis I plot the amount that the government will *raise* and *spend* on the federal loan portfolio for any value of g, accounting for heterogeneous risk types, maturity choices, and refinancing choices. Therefore, the net interest rate tax or subsidy provided by the government at any given value of g can be found by tracing out the vertical distance between the dotted revenue line and the solid cost line at that point. The value of g that will allow the government to break even on the federal portfolio will be the point where the dotted line and the solid effective cost line intersect.

Panel (a) plots the average cost of the portfolio absent all maturity and refinancing responses (using the maturity choices of borrowers charged risk-based interest rates). This is equivalent to a policy scenario in which the government forced individuals into a specific maturity, and shut down the refinancing channel. Because there are no maturity or refinancing responses, this cost stays constant for all values of g. If the government charged $g = 6.15\%$, the point where cost is equal to revenue, they would break-even on the portfolio.

Panel (b) plots the average cost of the portfolio allowing only for a maturity response, not a refinancing response. This is equivalent to our baseline scenario, in which all borrowers were pooled in the federal portfolio. Note that as the government charges a lower g, individuals extend their maturity which increases the average cost. As they increase g individuals decrease their maturity, which decreases the average cost. The break-even point is slightly higher, $g = 6.39\%$, then in the case where there was no maturity response.

Panel (c) plots the average cost of the portfolio allowing only for a refinancing response, not a maturity response. The average cost curve is always upward sloping in g, since as g increases the lowest risk individuals in the pool will refinance,

increasing the average cost of the remaining pool. This means that it will be impossible to break-even on the portfolio once the refinancing channel is open, but it is possible to minimize the size of the subsidy provided to the remaining borrowers. At $g = 7\%$, the subsidy is minimized at 0.22%; moving from $g = 6.39\%$ to $g = 7\%$ reduces the interest rate subsidy by $.12\%$ from 0.34% to $.22\%$.

Panel (d) plots the average cost of the portfolio allowing for both a refinancing response and a maturity response, the scenario closest to current policy. Individuals, especially higher risk individuals, choose a longer maturities under the flat, uniform g price schedule. This increases the effective subsidy size for smaller values of g, and suggests that the government could minimize the cost of this subsidy by increasing the uniform rate slightly above the rate that assumes away all behavioral responses, to 8.27%.

6   Conclusion

Risk based pricing has advanced in many lending and insurance markets, with real implications for sorting and surplus. In car insurance, firms now use data from tracking devices to cherry-pick the lowest risk drivers. In the private student loan origination market in 2011, 40% of new borrowers had FICO scores greater than 770, while less than 5% had scores below 670. At times, the government acts as a concurrent source of credit or a regulatory body: in the mortgage market, FHA-backed loan eligibility is predicated on risk-related factors like FICO score, with government-provided subsidies for lower income households. In the health insurance market, the government has limited the set of risk-related factors that can determine premiums. This paper provides a framework to analyze and inform the role of government in such settings.

I study how the evolution of risk-based pricing in the student loan market im-placts borrower welfare and government revenue. I show that without a private refinancing option, the governments' uniform interest rate policy achieves modest redistribution over income, but generates sizable distortions in borrowers' intertemporal consumption decisions.[29] By refinancing into a risk-based interest rate, the

(a) Costs vs. Revenue by g, no Responses

(b) Costs vs. Revenue by g, with Maturity Response

(c) Costs vs. Revenue by g, with Refinancing Response

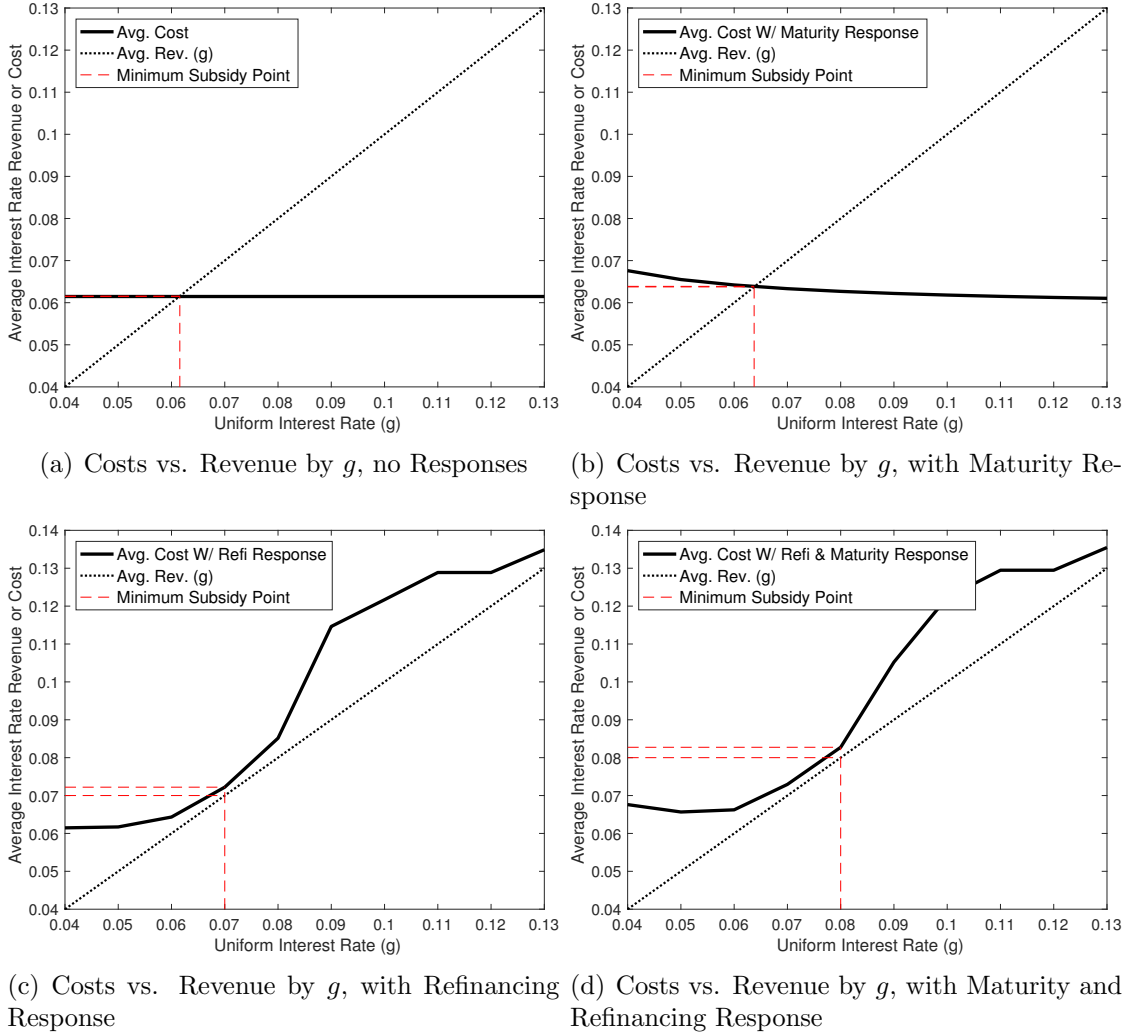(d) Costs vs. Revenue by g, with Maturity and Refinancing Response

Figure 6: Federal Loan Portfolio Cost vs. Revenue by Uniform Interest Rate, Accounting for Behavioral Responses

These graphs analyze how revenue and costs vary for the government with the uniform interest rate that they charge, g, and with various behavioral responses by borrowers. The horizontal axis represents the different values of a uniform interest rate g, that the government could *charge* and collect on federal loans. On the vertical axis, I show how the amount that the government will *raise* and *spend* on the federal loan portfolio, accounting for heterogeneous risk types, maturity choices, and refinancing choices, varies with g. The dotted lines, on the 45 degree angle, plots the uniform interest rate being charged. This is effectively the revenue the government will collect for a given value of g. The solid lines plot the average cost the government faces for a given value of g. Therefore, the net interest rate tax or subsidy provided by the government at any given value of g can be found by tracing out the vertical distance between the dotted line and the solid line at that point. The value of g that will allow the government to break even on the federal portfolio will be the point where the dotted line and the solid effective cost line intersect.

average low-risk borrower can increase surplus by \$1530, but the government will

need to subsidize the remaining high-risk borrower on average \$1507 to maintain

---

[29]While this analysis focuses on cross-sectional redistribution under a uniform rate, the policy also redistributes consumption longitudinally— all borrowers were liquidity constrained when beginning school and unable to secure private credit at price below the Federal rate. Thus even ex-post low risk borrowers benefit from the uniform rate in the interim. The welfare implications are more complex in light of this dynamic selection problem.

equity.

In the student loan space, a subsidy policy will stem unraveling *and* fulfill more precise redistributive motives – rather than implicitly "taxing" low risk borrowers under a break-even interest rate, the transfer could be funded by an income tax, by individuals who are not necessarily borrowers, and allow for intergenerational redistribution. In contrast, the findings suggests that a regulatory policy would primarily reduce efficiency, not inequity. I show that when firms price can only on FICO score, they reduce the gains to low risk borrowers substantially but do not reduce the extent of selection into the private market.

