Opportunity Unraveled: Private Information and the Missing Markets for Financing Human Capital*

Daniel Herbst†
University of Arizona

Nathaniel Hendren‡
Harvard University

August 29, 2021

Abstract

Investing in college carries high returns, but comes with considerable risk. Financial products like equity contracts can mitigate this risk, yet college is typically financed through non-dischargeable, government-backed student loans. This paper argues that adverse selection has unraveled private markets for college-financing contracts that mitigate risk. We use survey data on students' expected post-college outcomes to estimate their knowledge about future outcomes, and we translate these estimates into their implication for adverse selection of equity contracts and several state-contingent debt contracts. We find students hold significant private knowledge of their future earnings, academic persistence, employment, and loan repayment likelihood, beyond what is captured by observable characteristics. For example, our empirical results imply that a typical college-goer must expect to pay back $1.64 in present value for every $1 of equity financing to cover the financier's costs of covering those who would adversely select their contract. We estimate that college-goers are not willing to accept these terms so that private markets unravel. Nonetheless, our framework quantifies significant welfare gains from government subsidies that would open up these missing markets and partially insure college-going risks.

*This draft is preliminary and comments are welcome. We are grateful for the thoughtful comments from Zach Bleemer, Raj Chetty, Erzo Luttmer, Miguel Palacios, Constantine Yannelis, and seminar participants at the University of Arizona, the MIT Golub Center for Finance and Policy, the University of Bristol, ZEW Mannheim, and the International Institute of Public Finance. We acknowledge financial support from the National Science Foundation (Hendren) and the MIT Golub Center for Finance and Policy (Herbst).

†dherbst@arizona.edu

‡nhendren@fas.harvard.edu

1 Introduction

Investing in college carries persistently high returns to both individuals and society, but also carries a huge risk. Nearly half of all college enrollees fail to complete their degree. Conditional on completion, only 85% find work after graduation. Even by age 40, over 15% of college graduates have household incomes below $40,000 a year. Meanwhile, the government's financing options have done little to mitigate this risk. Roughly one in five borrowers default on their student loans within the first five years of repayment, most of whom never graduated college and cannot discharge their debt (Looney and Yannelis, 2015).[1]

Economists have long advocated for the development of education financing contracts that mitigate risk (Chapman, 2006; Barr et al., 2017; Palacios, 2004; Zingales, 2012). Most famously, Friedman (1955) writes:

> "[Human capital] investment necessarily involves much risk. The device adopted to meet the corresponding problem for other risky investments is equity investment...The counterpart for education would be to 'buy' a share in an individual's earnings prospects; to advance him the funds needed to finance his training on condition that he agree to pay the lender a specified fraction of his future earnings."

A handful of private companies and post-secondary institutions have attempted to put this theory into practice with state-contingent or equity-like contracts for college.[2] Yet despite persistent attempts by private firms, decades of academic advocacy, and increasing returns to education over time, there is no active private market for equity or state-contingent college financing. Instead, federally-backed debt remains the dominant form of financing higher education in the US.[3]

[1] Employment and completion statistics are calculated six years from enrollment using the Beginning Postsecondary Students (BPS) study, a representative sample of first-time college enrollees in 2012. The fraction finding employment after completion corresponds to the percent of graduates who find a job within six years of their initial enrollment in college (excluding those still enrolled after six years). Default rates are taken from Looney and Yannelis (2015), who find five-year default rates ranging from 18 to 28 percent for post-2003 cohorts. Household income among forty-year-old college graduates are calculated using Current Population Survey (CPS) from 2011 onward.

[2] Yale, Purdue, and the University of Utah have each launched "Income-Share Agreements" (ISAs) to qualifying students with varying success, along with several smaller colleges and trade schools (Ladine, 2001; Hartley, 2016; University of Utah, 2021; Cowen, 2019). Early attempts at ISAs from private corporations like My Rich Uncle ended in bankruptcy (Rudegeair, 2016), while present-day providers, like Lumni, Vemo, and Better Future Forward, have collectively sold at most a few thousand contracts to targeted populations (Berman, 2017; Kreighbaum, 2019).

[3] As of July 2021, over 40 million borrowers hold a total of $1.6 trillion in outstanding student debt (Department

What explains this absence of risk-abating alternatives to student loans? One possibility is that college-goers simply don't value the insurance provided by such contracts. An alternative possibility is that such contracts are valued, but cannot be profitably sold because the market has unraveled due to adverse selection. Distinguishing between these explanations is critical for determining whether and how the government should intervene in financial markets for higher education. This paper quantifies the extent to which adverse selection prevents private financial markets from offering risk-mitigating contracts for financing human capital investments and explores the welfare implications of government subsidies to these types of contracts.

We begin by considering a model of state-contingent financial contracts under private information. In the spirit of Akerlof (1970) and Einav et al. (2010), the model shows that market existence depends on two curves: (1) the "willingness to accept" (WTA) curve, which corresponds to the minimum amount a given individual is willing to accept today to forego their future outcome, and (2) the "average value" (AV) curve, which corresponds to the average outcome among those willing to accept less than a given individual for the contract. If the AV curve lies below the WTA curve for all individuals, the market will completely unravel: If the financier lowers their offer to the college-goer in an attempt to make a profit, it leads those with relatively better expected outcomes to exit the market. Lowering the offer can in turn lower the value of the contract to the financier. If the AV curve lies below the WTA curve, this process will continue until no contracts are bought or sold. Along the way, the model also illustrates why moral hazard responses do not in general cause market non-existence: The first dollar of a financial contract induces only a small behavioral response, which has a second order effect on profits in contrast to adverse selection's first-order effect.

The goal of our empirical analysis is to evaluate the model's unraveling condition in four hypothetical contract markets by estimating AV and WTA curves in a representative sample of first-year college students. We begin with a market for "earnings equity" contracts, in which financiers can buy "a share in the individual's earnings prospects," as envisioned by Friedman (1955). In addition, we also examine markets for three types of state-contingent debt contracts. First, we consider a

of Education, 2021). The financial hardships faced by these borrowers has prompted calls among policymakers for a sweeping cancellation of student-loan balances (Warren, 2020; Harrison, 2021).

"completion-contingent loan" market, in which borrowers who complete their degree repay more than those who drop out of college. Second, we consider a market for "employment-contingent" loans, which are forgiven if the borrower fails to find employment.[4] Third, we consider a limited-liability loan market in which delinquent debt can be fully or partially discharged. Examining these contracts, which we refer to as "dischargeable loans," can help determine whether private debt markets could profitably replace the current system of subsidized and fully enforceable federal student loans.[5]

Quantifying the threat of adverse selection is difficult in these settings because unraveled contracts are not readily observed. To solve this problem, we build on an approach in Hendren (2013, 2017) that uses subjective probability elicitations as noisy and potentially biased measures of beliefs about future outcomes. We use panel data from the Beginning Postsecondary Students study (BPS). In 2012, this survey asked over 20,000 first-year college students about a range of future expectations, including their likelihood of graduating, expected occupation, and expected salary. The BPS then gathered follow-up data on these students' academic, financial, and employment outcomes five years later. Importantly, this follow-up information includes realized outcomes corresponding to our hypothetical contract markets: college completion, employment, loan repayment, and salary. Finally, the data contain a rich battery of student-level demographic and college information. The inclusion of this *public* information allows us to control for observable characteristics and assess whether a financier could potentially price contracts using different sets of observables to mitigate adverse selection.