In addition to policy takeaways, this paper presents a novel micro-analysis of student borrowers. It shows there is demand for flexible repayment structures (like a maturity continuum) that allow households to distribute payments optimally over time, and considerable heterogeneity in borrowers' desire to lower monthly payments vs. interest rates. While this analysis focuses on how interest rates impact repayment decisions, student borrowers could also respond to interest rate levels at earlier steps in the borrowing process, for example when taking out debt or deciding whether to attend graduate school. The availability of better risk-based interest rates could change these decisions, and is an interesting area for future work.

References

[1] William Adams, Liran Einav, and Jonathan Levin. Liquidity constraints and imperfect information in subprime lending. *The American Economic Review*, 99(1):49–84, 2009.

[2] Sumit Agarwal, Gene Amromin, Souphala Chomsisengphet, Tomasz Piskorski, Amit Seru, and Vincent Yao. Mortgage refinancing, consumer spending, and competition: Evidence from the home affordable refinancing program. Technical report, National Bureau of Economic Research, 2015.

[3] Orazio P Attanasio, Pinelopi Koujianou Goldberg, and Ekaterini Kyriazidou. Credit constraints in the market for consumer durables: Evidence from micro data on car loans. *International Economic Review*, 49(2):401–436, 2008.

[4] Orazio P Attanasio and Guglielmo Weber. Consumption and saving: models of intertemporal allocation and their implications for public policy. *Journal of Economic literature*, 48(3):693–751, 2010.

[5] Christopher Avery and Sarah Turner. Student loans: Do college students borrow too much?or not enough? *The Journal of Economic Perspectives*, 26(1):165–192, 2012.

[6] Michael Carlos Best, James Cloyne, Ethan Ilzetzki, and Henrik Jacobsen Kleven. Interest rates, debt and intertemporal allocation: Evidence from notched mortgage contracts in the uk. Technical report, Working Paper, 2015.

[7] Harald Beyer, Justine Hastings, Christopher Neilson, and Seth Zimmerman. Connecting student loans to labor market outcomes: Policy lessons from chile. *American Economic Review*, 105(5):508–13, May 2015.

[8] M Kate Bundorf, Jonathan Levin, and Neale Mahoney. Pricing and welfare in health plan choice. *The American Economic Review*, 102(7):3214–3248, 2012.

[9] Caroline Carlin and Robert Town. Adverse selection, welfare and optimal pricing of employer-sponsored health plans. *U. Minnesota Working Paper*, 2009.

[10] Raj Chetty and Amy Finkelstein. Social insurance: Connecting theory to data. Technical report, National Bureau of Economic Research, 2012.

[11] Alma Cohen and Liran Einav. Estimating risk preferences from deductible choice. *The American economic review*, 97(3):745–788, 2007.

[12] James C Cooper, Luke Froeb, Daniel P O'Brien, and Steven Tschantz. Does price discrimination intensify competition? implications for antitrust. *Antitrust Law Journal*, 72(2):327–373, 2005.

[13] Susan Dynarski and Judith Scott-Clayton. Financial aid policy: Lessons from research. Technical report, National Bureau of Economic Research, 2013.

[14] Wendy Edelberg. Risk-based pricing of interest rates for consumer loans. *Journal of Monetary Economics*, 53(8):2283–2298, 2006.

[15] Liran Einav and Amy Finkelstein. Selection in insurance markets: Theory and empirics in pictures. *The Journal of Economic Perspectives*, 25(1):115–138, 2011.

[16] Liran Einav, Amy Finkelstein, and Mark R Cullen. Estimating welfare in insurance markets using variation in prices. Technical report, National Bureau of Economic Research, 2008.

[17] Liran Einav, Amy Finkelstein, and Jonathan Levin. Beyond testing: Empirical models of insurance markets. Technical report, National Bureau of Economic Research, 2009.

[18] Liran Einav, Amy Finkelstein, and Paul Schrimpf. Optimal mandates and the welfare cost of asymmetric information: Evidence from the uk annuity market. *Econometrica*, 78(3):1031–1092, 2010.

[19] Liran Einav, Mark Jenkins, and Jonathan Levin. Contract pricing in consumer credit markets. *Econometrica*, 80(4):1387–1432, 2012.

[20] Liran Einav, Mark Jenkins, and Jonathan Levin. The impact of credit scoring on consumer lending. *The RAND Journal of Economics*, 44(2):249–274, 2013.

[21] Amy Finkelstein and James Poterba. *Testing for Adverse Selection with" unused Observables"*. National Bureau of Economic Research, 2006.

[22] Christa Gibbs. Cfpb datapoint: Student loan repayment. *CFPB Office of Research, August*, 2017.

[23] David B Gross and Nicholas S Souleles. Do liquidity constraints and interest rates matter for consumer behavior? evidence from credit card data. Technical report, National bureau of economic research, 2001.

[24] Jonathan Gruber. A tax-based estimate of the elasticity of intertemporal substitution. *The Quarterly Journal of Finance*, 3(01):1350001, 2013.

[25] Ben Handel, Igal Hendel, and Michael D Whinston. Equilibria in health exchanges: Adverse selection versus reclassification risk. *Econometrica*, 83(4):1261–1313, 2015.

[26] Andrew Hertzberg, Andres Liberman, and Daniel Paravisini. Adverse selection and maturity choice in consumer credit markets: Evidence from an online lender. Technical report, Working Paper, 2015.

[27] Erik Hurst, Benjamin J Keys, Amit Seru, and Joseph S Vavra. Regional redistribution through the us mortgage market. Technical report, National Bureau of Economic Research, 2015.

[28] Matteo Iacoviello. House prices, borrowing constraints, and monetary policy in the business cycle. *The American economic review*, 95(3):739–764, 2005.

[29] David S Johnson, Jonathan A Parker, and Nicholas S Souleles. Household expenditure and the income tax rebates of 2001. *The American Economic Review*, 96(5):1589–1610, 2006.

[30] Dean Karlan and Jonathan Zinman. Expanding credit access: Using randomized supply decisions to estimate the impacts. *Review of Financial studies*, page hhp092, 2009.

[31] Dean Karlan and Jonathan Zinman. Long-run price elasticities of demand for credit: evidence from a countrywide field experiment in mexico. Technical report, National Bureau of Economic Research, 2013.

[32] Dean S Karlan and Jonathan Zinman. Elasticities of demand for consumer credit. *Yale University Economic Growth Center Discussion Paper*, (926), 2005.

[33] Lance Lochner and Alexander Monge-Naranjo. Credit constraints in education. Technical report, National Bureau of Economic Research, 2011.

[34] Lance Lochner and Alexander Monge-Naranjo. Student loans and repayment: Theory, evidence and policy. Technical report, National Bureau of Economic Research, 2015.

[35] Lance J Lochner and Alexander Monge-Naranjo. The nature of credit constraints and human capital. *The American Economic Review*, 101(6):2487–2529, 2011.

[36] Adam Looney and Constantine Yannelis. A crisis in student loans?: How changes in the characteristics of borrowers and in the institutions they attended contributed to rising loan defaults. *Brookings Papers on Economic Activity*, 2015(2):1–89, 2015.

[37] Deborah Lucas. Introduction to" measuring and managing federal financial risk". In *Measuring and Managing Federal Financial Risk*, pages 1–12. University of Chicago Press, 2010.

[38] Deborah Lucas. Credit policy as fiscal policy, 2012.

[39] Deborah Lucas. Valuation of government policies and projects. *Annu. Rev. Financ. Econ.*, 4(1):39–58, 2012.

[40] Deborah Lucas and Damien Moore. Guaranteed versus direct lending: the case of student loans. In *Measuring and managing federal financial risk*, pages 163–205. University of Chicago Press, 2007.

[41] Nuno C Martins and Ernesto Villanueva. The impact of mortgage interest-rate subsidies on household borrowing. *Journal of Public Economics*, 90(8):1601–1623, 2006.

[42] Daniel Paravisini and Antoinette Schoar. The incentive effect of it: Randomized evidence from credit committees. *NBER Working Paper Series*, page 19303, 2013.

[43] Jesse Rothstein and Cecilia Elena Rouse. Constrained after college: Student loans and early-career occupational choices. *Journal of Public Economics*, 95(1):149–163, 2011.

[44] J Brad Schwartz. Student financial aid and the college enrollment decision: The effects of public and private grants and interest subsidies. *Economics of Education Review*, 4(2):129–144, 1985.

[45] Lawrence H Summers. Capital taxation and accumulation in a life cycle growth model. *The American Economic Review*, 71(4):533–544, 1981.

[46] Nicholas Wonder, Wendy Wilhelm, and David Fewings. The financial rationality of consumer loan choices: Revealed preferences concerning interest rates, down payments, contract length, and rebates. *Journal of Consumer Affairs*, 42(2):243–270, 2008.

[47] Jonathan Zinman. Household debt: Facts, puzzles, theories, and policies. Technical report, National Bureau of Economic Research, 2014.

A   Modeling Borrower Delinquency

A.1   Borrower problem: Impact of delinquency risk

In this section I solve the borrowers' problem, allowing them to consider the impact of delinquency on their expected utility. Borrowers of risk type p and debt amount D choose maturity T to maximize:

$$\max_T \int_1^Q \mathbb{E}[u(c_t)]dt$$

where Q is the individuals' maximum age (i.e. finite). I define consumption as post-tax income w_t minus a student debt payment $d(T)$ if the loan is still being paid off – this implies that individuals are "hand-to-mouth" and not smoothing consumption through other debt or savings.[30]

The debt payment $d(T)$, that individuals make each period is a function of their total debt amount D, their chosen maturity, T, and the risk, maturity-specific interest rate schedule that they are offered, $r(T, p)$:

$$d(T) = T * D * \frac{r(T, p)}{(1 - (1 + r(T, p))^{-T})}$$

In the borrowers' maximization problem, agents can decrease their debt payment by increasing their maturity – i.e. $\frac{d\,d(T)}{dT} < 0$. This is the only "choice" variable that the borrower has – I assume that the interest rate schedule a borrower is offered is exogenous (determined by the lender) and therefore cannot be manipulated by the borrower to change their monthly payments or financing costs.

While borrowers can control their yearly payment level, they cannot control their variable, growing income stream which I parametrize as:

$$ln(w_t) = ln(w_0) + g * t + u_t$$

$$u_t \sim N(0, \sigma_u^2(p))$$

---

[30]I address this assumption empirically in the robustness section.

This specification says that each period income grows by g percent, and is hit with a one-period income shock u_t that has a risk-type specific variance.

Since income is variable, it is possible that a bad income shock could force consumption as currently defined $(w_t - d(T))$ to a negative level. Therefore I specify a certain minimum threshold (x) that consumption never falls below. If ever $(w_t - d(T) < x)$, an individual will not pay their entire debt payment in that period and instead consume x, creating a discontinuity in the consumption function. Specifically, the consumption function is given by:

$$c_t = (w_t - d(T)) * \mathbb{1}(w_t - d(T) \geq x) + x * \mathbb{1}(w_t - d(T) < x) \qquad \text{for } t \leq T$$

$$c_t = w_t \qquad \text{for } t > T$$

If a borrower ever misses a payment (or part of a payment) in a period, the borrower will have to continue making payments on the loan as scheduled for the remainder of the T periods *and* face a large utility penalty, which I denote as G, after $T + 1$. If G is large enough, this will rule out the possibility of "strategic" default on the borrowers' part.[31] If an individual never misses a payment, they will consume their entire income in every period after period T and not face a penalty.

I assume that both the firm and the borrower have symmetric expectations about the borrowers' income process, and therefore of $Pr(w_t - d(T) < x)$. In the following section, I empirically back out the $Pr(w_t - d(T) < x)$ that is implied from each borrower/maturity specific interest rate that is offered by the firm.

To simplify the maximization problem, I make a key assumption that the probability of delinquency does not change over time and is equal to the borrower's initial delinquency risk; i.e. $Pr(w_t - d(T) < x) = Pr(w_0 - d(T) < x) = P(T)$.[32] While $P(T)$ does not change over time, I still allow the probability of delinquency to change with the chosen maturity; specifically, $\frac{dP(T)}{dT} < 0$, since a longer maturity

---

[31]Specifically, G must be large enough that totally utility after missing a payment is smaller than total expected utility without missing a payment.

will lower the the monthly payment in any given period.

I also make the delinquency penalty G linearly additive to utility from consumption. These simplifications make the maximization problem:

$$= \max_{T} \int_{1}^{T} \overbrace{\mathbb{E}[u(w_t - d(T))|w_t - d(T) > x]}^{\text{EU in period if not delinq.}} * \overbrace{(1 - P(T))}^{\text{Prob. not delinq.}} + \underbrace{\overbrace{u(x)}^{\text{U if delinq.}} * \overbrace{P(T)}^{\text{Prob. delinq.}}}_{\text{}} \ dt +$$

$$\underbrace{\int_{T+1}^{Q} \overbrace{\mathbb{E}[u(w_t)]}^{\text{EU if never delinq.}} \ dt}_{\text{Remaining periods not paying back loan}} - \overbrace{G}^{\text{Penalty if ever delinq.}} * \overbrace{P(T) * T}^{\text{Prob. ever delinq.}}$$

Differentiating the borrower problem with respect to T, we can see the exact impact of a change in maturity on expected utility through the first order condition. Allowing delinquency to enter the problem also adds two additional terms to the simplified first order condition:

$$0 = \underbrace{\int_{1}^{T} \mathbb{E}[-u'(w_t - d(T))|w_t - d(T) > x] * \frac{dd(T)}{dT} * (1 - P)dt}_{\text{A: Increase in utility from having a slightly lower payment each period}}$$

$$+ \underbrace{\int_{1}^{T} (u(x) - \mathbb{E}[u(w_t - d(T))|w_t - d(T) > x]) * \frac{dP}{dT}dt}_{\text{B: Change in EU due to change in probability of being delinquent}}$$

$$+ \underbrace{(\mathbb{E}[u(w_{T+1} - d(T))|w_{T+1} - d(T) > x] * (1 - P) + u(x) * P) - \mathbb{E}[u(w_{T+1})]}_{\text{C: Decrease in utility from having to pay for one additional period}}$$

$$- \underbrace{G * (P(T) + T * \frac{dP(T)}{dT})}_{\text{D: Change in expected penalty payment due to longer loan period}}$$

An increase in maturity will now impact expected utility in four main ways:

- It lowers the size of the monthly payment, which increases utility while pay-

---

[32]Note that:

$$Pr(w_0 - d(T) < x) = Pr(w_0 < x + d(T))$$
$$= Pr(ln(w_0) + u_t < ln(x + d(T)))$$
$$= Pr(u_t < ln(x + d(T)) - ln(w_0))$$
$$= \frac{\Phi(ln(x + d(T)) - ln(w_0))}{\sigma_u(p)}$$

ing off the loan (Effect **A**). Because agents are hand-to-mouth, they can better smooth consumption with a lower monthly payment – there will be a small "jump" in utility when they finish paying off the loan. A second utility smoothing benefit comes from the fact that income is risky and a lower monthly payment provides some "insurance" value against more volatile income shocks. However, the interest rate $r(T)$ increases with maturity, so $d(T)$ is decreasing in T at a *decreasing* rate.

- It changes the length of repayment, which means the borrower has to pay the loan off for one additional period, and this lowers total expected utility (Effect **C**).

- It decreases the probability of delinquency, ($P(T)$, in any single period since there is a lower monthly payment. *During* repayment, this increases expected utility by lowering the chance that you consume at the threshold x (Effect **B**). *After* repayment, this increases expected utility by lowering the chance that you have to pay the penalty G (Effect **D**).

- It increases the probability of delinquency over the life of the loan since you are paying off over more periods. This increases the expected penalty G (Effect **D**).

The main specification captures the consumption smoothing effects, A and C, of a lower maturity; the delinquency threshold adds the additional effects, B and D. The FOC allows one to say something about the size and sign of these effects: the size of effect B is bounded mechanically – when $(w_t - d(T))$ is small, $(\mathbb{E}[u(w_t - d(T))|w_t - d(T) > x] - u(x))$ is close to zero. Conversely, when $(w_t - d(T))$ is large, there is a small probability of delinquency and $\frac{dP}{dT}$ is near zero. This means effect B is positive but small. Effect D has an ambiguous sign and size. As maturity increases, each period the chance of delinquency becomes smaller; yet, the number of periods during which delinquency can occur increases. This means that $(P(T) + T * \frac{dP(T)}{dT})$ could be positive or negative. The size of D also depends on the size of G, the delinquency "penalty", which is unobserved.

In the empirical section I test how sensitive estimates of γ are to the inclusion or exclusion of effect D. I calibrate $(P(T) + T * \frac{dP(T)}{dT})$ using the empirical delinquency probabilities implied by observed risk-based interest rates, and test the sensitivity of γ to a wide range of values of G. The results suggests that the estimates are not very sensitive to the exclusion of the delinquency channel – for this low risk group, the first order effect of a change in maturity seems to operate instead through the monthly payment consumption smoothing channel.

A.2 Lender Problem: Estimating Income Risk Implied by Interest Rates

Estimation of the borrower's problem with or without delinquency requires empirical inputs for σ_i^2 and possibly $P(T)$; however, our dataset only captures individuals when they first refinance their loans, providing no directly-observed longitudinal data on income volatility or delinquency rates. One can recover individuals' implied income volatility using the lender's problem, which links observed interest rates to delinquency probabilities and thus income risk.

The lenders offer risk, maturity specific interest rates $r(p, T)$ to each borrower that I observe in the data. I assume that lenders:

- Are perfectly competitive, and therefore set interest rates such that they are indifferent between lending to a risky borrower at maturity T and lending to a risk-free borrower at maturity T and interest rate $i(T)$.

    - I assume that this risk-free interest rate $i(T)$ also incorporates the other fixed costs that the company must incur when lending (i.e. origination costs and cost of capital), and therefore the only difference between $i(T)$ and $r(T)$ is the delinquency risk premium.

- Have a recovery rate of α if a payment is delinquent.