Our empirical approach to evaluating whether these markets have unraveled proceeds in four steps that rely on an increasing set of assumptions. First, we establish the presence of private information. We find that individuals' subjective elicitations hold significant predictive power for each of the four outcomes, even conditional on a wide set of observable characteristics. Under relatively weak assumptions, this finding means that individuals have information about their pay-

[4]This focus on college students prior to labor market entry contrasts with Hendren (2017) who conditions on people who are already employed.

[5]Unlike most private consumer debt, existing federal student loans are subsidized and rarely dischargeable. Even in bankruptcy, student borrowers are still liable for defaulted debt. Our dischargeable-loan exercise uses borrowers' private knowledge of repayment risk on their existing federal loans to simulate a private market for unsubsidized student debt that cannot always be collected.

offs from our hypothetical contracts beyond what is captured by their observable characteristics, demonstrating that these markets face the threat of adverse selection.

Second, we provide a lower bound on the size of the threat of adverse selection. In particular, we show that the distribution of predicted values of the outcome given elicitations and observables can generate a non-parametric lower bound on the difference between an actuarially fair contract and the average level of the AV curve. Using a machine-learning approach to estimate predicted values, we find significant frictions imposed by private information. For example, we find that the average individual expects to earn \$2,000 to \$4,000 more per year six years after enrollment than their observationally-identical peers who expect to earn less than them. Relative to a mean income of \$24,032, the average individual would have to accept at least a 10% loss in expected future income in order for the financier to earn a profit on their equity contract.Importantly, this bound places no assumptions on the nature of elicitation error in subjective elicitations about future outcomes.

Third, we develop a new approach to estimate the distributions of private beliefs about future outcomes and use this to generate a point estimate for the AV curve. Our strategy is similar to Hendren (2013, 2017), but expands upon them in two ways. First, we allow belief elicitations that do not correspond directly to the outcome. For example, we use measures of parental encouragement to stay in college and the strength of parental financial support to estimate the belief distribution about the likelihood of default. Second, we extend the method in Hendren (2013, 2017) beyond contracts over binary outcomes, like employment, to contracts over continuous outcomes, like income. To do so, we make novel use of the deconvolution methods in Bonhomme and Robin (2010). We show that one can flexibly estimate the true latent distribution of beliefs using information contained in two or more indirect elicitations, so long as those elicitations are only correlated through their relationship to true beliefs about the outcome.[6] The deconvolution enables us to non-parametrically decompose the observed ex-post variation in salaries into the component that is known to individuals when entering college and the component that is realized as individuals go through college and enter the labor market. The former generates adverse selection; the latter generates a risk premium that

[6]We take advantage of categorical questions to form novel instruments that are likely to satisfy this exclusion restriction. For example, individuals are asked both their expected future occupation and their salary in that occupation. To measure beliefs about future salary, we use both the expected future salary and a predicted future salary based on the average salary of individuals in the expected occupation.

students are willing to pay for an equity contract.

Across all four markets, we quantify large amounts of private information that would generate significant adverse selection. For the equity market, the median individual would need to be willing to accept a loss of 39 percent in expected future income for the firm to break even on their contract. This means that for every dollar of financing, they would need to pay back on average $1.64 in present value in order to cover the financier's cost from adverse selection. For completion-contingent debt contracts, we find that the median individual would need to expect to pay back at least $1.76 in present value for every dollar of financing they receive in college. For debt that is forgiven in unemployment, the median individual would need to be willing to expect to repay $1.19 in present value for every dollar of upfront financing. Finally, dischargeable loan contracts, which limit borrowers' liability in periods of delinquency, also show significant potential for market unraveling. In the absence of government subsidies, the median individual would need to accept a contract in which they repay $1.87 in present value for every dollar of upfront financing.

While our point estimates suggest considerable potential for adverse selection among borrowers with private knowledge of future outcomes, a large portion of the realized variation in these outcomes reflects ex-ante uncertainty, not privately-known heterogeneity. For example, our deconvolution exercise finds that roughly half of the variation in income after college is known to individuals at the time of enrolling in college—the other half reflects the realization of risk. This suggests that risk-averse enrollees would be willing to accept less than the actuarially fair value of their future outcomes to obtain financing for college. But is this willingness to accept low enough to prevent the market from unraveling?

In our fourth step, we use estimates of college-goers' ex-ante uncertainty to quantify the WTA curve and compare it against our AV curve estimates. Following the literature in optimal social insurance, we measure this reservation price using a calibrated coefficient of relative risk aversion combined with proxies for the difference in consumption across states of the world. Across all four markets, our estimates suggest that the WTA curve lies below the AV curve. No one is willing to accept an expected-income loss that is large enough to make private contract markets profitable.[7]

[7]Our baseline calibration assumes a coefficient of relative risk aversion of 2, but the WTA curve continues to fall

If the private market has unraveled, should the government create or subsidize risk-mitigating markets for human capital financing? Absent adverse selection, the fact that these markets do not exist would suggest college-goers do not value such contracts. But, our finding that individuals are willing to give up some of their expected outcome to mitigate the risk they face means that the market unraveling is leaving Pareto-improving exchanges on the table. Can the government facilitate these exchanges by subsidizing risk-mitigating financial contracts, or would it better serve college-goers by eliminating student debt all together (Warren, 2020; Harrison, 2021)? Should the costs of higher education be borne by high-income graduates, or by society at-large? In the final section of the paper, we compare the welfare impacts of these types of policies by translating our estimates into the implied marginal value of public funds (MVPF) for subsidizing equity contracts, state-contingent loans, and untargeted grants. The MVPF measures the benefits to individuals divided by the net cost to the government. Importantly, these costs include not only the subsidization of negative profits but also any fiscal externalities from behavioral responses that affect tax revenue. On the one hand, college financing opportunities may increase future earnings and thus tax revenue, effectively reducing the cost to the government. On the other hand, equity-like contracts may reduce future earnings (and tax revenue) because of their implicit tax on earnings. We show how one can draw on two estimates from existing literature to measure these components: (1) the effect of college financial assistance on future earnings and (2) the elasticity of taxable income with respect to the tax rate on earnings. These two forces go in opposite directions so that ex-ante the sign of the fiscal externality and more generally the desirability of risk-mitigating contracts is not obvious ex-ante.