From the borrower's problem, recall that a borrower will be delinquent on a loan in a given period if $w_t - d(T) < x$. This means that the probability that a borrower

is delinquent on a T maturity loan in a given period is:

$$P(T) = \frac{\Phi(ln(x + d(T)) - ln(w_0))}{\sigma_u(p)} \tag{1}$$

Given the delinquency probability and recovery rate, a lender's expected stream of payments on a loan of maturity T to risk type p is:

$$T * \left((1 - P(T)) * D * \frac{r(T, p)}{(1 - (1 + r(T, p))^{-T})} + \alpha * P(T) * D * \frac{r(T, p)}{(1 - (1 + r(T, p))^{-T})} \right)$$

The expected stream of payments on a risk-free loan of maturity T is:

$$T * \left(D * \frac{i(T)}{(1 - (1 + i(T))^{-T})} \right)$$

Therefore, lenders set $r(p, T)$ to satisfy the condition:

$$P(T) = \left(\frac{i(T)}{(1 - (1 + i(T))^{-T})} * \frac{(1 - (1 + r(T, p))^{-T})}{r(T, p)} - 1 \right) * \frac{1}{(\alpha - 1)} \tag{2}$$

I observe everything on the right-hand side of this equation, which allows me to calculate empirical delinquency probabilities for every individual in the dataset, which I call $P(\hat{T})$. Combining $P(\hat{T})$ with equation (1)[33] allows me to then calculate individual specific $\hat{\sigma}_i$:

$$\hat{\sigma}_i = \frac{\Phi(ln(x + d(T)) - ln(w_0))}{P(\hat{T})}$$

A.3   Calibration of Delinquency Probabilities and Income Volatility

I begin with the lenders' problem, which allows one to retrieve the delinquency probabilities implied by the risk-specific interest rates offered to each borrower. I use formula 2, plugging in the empirical values of $i(T)$, $r(T)$, α, and T for each borrower. I directly observe $r(T)$ and T for each borrower, and infer values for α and $i(T)$ from the firms' cost accounting documentation. Table **??** lists the summary statistics for these variables.

---

[33]in which I observe all quantities other than σ_i

The calibrated recovery rate $\alpha = .1$ may seem low, but note that this reflects the model's per-payment recovery assumption – i.e. if a borrower misses a single payment, the lender will recover α of that period's payment, and then continue to collect the remaining loan capital at the full rate. In reality, the lender often charges-off the entire loan to a collections agency, which means they may have a higher *average* recovery rate that applies to the entire loan principal, but a very low recovery rate in any single period (the model's α).

The "risk-free" rate of lending $i(T)$ captures various costs of lending for the firm at that maturity, including the cost of capital, the cost of customer acquisition, and the cost of servicing the loan. In fact, the only cost *not* included in this measure is the expected cost of non-repayment (i.e. delinquency). Therefore, $i(T)$ is the same for borrowers of all risk types, but does vary with the maturity of the loan. In contrast, $r(T)$ varies with both the risk type of the borrower and the maturity of the loan.

Given these values, and the assumptions of the lender model, the calculated values of $\hat{P}(T)$ range anywhere from .003 to .15 (see Figure **??** for a histogram of the calculated probabilities). While we cannot compare these implied probabilities to delinquency rates observed in this dataset, since the loans have long maturities and are still in the early stages of being paid off, they can be compared to delinquency rates for other historical student loan portfolios. For example, historical delinquency and default rates for graduate students in the government's Direct Loan portfolio range from 3-7%.

Observed Variables	Definition	Data Used	Mean	Median	SD
w_0	Annual Income	Initial Observed Post-tax Annual Income	80,980	70,080	43,947
$d(T)$	Yearly Debt Payment at Maturity T	Yearly Debt Payment at Maturity T	10,951	7,860	8,993
T	Loan Maturity	Maturity in Years	7	5	4
$r(T)$	Risk-based Interest rate at Maturity T	Observed $r(T)$ %	4.91	4.91	0.86
$i(T)$	Risk-Free Interest Rate at Maturity T	Lowest Risk Type $r(T)$ - ϵ in %	3.87	3.50	0.58

Calibrated Parameters	Definition	Calibration Value Used	Mean	Median	SD.
α	Recovery Rate on Per-Period Payment	0.1			
x	Min. Consumption Threshold	Yearly Housing Expense	32,868	28,902	19,475

Estimated Parameters	Definition	Mean	Median	SD.
$P(T)$	Per-Period Delinquency Probability	0.042	0.044	0.027
σ_u	Standard Deviation of $ln(w)$	0.19	0.15	0.18

Table 6: *Empirical Income Volatility Exercise: Calibration Values and Estimates*
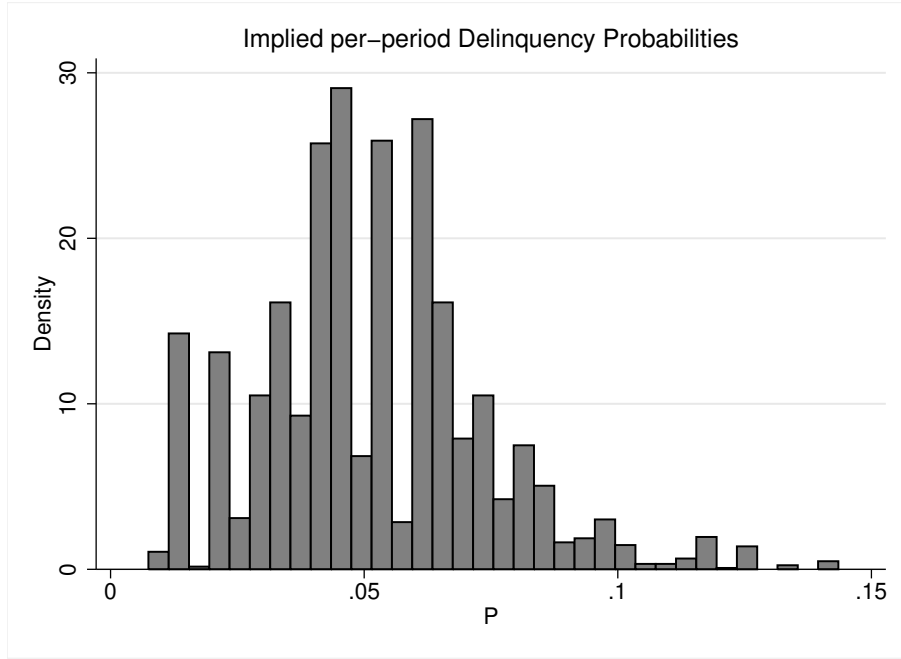
Figure 7: Distribution of calculated values of $\hat{P}(T)$

I next transition to the borrower problem, calculating the individual-specific income volatility that would generate the estimated values of $\hat{P}(T)$. This relies on equation 1, which relates the delinquency rate to the probability that an income shock will push log consumption below a certain minimum threshold. Specifically, under the assumption that income is log-normally distributed, equation 1 maps $\hat{P}(T)$ to σ_u.

For these calculations, I use post-tax annual income as an empirical measure of w_0, and individuals' verified yearly housing expenditures as a proxy for x, the minimum consumption threshold. This measure of x assumes that individuals will first pay their rent or mortgage before making their student debt payment. For the median borrower, this threshold is around 40% of their post-tax income or \$28,500. Note that as x becomes larger relative to w_0, σ_u will become smaller: a smaller shock to income is needed to push an individual into delinquency.

The resulting income variance estimates are described in the final row of table 6. On average, σ_u=.19, off a mean value of $ln(w)$ of 11.16. To make these numbers easier to interpret, figure 8 converts from logs to levels: for each level of post-tax yearly income, it plots the one and two standard deviation range calculated using
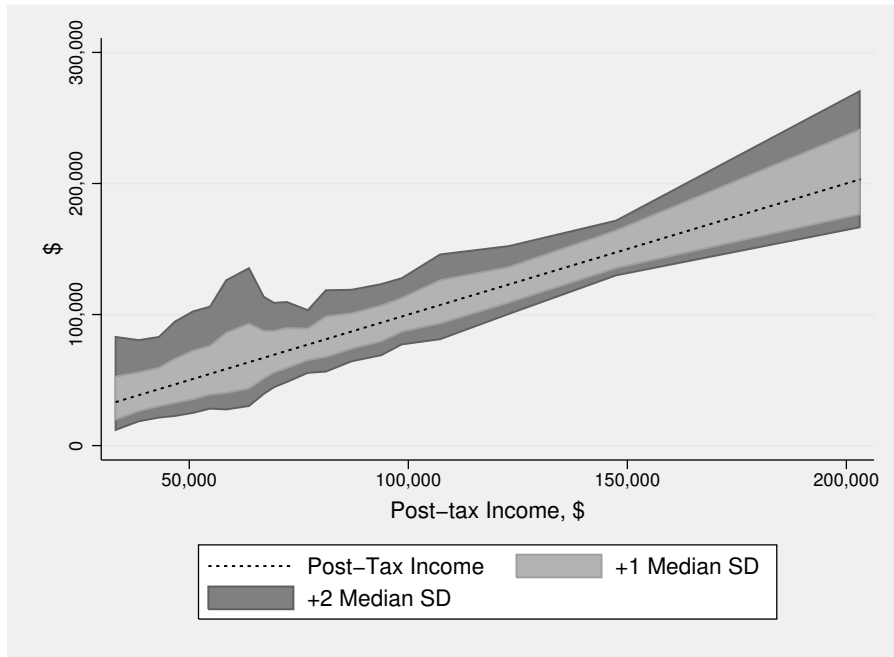
Figure 8: Implied post-tax income ranges, using median estimated income variance σ_u
For each post-tax yearly income level (x-axis), this graph plots the range +/- 1 standard deviation
and +/- 2 standard deviations around that income level, using the median income variance σ_u^2
estimated from the lender's problem. Given the log-normal income assumption, there is a 68
percent chance that income each year will fall within the 1 S.D. range, and 95 % chance that
income each year will fall within the 2 S.D. range

the median value of σ_u for individuals in that income category. This graph says,
for instance, that the median individual initially making \$50,000 has a 68% chance
of making anywhere between \$40,000 and \$82,000 in each of the following years.