Our baseline our results suggest significant value to moving from debt to risk-mitigating financial contracts. For the equity contract subsidy, we estimate a risk premium that is four times larger than the fiscal externality induced by the higher implicit tax rate.[8] This means that in general we find higher welfare impacts of government subsidy of equity contracts that require individuals to

below the AV curve for levels of risk aversion as high as 4. We also conduct a robustness analysis in Section 7.3 that allows for heterogeneity in risk aversion across the population and show that this has very minimal impact on our results.

[8]This calibration assumes a coefficient of relative risk aversion of 2 and an elasticity of taxable income of 0.3. The risk premium continues to exceed the fiscal externality even for an elasticity of taxable income of 1.

pay back some of their financing in adulthood as opposed to grants that do not require repayment. In particular, we estimate an MVPF of the equity contract of 1.86, but an MVPF of a grant of 1.17.[9] Asking more successful graduates to provide larger repayments delivers a higher welfare gain than free grants.

In summary, our paper makes two primary contributions. First, we provide the first empirical evidence that private information prevents financial markets for human capital investments of the type envisioned by Friedman (1955). Second, we quantify the welfare gains to government subsidies of risk-mitigating financial contracts contracts. Even though our results suggest these contracts are not profitable for private firms, our estimates lend support to policies that would expand government subsidies for equity contracts and other financial products that insure individuals against the risk of their human capital investments.

This paper relates to a large literature on higher education financing. Beginning with Friedman (1955), many researchers have documented the theoretical benefits of more "equity-like" instruments for human capital investments (Chapman, 2006; Barr et al., 2017). More recently, empirical work has focused on public income-driven repayment (IDR) plans for student loans, including their labor supply effects (Chapman and Leigh, 2009; Britton and Gruber, 2019; Abraham et al., 2018; Field, 2009), liquidity benefits (Herbst, 2020), and barriers to enrollment (Abraham et al., 2018; Cox et al., 2018; Mueller and Yannelis, 2019). Herbst (2020) and Karamcheva et al. (2020) document negative selection into IDR plans, showing that high-balance, low-income borrowers are more likely to opt into IDR. While these plans often carry debt-forgiveness provisions, their primary function is to transfer payments within individuals over time (Herbst, 2020). By contrast, we focus on risk-mitigating contracts aimed at pooling risk across college-goers.[10] Finally, our paper relates to Mumford (2020), who finds that participants in an income-share agreement at Purdue are more likely to major in lower-income fields and take lower-paying jobs after graduation.[11]

[9] Beyond equity contracts, we find a relatively high MVPF for contracts that require repayment only if the borrower finds employment (MVPF of 1.42), but we find a lower MVPF for contracts that expand limited liability of student loans (MVPF of 0.79).

[10] Many IDR plans incorporate some amount of risk-pooling through debt-forgiveness thresholds, which forgive outstanding balances after twenty or twenty-five years in repayment. Our results suggest expanding such provisions may carry large welfare gains, as they can serve as a government-subsidized substitute to unraveled earnings-equity contracts.

[11] Mumford (2020) interprets these patterns as the result of moral hazard, whereas our analysis suggests they may

Our paper is also related to studies documenting worker expectations about future outcomes. In particular, our results are consistent with earlier research that has documented that college-goers' subjective expectations about future incomes are (imperfectly) predictive of their future incomes (Conlon et al., 2018; Gong et al., 2019). Relative to this literature, our core contribution is to quantify the implications for the workings of private markets and to quantify the welfare implications of government intervention in these markets.

Our study also complements a literature in macroeconomics studying the implications of financial market incompleteness (see e.g. Brunnermeier and Sannikov (2013)) and optimal government policy towards human capital (Jacobs and van Wijnbergen (2007); Stantcheva (2017)). Relative to this literature, we provide a microeconomic foundation for the absence of private financial markets for human capital. We also illustrate how one can use a sufficient statistics approach to assess the desirability of government policies that subsidize financial contracts for human capital investment even when such contracts do not currently exist.

Finally, our paper contributes to a large empirical literature on information asymmetries and adverse selection in a variety of financial markets, including mortgages (Stroebel, 2016; Gupta and Hansman, 2019), auto loans (Adams, Einav and Levin, 2009; Einav, Jenkins and Levin, 2012), credit cards (Ausubel, 1999; Agarwal, Chomsisengphet and Liu, 2010), personal loans (Dobbie and Skiba, 2013; Karlan and Zinman, 2009), to name only a few.[12] Our empirical framework builds upon insights from Einav et al. (2010), in particular, who use price variation to estimate welfare implications of adverse selection in health insurance markets. We provide a new method to estimate similar sufficient statistics without observing prices of contracts; we also extend their normative framework using the MVPF approach to account for fiscal externalities from changes in earnings in response to the contract.

The rest of this paper proceeds as follows. We begin in Section 2 with a theoretical model of human capital contract markets and develop the no-trade condition under which such markets completely unravel. We then outline our data in Section 3, and we provide reduced form evidence

be due to adverse selection.

[12] See Einav et al. (2021) for a recent review of papers documenting and quantifying selection in credit and insurance markets.

of college-goers private information about future outcomes in Section 4. In Section 5, we estimate a lower bound on the magnitude of this private information—the average difference between an individual's own expected outcome and the average value of those who would adversely select their contract. In Section 6, we provide point estimates for the average value curve. In Section 7, we provide calibrated estimates of the willingness to accept curve and formally test the unraveling conditions. Section 8 discusses the welfare impact of government subsidies for risk-mitigating college financing products.

2 Model of Market Unraveling

We begin by developing a model of market unraveling of risk-mitigating contracts for financing human capital investment. We use the model to derive estimable statistics characterizing when the markets exist. Along the way, we also clarify why adverse selection plays a unique role, relative to moral hazard, in affecting market existence. We later use the model to analyze the welfare impact of government subsidies that would help open up these markets.

Consider a population of college-goers facing an uncertain future outcome, y. We consider a financial contract that provides a payment $\lambda\eta$ today in exchange for a repayment of ηy when y is realized. The size $\eta \in [0,1]$ measures the fraction of the future outcome, y, that the individual agrees to repay. The price $\lambda \geq 0$ represents the amount today that the individual can receive per unit of y that is pledged for repayment. The outcome, y, can be either continuous or discrete. For example, if y is post-college income, the contract corresponds to the equity share envisioned by Friedman (1955). Alternatively, setting y to be an indicator of college completion corresponds to a state-contingent debt contract that requires repayment only if one graduates.