A.4 Calibration of Income Growth Rates

In addition to calibrating σ_u, the borrowers' maturity problem also requires
calibration of income growth rates (g_i). While one would ideally use panel data to
directly observe the income growth of my borrowers, the long time horizon of the
debt contracts (up to 20 years) is a limiting factor. I instead use a cross section of
observationally similar individuals at various ages to estimate a pseudo age-income
profile.

The dataset I use to estimate these cross-sectional profiles contains individuals
who are similar to my refinancing applicants in many important respects (high in-
come, high FICO, mainly graduate degree recipients), but who are applying instead

for small personal loans rather than applying to refinance student debt. This distinction is important when estimating cross-sectional age profiles – if I instead used a cross-section from the student loan borrower population, one might worry that individuals refinancing student debt at age 40 have very different income trajectories than those refinancing at age 30. Here the worry is that individuals borrowing small amounts ($5,000 - $15,000) at different ages have fundamentally different earning trajectories. While this selection concern is valid, one must weigh it against the fact that this population is similar to my borrowers in many unique respects that would be difficult to find and match to in a survey dataset like the CPS. These include both tangible characteristics, like degree type, income level, or FICO score, as well as intangible characteristics. For example, my population is refinancing with a new internet-based bank, which makes them potentially different, or more tech savvy, then a population that uses only traditional banks. Furthermore, because my sample has a high socioeconomic status, they make up only a small percentage of most representative survey samples.

This dataset contains approximately 250,000 borrowers who range in age from 20 to 50. Figure 18 gives a sense of what the cross-sectional income trajectories look like for individuals in this sample – it plots yearly post-tax income after separating individuals into 4 degree levels: associates, bachelors, masters, and professional.

Using this sample, I estimate degree-specific growth rates that match the log-income parametrization of my model with the following regression:

$$ln(y_i) = \beta_0 + \beta_1 * age_i + \gamma_0 * degree_i + \gamma_1 * age_i * degree_i + e_i$$

Here $ln(y_i)$ is log post-tax yearly income, and $degree_i$ are dummies indicating highest degree level. The coefficient $\hat{\beta}_1$ estimates the average yearly growth rate of log income over the life-cycle, while $\hat{\gamma}_1$ allows this growth rate to deviate by degree type. The estimated growth rate for any borrower in my student loan dataset is thus given by $\hat{g}_i = (\hat{\beta}_1 + \hat{\gamma}_1)$.[34] Figure 9 below shows estimates of $\hat{\beta}_1$ and $\hat{\gamma}_1$.

---

[34]While we would like to estimate g_i that vary by many characteristics, from number of dependents to FICO score, the use of cross-sectional data only allows for comparison on *time-invariant* characteristics like degree type, or occupation.
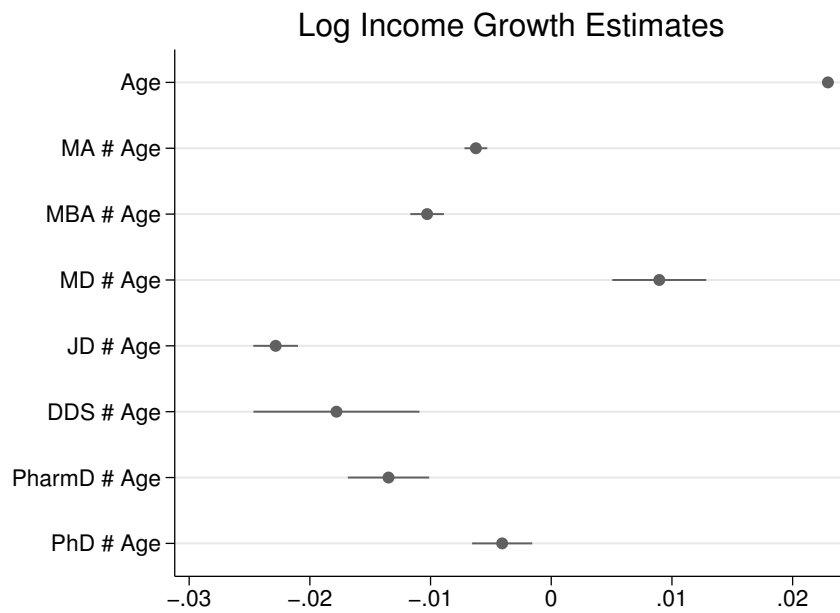
Figure 9: Estimated log income growth rates

B Derivation of Analytical First Order Condition:

Analytical Estimation:

When choosing a term, individuals chose T to maximize the discounted stream of yearly utility, which lead to the first order condition:

$$E[\sum_1^T \beta^t \frac{\partial d}{\partial T}(w_{it} - d_i)^{-\gamma}] = E[\beta^{T+1}(-d_i)(w_{iT})^{-\gamma}]$$

$$s.t. \ d_i = T * D_i * \frac{r(T, p_i)}{(1 - (1 + r(T, p_i))^{-T})}$$

d_i represents the yearly payment for individual i at term T, and $r(T, p_i)$ is the risk, term specific interest rate faced by individual i at term T.

I assume that log income follows the unit root process:

$$ln(w_{it}) = ln(w_{it-1}) + (X_i'\mu) + u_{it}$$

where $X_i'\mu$ is a yearly growth rate specific to observable characteristics and

$$u_{it} \sim N(0, \sigma_u^2)$$
$$\sigma_u^2 = (\omega - v * p_i)^2$$

is a individual-specific yearly income shock that is allowed to be a function of observable risk type p_i.

We observe starting income levels w_{i0}. This means that we can express income at time t as:

$$ln(w_{it}) = ln(w_{i0}) + t * (X_i'\mu) + \sum_1^t u_{it}$$
$$w_{it} = w_{i0} * e^{t*(X_i'\mu) + \sum_1^t u_{it}}$$

56

If we return to the uncertain portions of the right hand side of our first order condition, $E[(w_{it} - d_i)^{-\gamma}]$, note that we can rewrite the expected marginal utility as the marginal utility of a certainty equivalent given by:

$$E[(w_{i0} * e^{t*(X_i'\mu)} * e^{\sum_1^t u_{it}} - d_i)^{-\gamma}] = (w_{i0} * e^{t*(X_i'\mu)} * e^{\pi_{it}} - d_i)^{-\gamma}$$

where π_{it} is the certain amount an individual would have to be given in that period to make their certain utility equivalent to the expected utility. Specifically:

$$\pi_{it} = \frac{1}{2} * t * \sigma^2 [1 - (1 + \gamma)\frac{w_{i0} * e^{t*(X_i'\mu)}}{w_{i0} * e^{t*(X_i'\mu)} - d_i}] \qquad \text{for t<T+1}$$

$$\pi_{it} = \frac{1}{2} * t * \sigma^2 (-\gamma) \qquad \text{for t≥T+1}$$

To derive π_{it}, note that we can write:

$$E[u'(w_{it})] = E[u'(w_{i0} * e^{t*(X_i'\mu)} * e^{\sum_1^t \sigma\epsilon_{it}})]$$

where $\epsilon_{it} \sim N(0, 1)$. We want to find the value of $\pi(\sigma)$ that allows us to write:

$$E[u'(w_{i0} * e^{t*(X_i'\mu)} * e^{\sum_1^t \sigma\epsilon_{it}} - d_i)] = u'(w_{i0} * e^{t*(X_i'\mu)} * e^{\pi(\sigma)} - d_i)$$

For simplicity, start with the case of no income growth in period 1.

$$E[u'(w_{i0} * e^{\sigma\epsilon_{i1}} - d_i)] = u'(w_{i0} * e^{\pi(\sigma)} - d_i)$$

We first take the derivative of this expression w.r.t. σ:

$$E[w_{i0} * \epsilon * e^{\sigma\epsilon_{i1}} u''(w_{i0} * e^{\sigma\epsilon_{i1}} - d_i)] = \pi'(\sigma)w_{i0} * e^{\pi(\sigma)} u''(w_{i0} * e^{\pi(\sigma)} - d_i)$$

At $\sigma = 0$ this becomes zero since $E[\sigma\epsilon] = 0$ and thus $\pi'(0) = 0$.