We imagine that we have conditioned on observable characteristics so that individuals are observationally identical in the model. But, individuals may have private information about their own future outcomes. We index the population by θ and denote their beliefs about y by the c.d.f. $\mathcal{F}_\theta(y)$ and mean μ_θ. Analogously, let $F(y|\theta)$ denote the cross-sectional realized c.d.f. of y given θ. We assume the following:

Assumption 1 *Beliefs are unbiased:*

$$\mathcal{F}_\theta(y) = F\left(y|\theta\right).$$

This assumption means that if we hypothetically isolated all people with a privately known belief, θ, the realized cross-sectional distribution of y would correspond to their ex-ante distribution of beliefs about y. Assumption 1 is a notion of unbiasedness that is weaker than traditional notions of rational expectations, as we do not require that θ incorporates any notion of all available information in the economy.[13] This assumption also implies that $\mu_\theta = E\left[y|\theta\right]$, so that individuals' expected y corresponds to the average realized y for those with beliefs θ.[14]

In this environment, we ask when a contract (i.e. a pair of η and λ) can generate non-negative profits. Suppose that individuals are offered a small contract $d\eta$ at price λ. Let $u\left(c_1,c_2,a\right)$ denote individuals' utility over consumption in period 1 when financing is provided, c_1, consumption in period 2 when y is realized, c_2, and a vector of other actions a taken in periods 1 and 2 (e.g. choice of college major or how many hours to work after college). Utility maximization implies that a type θ will accept the small contract if and only if

$$\lambda u_1(\theta) - E\left(yu_2|\theta\right) \geq 0. \tag{1}$$

where $u_1 = \frac{\partial u}{\partial c_1}$ and $u_2 = \frac{\partial u}{\partial c_2}$ are the marginal utilities of income in period 1 and 2, respectively, evaluated in the status quo world with no contract ($\eta = 0$).[15] The first term reflects the benefit of receiving $\$\lambda$ in period 1, which is measured by the marginal utility of consumption of a type θ in period 1. The second term is the disutility from future repayment, which equals the expected marginal utility of $\$y$ of consumption of type θ in the second period. Note that the envelope theorem implies that equation (1) is not affected by the impact of the contract choice on behavior, a. Because

[13] Nor will we require individuals to be able to perfectly report their true beliefs on surveys. This approach contrasts with papers that treat elicitations as a direct measurement beliefs (Wiswall and Zafar, 2021; Arcidiacono et al., 2020).

[14] In Section 7.3 we discuss how biased beliefs would affect our results and demonstrate robustness to many such biases, particularly overconfidence.

[15] Equation (1) is implied by the envelope theorem. See Milgrom and Segal (2002) for a wide class of primitive assumptions that ensure the envelope theorem holds even if the choice of a is discrete (e.g. college major) or from non-convex constraints (e.g. extensive margin labor supply decisions).

we consider a small contract, $d\eta$, these marginal utilities are evaluated using allocations, (c_1,c_2,a), in the status-quo world in which $\eta = 0$. But, in this status quo world, we assume individuals can borrow through existing student loan markets at an interest rate R between periods 1 and 2. The Euler equation from the choice of traditional student loans implies $\frac{E[u_2|\theta]}{u_1(\theta)} = R$. We discuss the impact of rationing and credit constraints in the traditional student loan market in Section 7.3.

We define the *willingness to accept*, $WTA(\theta)$, for a type θ to be the price in period 2 present value that a type θ is willing to accept to give up a small portion of their future y,

$$WTA(\theta) = \frac{E[yu_2|\theta]}{E[u_2|\theta]} \tag{2}$$

All types θ for whom $WTA(\theta) \leq R\lambda$ will accept the contract. Without loss of generality, we order types $\theta \in [0,1]$ in ascending order according to their willingness to accept in the status quo world, $WTA(\theta)$.[16] This means that a particular value of θ corresponds to the fraction of the market that would accept the financial contract when the financier offers a price λ.

An individual's willingness to accept can be expressed as the product of two terms:

$$WTA(\theta) = E[y|\theta] + cov\left(y, \frac{u_2}{E[u_2|\theta]}|\theta\right). \tag{3}$$

Those with higher expected incomes naturally require a higher price in order to be willing to accept the contract. But, the $cov\left(y, \frac{u_2}{E[u_2|\theta]}|\theta\right)$ reflects the risk discount individuals are willing to accept due to the insurance value of the contract. The covariance term is negative if higher realizations of y correspond to lower marginal utility states of the world, so that individuals are willing to accept less than the actuarially fair value of their future incomes, $E[y|\theta]$.

Because a risk averse borrower's willingness-to-accept, $WTA(\theta)$, is higher than the amount they expect to pay, $E[y|\theta]$, a financier who could observe θ could profitably sell them a contract at any price λ such that $R\lambda \in [WTA(\theta), E[y|\theta])$. However, because types θ are unobservable to the financier, they cannot control which types purchase the contract. So instead of the *marginal*

[16] This is without loss of generality because we consider small contracts that themselves do not affect the willingness to accept.

value of contracting with a single type, $E[y|WTA(\theta) = R\lambda]$, they must consider the *average* value of all others who accept the contract, $E[y|WTA(\theta) \leq R\lambda]$. Their realized profits measured in the second period are therefore given by

$$\Pi(\lambda) = \Pr\{WTA(\theta) \leq R\lambda\}(E[y|WTA(\theta) \leq R\lambda] - R\lambda), \tag{4}$$

where $\Pr\{WTA(\theta) \leq R\lambda\}$ is the fraction of the market that purchases the contract and R is the interest rate faced by the financier. For those who purchase the contract, the financier pays λ in the first period and receives the average y of those who purchase in the second period, $E[y|WTA(\theta) \leq R\lambda]$. We assume in our baseline specifications that the interest rate faced by the financier is the same as the interest rate faced by individuals.[17] We analyze the impact of differential interest rates on our results in Section 7.3.

Notably absent from equation (4) are the changes in college-goer's behavior, a, as a result of obtaining the financing contract. On the one hand, providing additional financing might cause the individual to take actions that increase y (e.g. complete college); on the other hand, being forced to repay a fraction of y may cause individuals may take actions that decrease y (e.g. reduce labor supply—commonly called "moral hazard"). Importantly, however, these behavioral responses do not have first-order effects on the financier's profits for a small contract, $d\eta$.[18] This insight, first noted by Shavell et al. (1979), implies that behavioral responses like moral hazard can attenuate the gains to trade, but cannot provide a singular theoretical explanation for the absence of a market. By contrast, even a small "$d\eta$-amount" of state-contingent financing can be adversely selected by strictly worse risks, so that private information imposes a first order cost on a financier's profits.

To quantify these costs, let θ_λ denote the type that is indifferent to accepting the contract at price λ, $WTA(\theta_\lambda) = R\lambda$. We define the *average value* curve, $AV(\theta_\lambda)$ as the average value of y for

[17] On the one hand, one might expect the interest rate faced by firms to be higher because student loan interest rates are implicitly subsidized by the government. On the other hand, financiers may have better access to efficient financial markets. Assuming similar interest rates therefore forms a natural benchmark case.