We next take the second derivative of this expression w.r.t. σ, and evaluate it at

57

$\sigma = 0$:

$$E[\epsilon^2 u''(w_{i0} - d_i) + \epsilon^2 w_{i0} u'''(w_{i0} - d_i)] = \pi''(0) u''(w_{i0} - d_i)$$

$$\pi''(0) = [1 + w_{i0} \frac{u'''(w_{i0} - d_i)}{u''(w_{i0} - d_i)}]$$

Under the assumption of CRRA utility, this becomes:

$$\pi''(0) = [1 + w_{i0} \frac{u'''(w_{i0} - d_i)}{u''(w_{i0} - d_i)}]$$

$$= [1 - (1 + \gamma) \frac{w_{i0}}{w_{i0} - d_i}]$$

We now have a value for $\pi''(0)$. This is helpful when evaluating a Taylor expansion of $\pi(\sigma)$:

$$\pi(\sigma) \approx \pi(0) + \pi'(0)\sigma + \frac{1}{2}\sigma^2 \pi''(0)$$

$$\pi(\sigma) \approx \frac{1}{2}\sigma^2 [1 - (1 + \gamma) \frac{w_{i0}}{w_{i0} - d_i}]$$

Therefore our analytical estimating moment becomes:

$$g_i(\theta) = \sum_1^T \beta^t \frac{\partial d}{\partial T}(w_{i0} * e^{t*(X_i'\mu)} * e^{\pi_{it}} - d_i)^{-\gamma} - \beta^{T+1}(-d_i)(w_{i0} * e^{(T+1)*(X_i'\mu)} * e^{\pi_{i(T+1)}})^{-\gamma}$$

To estimate the model, I use nonlinear least squares, choosing the parameters that minimize the quadratic form:

$$b = \arg\min_\theta g_i(\theta)' g_i(\theta)$$

C Evidence of Price Competition in the Private Sector

In the counterfactual, I make the assumption that the private refinancing market is perfectly competitive. Importantly, this assumption means that there will be no changes in producer surplus ($PS = 0$) during the counterfactual, only consumer surplus. On the surface, refinancing market displays most features of perfect competition, including rapid entry into the industry by many firms, and little product differentiation. It is also very easy for consumers to price shop and compare across refinancing firms, due to their online nature and the fact that they all offer quick, personalized price "quotes".

In this section I provide empirical evidence of very strong price competition by estimating empirical refinancing elasticities with respect to quoted APR. I estimate the model using the logistic regression framework:

$$Pr(R_i = 1) = X_i'\mu - \beta r(10, p)_i + v * p_i + \epsilon_{ij}$$

The dependent variable (R_i) indicates whether an individual refinanced with this specific refinancer after seeing an interest rate quote. The dependent variable of interest is the 10 year fixed APR ($r(10, p)_i$) an individual was offered. I again isolate and use only the firm-conducted price changes as a source of $r(10, p)_i$ variation, controlling directly for risk type in the regression. The price elasticity is therefore identified only off of within risk type variation that happened at this specific refinancing firm.

If price competition is high, then a small increase in $r(10, p)_i$ at this single firm should create a large decrease in $Pr(R_i = 1)$ at this firm. Using β to measure price competition assumes that the within-risk type price shifts were made independently of price changes at other firms and the Direct Loan program. This assumption seems reasonable, since the price changes were conducted to collect information on borrower elasticities, and not for competitive reasons. If they were made in tandem with other firms, this will bias our elasticity downwards.

The results of the main specification are shown in column 1 of Table 7, expressed as both elasticities and marginal effects. They provide evidence of strong price

Table 7: Extensive Margin Refinancing Elasticities

	Within Risk Type Variation Only*	Within and Over Risk Type Variation**
Elasticity	-5.63	-5.33
$dy/dx \times (x/y)$	(.182)	(.070)
Marginal Effect	-.00103	-.00098
dy/dx	(.000033)	(.00001)
Variable	*Mean*	*Std. Dev.*
$y=$ Prob. Refinance	.05	
$x = $ 10 Yr Fixed APR (bps)	562.56	57.42
N	202930	202930

Standard errors in parentheses. Specification are logistic regressions of dummy variable if individual fully refinances their loan at this company on the 10 year fixed APR they were offered. Results are expressed both as marginal effects, the unit change in y for a unit change in x, and as elasticities, the proportional change in y for a proportional change in x. Means and std. of the dependent and independent variable are shown in the bottom rows.

*Regression controls for risk type at the level of within risk type price variation (categorically). Therefore, variation in the 10 yr fixed APR comes only from over-time shifts in the price schedule.

**Regression does not control for risk type. Therefore, variation in the 10 yr fixed APR comes from *both* over-time shifts in the price schedule and risk-based prices.

competition: when the interest rate at this one firm increases by 1%, borrowers are 5.6% less likely to refinance at this specific firm.

Column 1 uses only within risk-type price variation, where as the regression in column 2 does not control for risk type and therefore uses both within and across risk type price variation. It is interesting to note that both find a very similar elasticity, which suggests that the extensive margin sensitivity to refinance does not differ over risk type. This helps support our assumption during the maturity analysis that the results were not driven by compositional changes; i.e. changes in the sample were not correlated with taste for maturity.

D Additional Figures and Tables

Table 8: Budget Lifetime Default Rates

Year Loan Enters Repayment	2007	2008	2009	2010	2011
2-Year Non-Profit/Public	32.10%	31.60%	31.10%	31.40%	33.80%
2-Year Proprietary	47.30%	48.60%	49.00%	48.40%	49.40%
4-Year Freshmen & Sophomores	24.70%	24.00%	23.60%	24.20%	25.40%
4-Year Juniors & Seniors	12.40%	12.30%	12.10%	11.90%	13.00%
Graduate Students	6.20%	6.20%	6.10%	6.10%	6.40%
Weighted Average	*15.90%*	*16.50%*	*17.30%*	*17.60%*	*18.40%*

Source: U.S. Department of Education (based on figures published in fiscal year 2014)

Figure 10: Description of Federal Loan Repayment Plans

	Eligible Loan Types	Term	Monthly Payment
Standard Fixed Repayment	• Direct Subsidized and Unsubsidized Loans • Subsidized and Unsubsidized Federal Stafford Loans • All PLUS loans	• Up to 10 Years	• Fixed amount, at least $50 per month
Graduated Repayment	• Direct Subsidized and Unsubsidized Loans • Subsidized and Unsubsidized Federal Stafford Loans • All PLUS loans	• Up to 10 Years	• Payments are lower at first and then increase every two years
Extended Repayment	• Direct Subsidized and Unsubsidized Loans • Subsidized and Unsubsidized Federal Stafford Loans • All PLUS loans	• Up to 25 Years	• Payments may be fixed or graduated
Income Based Repayment	• Direct Subsidized and Unsubsidized Loans • Subsidized and Unsubsidized Federal Stafford Loans • All PLUS loans made to students • *Consolidation* Loans (Direct or FFEL) that do not include Direct or FFEL PLUS loans made to parents	• Up to 25 Years • If loan not repaid in full after making the equivalent of 25 years of qualifying monthly payments, any outstanding balance will be forgiven.	• Maximum payment is 15% of discretionary income 　o Discretionary income is the difference between your AGI and 150% of the poverty guideline • Payments change as income changes • Individual needs to demonstrate partial financial hardship to qualify
Pay as You Earn Repayment	• Direct Subsidized and Unsubsidized Loans • Direct PLUS loans made to students • Direct Consolidation Loans that do not include (Direct or FFEL) PLUS loans made to parents • Must be a new borrower on or after Oct. 1, 2007, and must have received a *disbursement* of a Direct Loan on or after Oct. 1, 2011.	• Up to 20 Years • If loan not repaid in full after making the equivalent of 20 years of qualifying monthly payments, any outstanding balance will be forgiven.	• Maximum payment is 10% of discretionary income • Payments change as income changes • Individual needs to demonstrate partial financial hardship to qualify

Source: https://studentaid.ed.gov/repay-loans/understand/plans

This table describes the various repayment plans available for Federal Direct Loans as of 2015.

Table 9: Impact of Debt, Income, APR, and Risk on Term

	(1)	(2)	(3)
ln(Income)	-17.86***	-15.94***	-16.73***
	(2.349)	(2.385)	(2.381)
ln(D)	34.35***	34.05***	34.38***
	(1.083)	(1.083)	(1.084)
Age	1.016***	1.072***	1.039***
	(0.122)	(0.122)	(0.122)
Home Owner	8.294***	8.636***	8.611***
	(1.722)	(1.715)	(1.719)
Variable Rate	2.570	3.162*	2.683*
	(1.632)	(1.626)	(1.627)
Avg. APR	2143.4***	-2355.5***	1217.2
	(198.4)	(847.8)	(916.0)
Risk Score		-15.69**	5.286
		(6.890)	(7.364)
Avg. APR * Risk Score			-306.8***
			(116.1)
Constant	-229.9***	124.6*	-130.9*
	(32.31)	(67.15)	(67.00)
σ			
Constant	58.53***	58.18***	58.36***
	(0.638)	(0.634)	(0.636)
N	11663	11663	11663

Standard errors in parentheses

* $p < 0.10$, ** $p < 0.05$, *** $p < 0.01$

This table displays results from a series of Tobit regressions of price and borrower characteristics on term choice (which is truncated at 60 and 240 month). The first specification pools both risk and temporal variation in the Avg. APR variable – because higher risk borrowers (who face higher APRs) prefer longer loans, this regression suffers from omitted variable bias. It seems as though higher APRs drive individuals to increase their term choices. Specification (2) controls directly for risk score and therefore the only remaining variation in avg. APR comes from temporal price changes that were independent of borrower characteristics. Specification (3) allows risk score and price to interact, thereby allowing different risk types to have different price sensitivities.

Table 10: Test of Extensive Margin Response and Changes in Borrower Composition

	Observables over Price Regimes		
	Avg. APR		
	Coeff.	SE	t
ln(Income)	-.0000367	.0000362	-1.01
ln(Debt)	-9.55e-06	.0000161	-0.59
ln(Savings)	.0000162	.0000168	0.96
Mortgage	1.52e-06	.0000221	0.07
Age	1.24e-06	1.98e-06	0.63
$F(5, 11663) = .99$			
	Response of $\hat{T}$ to APR*		
	$\hat{T}$		
	Coeff.	SE	t
Avg. APR	363.09	341.98	1.06
N	11663		

The upper section of this table tests whether four important observable characteristics, income, debt, FICO, and savings, are predicted by the price regime shifts. These insignificant results show that price changes did not cause any differential attrition across observable characteristics: while characteristics like income and FICO did vary over price regimes, this variation was not correlated with the price level. The lower section predicts individuals? maturity choices, $\hat{T}$, using all observable characteristics other than APR, and tests whether this variable is predicted by the price regime shifts. $\hat{T}$ is the term predicted using all observables *except* price.
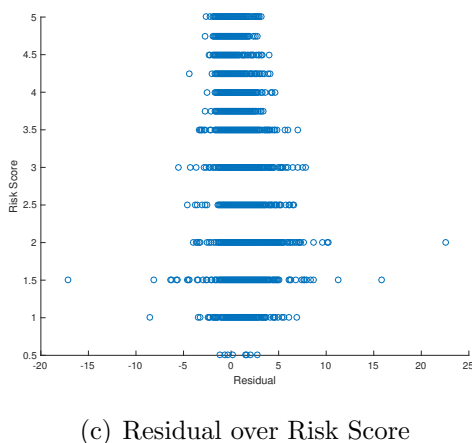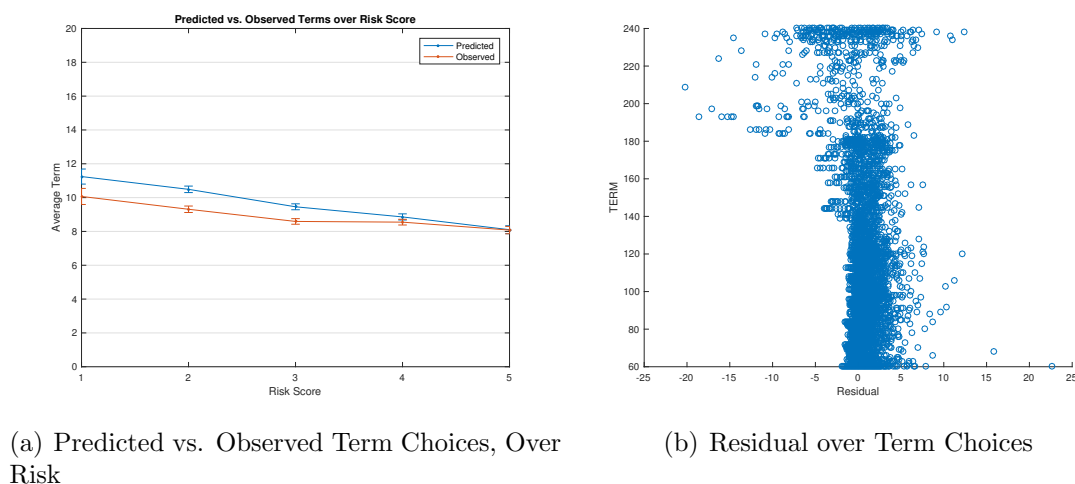

(a)   (b)

Figure 11: Comparison of Applicant Pool to Nationally Representative Sample of Graduate Student Borrowers

This figure compares the student loan amount and income quantiles of my applicant pool to those in a nationally representative sample of graduate student borrowers. The two populations look very similar, which suggests that the individuals I observe are similar to graduate students with Federal Loans.

(a) Across Risk Price Variation



(b) Within Risk Price Variation
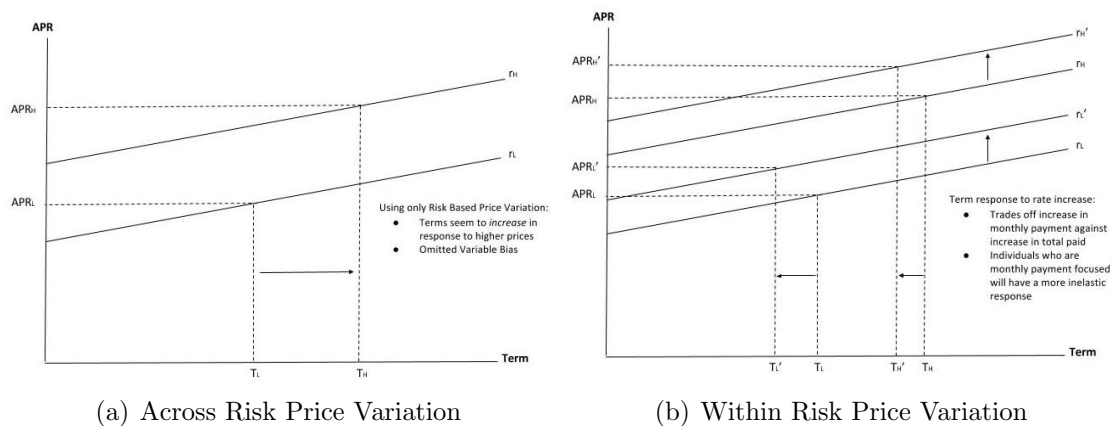
Figure 12: Using Across vs Within Risk Price Variation to Identify Term Elasticities

This figure shows the importance of only using within risk price variation to identify price elasticities. Panel (a) shows the term choices of two different risk types facing risk-based price variation – despite facing higher interest rates, the riskier type chooses a longer loan due to other omitted factors. This makes it seem as though term demand is increasing in interest rates. However, panel (b) shows that when prices increase *within* risk type, term demand actually decreases with interest rate.

(a) Predicted vs. Observed Term Choices, Over Risk

(b) Residual over Term Choices


(c) Residual over Risk Score

Figure 13: Model Fit

These figures analyze the model fit, comparing observed and predicted term choices, as well as the model residual over term and risk score. They show that in general the model slightly overpredicts terms, but otherwise seems to perform well. All counterfactual exercises use these predicted term choices as a comparison point.
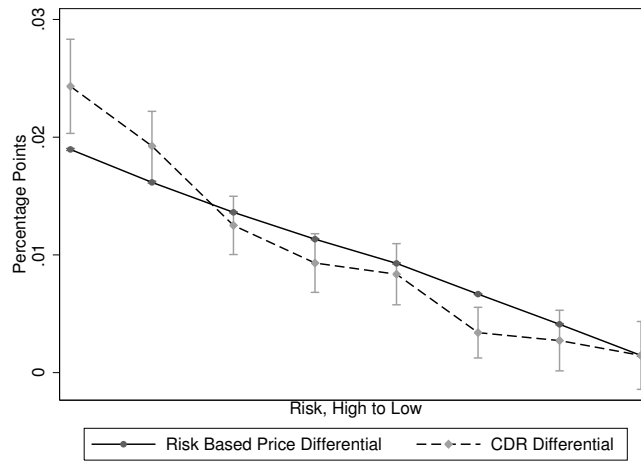
Figure 14: Cost Differential Relative to Lowest Risk Rating

The CDR is calculated and published by the Federal government at the school level, and reflects the student loan default rate of a cohort of students from that school after 3 years of completion. It is a much cited measure of expected costs used by Federal loan program. This figure compares the difference in the CDR between the highest and lowest risk types in my sample (which is roughly 2 percentage points) to the spread in their risk-based interest rates, and shows that private sector risk scores are highly correlated with the CDR.
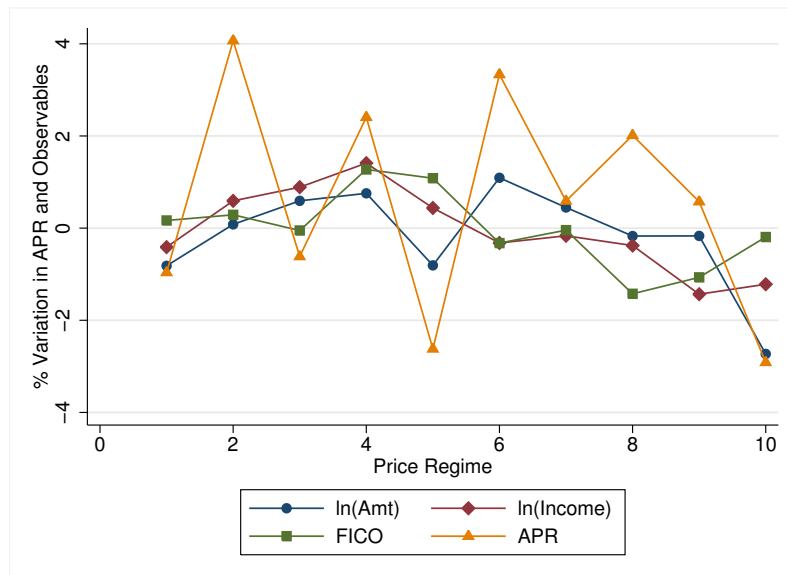