[18] While these behavioral responses have only second-order effects on a private financier's profits, they may have first-order effects on government tax revenue (Hendren and Sprung-Keyser, 2020). These externalities will play an important role in the welfare analysis in Section 8.

those who are willing to accept the contract at price λ:

$$AV(\theta_\lambda) = E[y|WTA(\theta) \leq WTA(\theta_\lambda)] \tag{5}$$

In general, the AV curve depends on the joint distribution of the expected outcome of each type θ, $\mu_\theta = E[y|\theta]$, and their willingness to accept, $WTA(\theta)$. For our baseline results, we derive a simpler expression for the AV curve by abstracting from heterogeneity in willingness to accept conditional on the expected outcome:

Assumption 2 *$WTA(\theta) > WTA(\theta')$ if and only if $\mu_\theta > \mu_{\theta'}$.*

Under this assumption, the average outcome of those who purchase at price λ is equal to the average outcome of those who expect to have lower outcomes than the person who is indifferent to the contract: $E[y|WTA(\theta) \leq R\lambda] = E[y|\mu_\theta \leq \mu_{\theta_\lambda}]$, where θ_λ is the type that is indifferent to the contract, $WTA(\theta_\lambda) = \lambda$. This means that the average value curve is given by the expected y for those with weakly lower beliefs:

$$AV(\theta_\lambda) = E[y|\mu_\theta \leq \mu_{\theta_\lambda}]. \tag{6}$$

Assumption 2 allows the average value curve to be estimated solely using information on the distribution of beliefs, μ_θ. In Section 7.3, we relax Assumption 2 and conduct simulations that introduce varying amounts of preference heterogeneity. We show that even large amounts of preference heterogeneity does not meaningfully affect our empirical conclusions.

Given the definition of the average value curve, we can write $\Pi(\lambda)$ in equation (4) as

$$\Pi(\lambda) = \Pr\{WTA(\theta) \leq \lambda\}(AV(\theta_\lambda) - R\lambda) \tag{7}$$

Using the identity $WTA(\theta_\lambda) = R\lambda$, the market will not be profitable at any price λ if and only if

$$AV(\theta) < WTA(\theta) \ \forall\theta \tag{8}$$

Unless someone is willing to accept a price corresponding to the pooled outcomes of those with lower expected outcomes than themselves, the market will unravel. Inequality (8) characterizes when the financier can profitably sell a small contract, $\eta \approx 0$. In a wide class of models, the profits to the financier are concave in (λ,η) so that the marginal profits to the financier are declining in the size of the contract, η, making larger contracts unprofitable as well (Hendren (2017)).[19] In this sense, the average value and willingness to accept curves characterize when a market can exist.

Example Figure 1 provides an illustrative example for the equity market case when y is post-college salary. The vertical axis presents the $AV(\theta)$, $WTA(\theta)$, and $E[y|\theta]$ curves as a function of the fraction of θ on the horizontal axis. We assume individuals' privately expected post-college salaries are uniformly distributed between \$20,000 and \$80,000. The blue line plots these quantiles of $E[y|\theta]$, which is linear because of the uniform distribution. The red line below $E[y|\theta]$ plots the $WTA(\theta)$ curve. For market existence, the question is whether firms can make a profit by offering contracts at price λ.

Figure 1A depicts a scenario where the financier can make a profit, and Figure 1B depicts a scenario where the market unravels. In Figure 1A, individuals are willing to accept less than the \$35,000 necessary for a market to be profitable when $\theta = 0.5$. But, in Figure 1B, the financier must set a price of \$40,000 in order to have 50% of the market accept the contract, so that they would lose \$5,000 per person who accepts. The financier could lower their offer to \$35,000 to attempt to break even, but doing so would further shrink the fraction of the market accepting the contract, rendering that contract unprofitable as well. Because no one is willing to accept the average value of the worse risks than them, the market unravels.

3 Data and Summary Statistics

Our goal is to empirically evaluate the unraveling condition given by (8). The data we use come from the 2012/2017 Beginning Postsecondary Students (BPS) longitudinal study, a dataset from

[19] See Hendren (2013) for a discussion of why equation (8) also characterizes market existence when financiers can use menus of contracts instead of a single contract, (η,λ).

the National Center for Education Statistics (NCES). The BPS data consist of administrative student loan and financial aid records linked to survey responses for a nationally representative sample of entering first-time college students in 2012, followed up in 2014 and 2017. The BPS data include three types of variables that are critical to our strategy. First, the panel aspect of the survey includes ex-post *realized outcomes, y,* corresponding to our hypothetical contracts—earnings, degree completion, employment status, and loan-repayment status. Second, the 2012 survey includes *private survey responses* related to individuals' future outcomes (e.g., subjective expectations of post-college earnings), which we use as noisy and potentially biased measures of individuals' true beliefs about future outcomes.[20] Third, the survey includes a rich body of *public information* from linked survey and administrative data, which we use to simulate the observable information upon which hypothetical financiers might set contract terms. Summary statistics are provided for key outcomes and elicitations in Table 1, and for public information in Table 2.

Outcomes, y, for the Four Hypothetical Markets We use the BPS data to assess whether the market has unraveled for four types of financial contracts whose corresponding outcomes, y, are observed in the 2017 BPS data. First, we consider an equity market in which individuals repay a fraction of their annual post-college earnings in 2017, $y = $*2017 Salary*. Figure 2A reports the distribution of post-college salary in 2017.[21] The average salary six years after enrollment is \$24,032, with a standard deviation of \$25,376.[22] Over 40 percent of those with positive earnings report annual salaries less than \$25,000.

We also consider three state-contingent debt markets. First, Figure 2B shows that in 2017, only 52 percent of 2011 enrollees have completed their degree; we consider a market in which borrowers only repay in the event of completion, so that y is defined to be an indicator of completion, $y = 1\{Completion\}$. Second, only 73 percent of 2011 college enrollees are employed in 2017; we

[20] We detail below how we do not require these elicitations to correspond to "true" beliefs entering individuals' maximization programs in Section 2.

[21] Respondents could report earnings in annual, monthly, weekly, or hourly amounts. To construct annual salary, the BPS included annual amounts as reported, multiplied monthly amounts by 12, multiplied weekly amounts by 52, and multiplied hourly amounts by 52 times the number of hours the respondent reported working at that job per week.

[22] Employment and salary outcomes are excluded for the 22 percent of the sample still seeking a degree.

consider a market in which borrowers only repay if they are employed, $y = 1\{Employed\}$.

Finally, we consider a market for dischargeable debt. Figure 2C shows a pie chart of the most severe loan status among student borrowers between 2012 and 2017.[23] By 2017, over two-thirds of student borrowers had experienced at least one delinquency since leaving college.[24] A full 18.4% of borrowers have already defaulted on their debt, indicating they have made no payments on their student loan for at least 270 days. The average earnings of those currently in default in 2017 is $25,000, while the average debt burden is $9,500. This debt overhang represents a real risk to borrowers, as defaulted student debt cannot be discharged in bankruptcy and often carries severe penalties. We therefore consider a market for debt that is only repaid in the event one avoids delinquencies on their current body of student loans, $y = 1\{Not\ delinquent\}$. This last contract can be thought of as debt that is dischargeable in the event of financial distress, as proxied by delinquency on existing student debt.

Private Information, Z Our empirical approach will use private information contained in survey responses to identify the distribution of individuals' beliefs about future outcomes and eventually estimate the WTA and AV curves. The principal source of this private information, Z, is subjective elicitations concerning expected outcomes.[25] Table 1, panel B reports the summary statistics for these subjective elicitations. In 2012, the survey asked new college enrollees a battery of subjective elicitations concerning uncertain outcomes, including their likelihood of degree completion, expected post-college occupation, expected salary in that occupation, and their expected salary if they did not go to college. It also asks a set of difficult-to-publicly-verify variables such as the extent to which their parents support them through encouragement and financial support—we also utilize these variables as subjective information that a financier would not be able to ascertain.[26]

The responses provide suggestive evidence that they perceive significant uncertainty about their post-college experiences. Appendix Figure A1 shows the histogram of reported likelihoods of com-

[23] We exclude those who are still enrolled in a degree program and therefore do not require repayment.

[24] A student loan is considered delinquent as soon as the borrower misses a payment, though loan servicers will often only record delinquencies if payments are not received within a week or two.

[25] Throughout this paper, we use upper-case letters (Z,X,Y) to denote vector-valued objects and lower-case letters (z,x,y) to denote single-dimensional objects.

[26] Appendix B provides the survey question text for each of the subjective elicitations used in our analysis.

pleting college. The average belief is 93 percent, with a standard deviation of 18 percent—a significant fraction of the population enters college uncertain about whether they will complete their degree.[27] Individuals are even more uncertain about completing on time: Table 1 panel A shows that the average belief about completing on time is 84 percent, with a standard deviation of 20 percent. However, going forward we do not assume that these subjective responses, Z, correspond to the "true" subjective expectations that individuals' utility maximization program in Section 2. Rather, our approach will treat these elicitations as noisy and potentially biased measures of true beliefs.

Observable Information, X A key benefit of the BPS is that it is linked to FAFSA records, administrative high school and college records, administrative loan data, and a battery of survey data. Table 1 reports the set of observable variables we consider in our analyses, and Table 2 reports the summary statistics for the non-categorical outcomes (Appendix Table A1 lists all variables). We classify these observables into five groups: (1) institutional characteristics, which includes enrollment size of the institution, admission rate, tuition, degree offerings, urban versus rural location, demographic compositions, and test scores of the entering class; (2) academic program characteristics, which includes the degree type (BA vs AA), field of study, and the college-goer's age at enrollment; (3) high school performance measures, which include high school GPA and SAT/ACT scores; (4) demographic information, which includes citizenship status, marital status, number of children, state of residence prior to enrollment; and (5) parental characteristics, including annual income, expected family contribution (EFC) from the FAFSA, number of children, and marital status.[28] This large body of information allows us to assess whether markets would unravel even if financiers could use this information to price contracts.

[27] Nonetheless, enrollees report an expectation of high returns to college: the average expected salary is $64,000, in contrast to the average expected salary if they had not completed college is $17,400.

[28] Categorical variables are simplified to binary indicators in Table 1 (e.g., STEM indicator in lieu of field of study). We also report summary statistics for race and gender, however we do not include these variables in our demographics controls because they are protected classes and cannot be used in pricing or screening for financial products. However, in Section 5 we show their inclusion does not significantly affect our results. A full description of each category's variables is provided in Appendix Table A1.

4 Identifying Private Information About Risk

Satisfying the market unraveling condition (8) requires the existence of private information about outcomes, y, among observationally equivalent borrowers. We begin our empirical analysis by assessing whether individuals possess such private information. To test for this, we retain our unbiased beliefs assumption (Assumption 1) from Section 2.[29] We also follow Hendren (2013) and impose an additional assumption that the information contained in elicitations Z can contain *no more information* about y than knowledge of one's own type, θ:

Assumption 3 *Z contains no more information than θ about y,*

$$F\left(y|\theta,X,Z\right) = F\left(y|\theta,X\right). \tag{9}$$

This assumption is relatively weak, as it is hard to imagine how one could say something more than what they know on a survey. The implication of this assumption is that if two people with different elicitations have, on average, different average outcomes, then individuals have private information. Formally, Assumptions 1 and 3 imply

$$E\left[E\left[y|X,\theta\right]|X,Z\right] = E\left[y|X,Z\right] \tag{10}$$

so that the distribution of true beliefs about y, $E\left[y|X,\theta\right]$, is a mean-preserving spread of the distribution of predicted values, $E\left[y|X,Z\right]$.[30] If Z is predictive of y, then so are true beliefs, θ. This motivates a simple test for the presence of private information: regress y on Z controlling for observables, X:

$$y_i = \Gamma Z_i + \beta X_i + \epsilon_i. \tag{11}$$

Rejecting the null hypothesis that $\Gamma = 0$ implies the existence of private information under Assumption 3. Importantly, this test does not require the elicitation, Z, to be directly concerned

[29] In our empirical context, this assumption implies that if we could hypothetically take a group of individuals with same beliefs, their realized distribution of outcomes (of y conditional on θ) would correspond to a type θ's ex-ante beliefs about y.

[30] This result is shown in Hendren (2013) for the case when y is binary.

with the outcome, y. While closely related survey questions are likely to elicit responses with more predictive content about y than questions concerning a completely different topic, a close relation between Z and y is not formally necessary—the predictive content in the elicitations provides a one-sided test for the presence of private information.

We begin this analysis in Figure 3 with a simple binned scatter plot of the univariate relationship several survey responses and each of our four primary outcomes of interest. In Panel A, we plot employed individuals' log realized salary in 2017 against the log of their subjective "expected future salary" measured in 2012. Those who report higher expected salaries in 2012 on average have higher salaries in 2017. In Table 3, we explore how the inclusion of control variables affects this relationship. Each column reports the slope coefficient for an increasing set of control variables corresponding to categories outlined in Section 3.[31] The first column includes no controls; the second column includes controls for institutional characteristics of the college in which the student is enrolled; the third column adds controls for the students' academic program in the college, the fourth column adds controls for the student's high school GPA and SAT scores, the fifth column adds controls for student demographic characteristics, and the sixth column adds controls for parental information.