*Figure 15: Variation in Observable Characteristics over Time*

This graph shows changes in three important observable characteristics, income, debt amount, and FICO score, over 10 price regimes. While there are differences across price regimes, it is comforting to note that there are no obvious monotonic trends in these three variables and that they are not correlated with the exogenous price shifts.
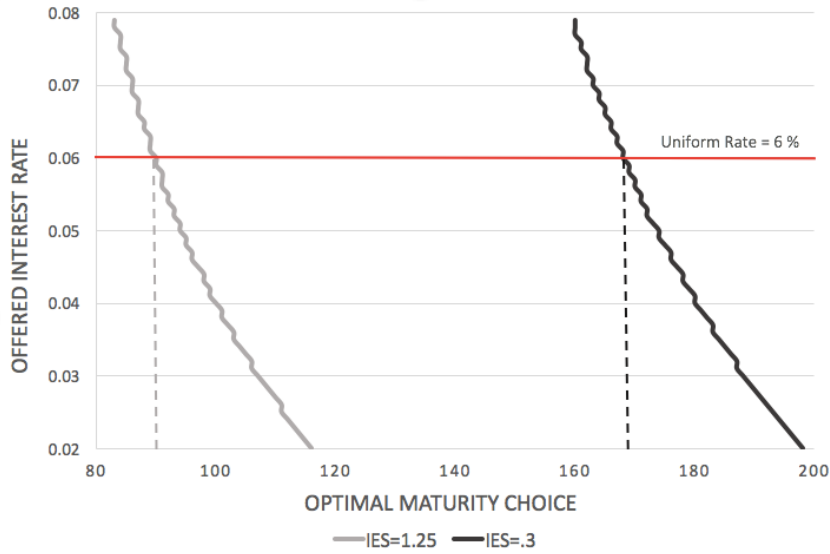
Figure 16: Simulated Relationship between Optimal Term Choice and offered Interest Rate
This figure shows how the simulated optimal term choice varies with interest rate levels for two values of the IES–
as the level of interest rates increases, the optimal maturity choice decreases. For any given interest rate,
individuals with a higher intertemporal elasticity of substitution will choose shorter optimal maturity. The graph
shows, for example, the optimal maturity choices of two otherwise identical individuals when facing a uniform
interest rate of 6% are almost 70 months apart. The simulation is for individuals with a loan amount of $90,000
and annual income of $50,000.
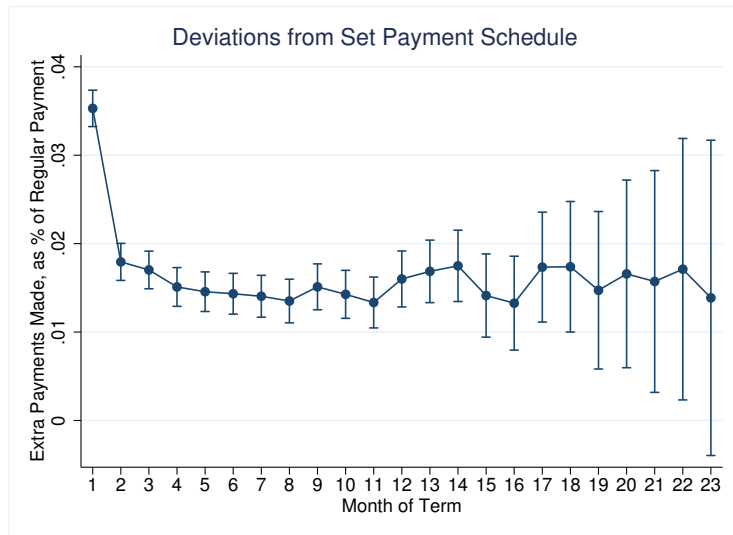
Figure 17: Average Size of Extra Payments Over Time Made by Borrowers
I look at payment patterns over time within my sample of refinancers – in other words, do any individuals change
their payment level over time permanently, or do they systematically make higher or lower payments on their debt.
I find that there are some extra payments in the data, but they are small and do not vary systematically over
time. This supports our model's assumption that borrowers make a term choice in year 1 to maximize expected
utility over the life of the loan and are not in fact choosing a monthly payment to fit their *current* income level,
with the intent to refinance and change term yet again in the future when their income level changes.

66
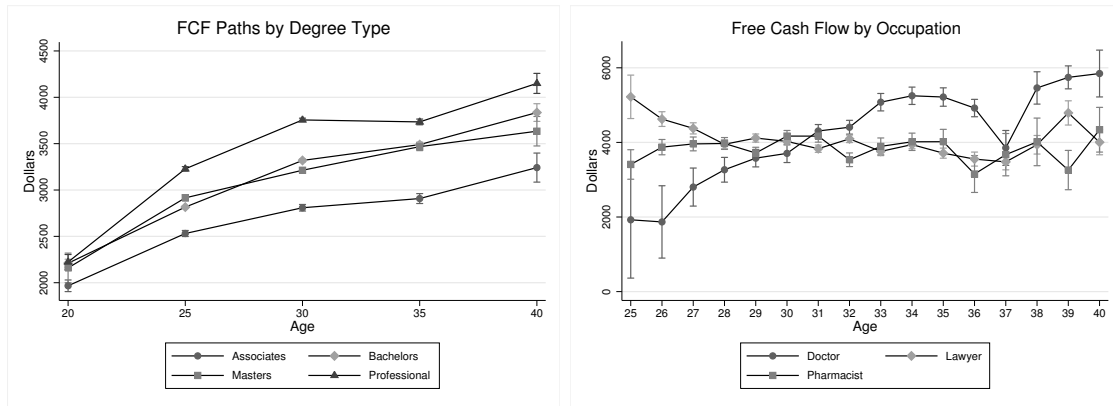
(a) Monthly FCF, by Age and Degree   (b) Monthly FCF, by Age and Occupation

Figure 18: Cross-Sectional Age Earnings Profiles, by Degree and Occupation

Here I plot free cash flow paths for individuals with different degree types and occupations. There are notable differences in both the level and the changes in free cash flow over the lifetime for each of these groups.

	Initial Monthly Payment			
	Initial Auto	*Initial Real Estate*	*Initial Non-Real Estate*	*Initial Credit Card*
Mean	449.32	1927.72	1104.31	94.37
Median	386.00	1698.00	882.00	52.00
IQR (25,75)	246.50	1210.00	897.00	84.00
	Change in Monthly Payment			
	Auto Change	*Real Estate Change*	*Non-Real Estate Change*	*Credit Card Change*
Mean	27.15	63.64	-181.06	-12.74
Median	0.00	0.00	-75.00	0.00
IQR (25,75)	0.00	37.00	490.00	52.00

Figure 19: Levels and Changes in Other Monthly Payments Before and After Refinancing

My model assumes that borrowers are not readjusting on other financial margins when refinancing. In other words, contemporaneous savings and debt decisions are assumed to be exogenous, predetermined, and unaffected by maturity and refinancing decisions. Here I test this assumption by looking at borrowers' other monthly payments before and after refinancing, and measuring whether they adjust immediately during refinancing. This table describes changes in other monthly payments (mortgages, auto loans, credit cards, etc) before vs. after refinancing for individuals who had positive monthly payments to begin with, and shows that for the vast majority of borrowers these stayed constant. This makes sense, since many of these payments are fixed installments, and it would take active work on the borrower's part to readjust.
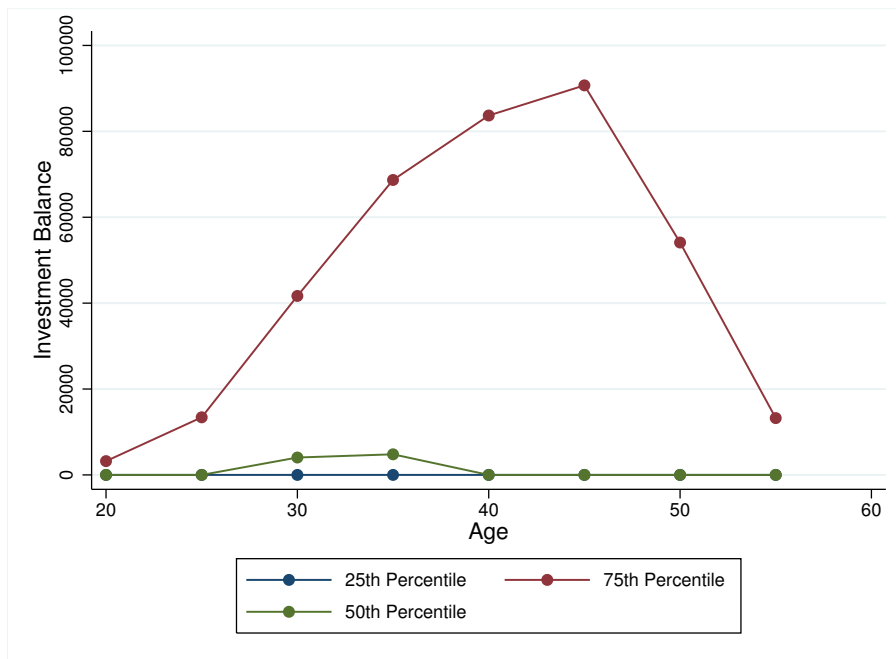
Figure 20: Investment Balances over Lifetime

My model defines yearly consumption as income minus the student debt payment; in reality individuals may also be making savings decisions that could impact their maturity choices. I can observe the savings and investment behavior of borrowers in my sample: because individuals in my sample are young, they have relatively low levels of savings to begin with. Slightly under 40% have a formal retirement savings account – for example 25% have a 401k, with a median balance of $24,000. The number of individuals with investment holdings increases with age. This figure shows that while the median borrower continues to not have substantial savings through age 60, the 75th percentile has accumulated over $80,000 by age 50. However, 90% of my borrowers are under 40 years old, and therefore even the most active savers have investment holdings that are much smaller than their student debt amount.