The table shows that predictive content of the elicitation remains conditional on a range of observables used as public information. Controlling for institutional characteristics, academic program characteristics, and high school performance characteristics leads to a slope of 0.043 (s.e. 0.0160). Further adding controls for all of our observables leads to a slope of 0.031 (s.e. 0.0158).

Individuals also have private information about their likelihood of degree completion. Figure 3B displays the relationship between six-year graduation rate and respondents' stated likelihood in 2012 of completing their degree "on-time." Those who say they are more likely to complete on time are indeed more likely to do so. Table 4 shows how this slope changes with the inclusion of controls and illustrates that the predictive content remains: including all of our sets of public information controls yields a slope of 0.033 (s.e. 0.0022).

Next, we explore the presence of knowledge about future employment. Individuals are not directly asked about their likelihood of being employed. But, they are asked about other expectations

[31] Appendix Table A1 provides a full list of variables corresponding to each control category.

about their labor market opportunities - in particular, individuals are asked about their expected salary if they were not going to college. In Figure 3C, we show that the likelihood that students are employed in 2017 is increasing in this expected salary if they were not attending college. The predictive content in this elicitation is of course likely to be smaller than what would be obtained if individuals had been asked about their likelihood of being employed; but nonetheless our approach can still identify the presence of private information. In Table 5, we show that this predictive content remains after including controls for public information. The slope is 0.020 (s.e. 0.0107) using controls for institutional, academic program, and high school performance characteristics and 0.017 (s.e. 0.0106) when further including demographic and parental characteristics.

Finally, we explore college-goers knowledge about their likelihood of remaining current on their federal student loan obligations. Here again, individuals are not directly asked their likelihood of delinquency; instead, we exploit here the fact that individuals are asked about their parent's level of financial support for college. Figure 3D shows that student borrowers who report greater parental encouragement for college on a 1–5 scale are more likely to be current on their federal student loans (no delinquencies, defaults, or forbearances) through 2017. Table 6 shows that this pattern remains even after including our wide array of control variables, with a slope of 0.031 (s.e. 0.0049) when including institutional, academic program, and high school performance characteristics, and 0.029 (s.e. 0.0048) when further adding controls for demographics and parental characteristics. In summary, individuals have significant amounts of information about their future outcomes that would be difficult for financiers to obtain.

5 Lower Bounds on the Average Value Curve

How "much" private information do individuals have about future outcomes? In this section, we move from rejecting the null of no private information to providing a lower bound on the frictions imposed by adverse selection. We measure these frictions using the difference between the actuarially fair price for a type θ's expected outcomes, $E[y|\theta]$, and the average value of those with weakly

lower expected outcomes, $AV(\theta)$,

$$m(\theta) = E[y|\theta] - AV(\theta)$$

Following Hendren (2013), we define $m(\theta)$ as the *magnitude of private information.*

We do not observe θ and so cannot construct a direct estimate of $m(\theta)$. But, we can use the empirical expectation of outcomes, y, given observable characteristics, X, and private elicitations, Z, to construct a lower bound on the average magnitude of private information, $E[m(\theta)]$. We exploit the result in equation (10) that the distribution of true beliefs about y are a mean-preserving spread of the distribution of predicted values given the elicitations. For each individual, i, we define r_i as the difference between their predicted outcome conditional on both observable information and private elicitations, $E[y|X = X_i, Z = Z_i]$, and their predicted outcome given only observable information, $E[y|X = X_i, Z = Z_i]$:

$$r_i \equiv E[y|X = X_i, Z = Z_i] - E[y|X = X_i]. \tag{12}$$

The values of r_i deviate from zero to the extent to which the elicitations predict the outcome, y, conditional on observable characteristics. For each i, we calculate the average value of r_j among all individual's, j, with $r_j < r_i$:

$$m_i^Z \equiv E[-r|r < r_i]. \tag{13}$$

The value of m_i^Z provides the mean difference between privately expected outcomes and average value to insurers in a world where borrowers' private knowledge was limited to *only* the information contained in Z. Hendren (2013) shows that the population average of these individual estimates of m_i^Z form a lower-bound on the magnitude of private information:

$$E_\theta[m(\theta)] \geq E_i\left[m_i^Z\right]. \tag{14}$$

The left-hand side of equation 14 is the (unobserved) average difference between the outcome $E[y|\theta]$

and the average value curve, $AV(\theta)$, over all types, θ. The right-hand side is a lower bound that can be estimated using the distribution of predicted values of y given X and Z.

Importantly, note that we have not assumed anything about the measurement error embodied in the elicitations or the relationship between true beliefs and elicitations. But, the bound in equation (14) is of course tighter the greater the predictive content in Z. While the analyses in Tables 3 through 6 use two elicitations in each context, there are a large number of elicitations that could have been used to predict y, as well as any number of non-linear transformations of those variables. To address this, we estimate $E[y|X]$ and $E[y|X,Z]$ using the following machine-learning procedure. First, we use a random-forest algorithm to predict each outcome y from the set of public information, $\{X\}$. Then we apply the same algorithm to the set of both public and private information, $\{X,Z\}$.[32] Finally, we repeat this procedure for five different specifications of $\{X\}$: (1) a benchmark case with no public information ($E[y|X] = E[y]$), (2) allowing $\{X\}$ to include only institutional and academic characteristics, (3) adding performance and demographic characteristics, (4) adding parental background characteristics, and (5) adding race and gender.[33] Variables corresponding to each category of public information listed in Appendix Table A1. For Z variables, we include all elicitations used in Tables 3 through 6, plus any additional observable variables that are not included in the definition of public information. This means that we allow $\{X,Z\}$ to include all elicitations and observable information in our dataset so that $E[y|X,Z]$ does not vary across specifications.[34] Appendix Table A2 reports out-of-sample performance statistics from the random forest estimates of $E[y|X]$ and $E[y|X,Z]$.[35] Consistent with the results in Tables 3-6, we find that individuals' beliefs have strong predictive content beyond observable information in all four settings; the predictive metrics improve when adding subjective elicitations in column 4. For example, adding elicitations to predict salary increases the out-of-sample R–squared from .0766

[32] Details on our machine-learning procedure are provided in Appendix D.

[33] Race and gender are protected class, so they cannot be legally used in pricing or screening for financial products. We include them in our final specification for completeness.

[34] In other words, we allow private information, Z, to include not only elicitations data listed in Appendix B, but also any observable variables not included in the specified set of public information, X. While such variables can be plausibly designated as private information—individuals can observe their own SAT scores even if financiers cannot—limiting Z to elicitations-only generates qualitatively similar results. See Appendix Table A3.

[35] For binary outcomes, we calculate Pseudo R-squared, model accuracy, and area under the ROC curve. For salary, we calculate R-squared, root-mean-square error (RMSE), and mean absolute error (MAE). Standard errors for each statistic are calculated from 1,000 bootstrap replications of the holdout sample.

to 0.0966.

We use the out-of-sample random-forest predictions as our estimates of $E[y|X]$ and $E[y|X,Z]$ to form our empirical estimates of m_i^Z in equation for each individual with a given value of X_i and Z_i. Table 7 reports the estimates of the average, $E\left[m_i^Z\right]$, which forms the lower bound for $E[m_i]$. The first row considers the equity market case when y is salary, and the columns provide estimates for an increasing set of public information.[36] Without conditioning on observable characteristics, the average college-goer expects to earn at least \$5,765 more than their peers who expect to earn less than them. Conditioning on institutional and academic characteristics, this difference is reduced to \$5,314; it remains \$2,907 even conditional on observables such as parental characteristics that would likely be difficult for a financier to use to price the contract. Relative to a mean earnings of \$24,032, these results mean that the average individual would have to be willing to accept a valuation that is at least 10–24% lower than their expected future income in order for the financier to earn a profit.

The second row of Table 7 reports the estimates for the state-contingent contract in which individuals repay only if they complete college. Across the increasing controls for public information, we find that the average college-goer has a completion probability that is 11–20pp higher than those who are observationally identical but privately believe they are less likely to complete college. Compared to the mean completion of 51%, this means college-goers would have to be willing to accept a valuation that is at least 22–39% lower than their actuarially fair value for this market to exist.

The third row of Table 7 considers the state-contingent contract that requires repayment only if individuals are employed. We find that the average college-goer has a likelihood of employment that is 5–11pp higher than those who are observationally identical but privately believe they are less likely to find employment. Compared to a mean employment likelihood of 73%, this means the average college-goer would have to be willing to accept a valuation that is at least 6–15% lower than their actuarially fair valuation for this market to exist.

[36] Note that for the equity contract, equation (14) is written in terms predicted salary level, including the likelihood of being unemployed and earning zero. We transform predicted employment and predicted log earnings conditional on employment into predicted unconditional level earnings before we calculate r_i. Details are provided in Appendix D.

Finally, the fourth row of 7 considers the contract requiring repayment only in the state of the world that the individual is not delinquent on their student debt obligations. On average, college-goers with a greater expectation of timely repayments are 4–13pp less likely to fall delinquent than their observationally-identical peers. Compared to a mean delinquency rate of 69%, this means individuals would have to be willing to accept a valuation that is 4–42% lower than their actuarially fair price for this market to exist.

Together, the results provide evidence that college-goers' private information imposes a significant barrier to the workings of state-contingent college financing markets. Importantly, we arrive at these lower bounds without assumptions about the nature of the measurement error in the elicitations. But, these provide only lower bounds on the average deviation of the average value curve from individuals' own expected outcomes. In the next two sections, we turn to a method to provide point estimates for the AV and WTA curves.

6 Point Estimation of Average Value Curve

This section constructs point estimates of the AV curves. The key input into these curves is the cross-sectional distributions of beliefs about y, $\mu_\theta = E[y|\theta]$, conditional on observables, X. We develop a new approach to estimate these cross-sectional distributions, which makes use of recent advances in non-parametric deconvolution methods (Horowitz and Markatou, 1996; Li and Vuong, 1998; Bonhomme and Robin, 2010) and measurement error models (Hu and Schennach, 2008).

6.1 Estimating the Distribution of Beliefs about Expected Salary

We begin by illustrating our approach to estimating the distribution of $E[y|\theta]$ for the equity market case in which y is salary. It is helpful to split the distribution of these salaries into two components:

$$E[y|\theta] = \Pr\{y > 0|\theta\} E[y|\theta, y > 0]$$

where $\Pr\{y > 0|\theta\}$ is the probability that a type θ is employed and $E[y|\theta,y > 0]$ is their salary conditional on employment. We identify these two cross-sectional distributions separately, as this enables us to use a log parameterization for salary conditional on $y > 0$. Moreover, our method for identifying the distribution of $\Pr\{y > 0|\theta\}$ will be the same method we use for other binary outcomes.

Beliefs about Salary Conditional on Employment For salary conditional on employment, we follow a common specification in existing literature motivated by empirical work by Guvenen (2007) and others that uncertainty about future salary is realized as a proportional shock. In particular, we let $\tilde{\mu}_\theta$ denote the expected log salary, $\tilde{\mu}_\theta = E[log(y)|\theta]$. We consider a specification in which the realization of incomes, y_i, for each person in the data, i, follows the functional form:

$$log(y_i) = \tilde{\mu}_{\theta(i)} + \epsilon_i. \tag{15}$$

where $log(y_i)$ is the log of an individual's realized salary, $\tilde{\mu}_{\theta_i}$ is individual i's expected log income, $\tilde{\mu}_{\theta(i)} = E[log(y)|\theta = \theta(i)]$ and ϵ_i is the realization of an idiosyncratic shock. Note that because we have taken logs, recovering the distribution of $E[y|\theta,y > 0]$ will require us to estimate the marginal distributions of both $\tilde{\mu}_\theta$ and ϵ.

Next we turn to the elicitations. Recall that the BPS asks people their expected salary in two components, which will be useful for identifying the model. First, individuals are asked their expected tasks. These coded into occupation codes by the BPS, which we take at the 3-digit level. Next, people are asked conditional on working in that occupation, what is the salary they expect. We let z_i denote the log of individual i's reported expected salary. We assume these reports are related to true beliefs through an equation

$$z_i = \alpha + \gamma\tilde{\mu}_{\theta(i)} + \nu_i \tag{16}$$

where ν_i is an idiosyncratic error term. The specification in equation (16) allows the elicitations to be noisy and potentially biased measures of true beliefs. Moreover, the coefficient γ is allowed

to deviate from $\gamma = 1$, in contrast to the specifications in Hendren (2013; 2017). This allows the model to account for the fact that $\tilde{\mu}_\theta$ corresponds to salary in 2017, but in the survey response they might be thinking about salary in perhaps a later period of their lifetime.

To identify the coefficients α and γ, we use the information contained in the second elicitation, z'. The key identification condition that enables us to identify the distribution of beliefs is that the elicitation is correlated with z only through its correlation with beliefs about y.

Assumption 4 *There exists a second elicitation, z'_i, such that $z'_i \perp \nu_i$, so that z'_i is correlated with z_i only through beliefs about y.*

To plausibly satisfy Assumption 4, we take advantage of the sequential elicitation procedure on the survey and use an assumption that the measurement error in the categorical questions about expected occupation is orthogonal to the measurement error in salary conditional on employment within one's expected occupation. Specifically, for each individual i, we construct z'_i as the log average salary of college graduates (j) who had worked in individual i's expected occupation (occ_i) as of the 2012 freshman-year survey:

$$z'_i = log \frac{1}{N^{BB}_{occ_i}} \sum_{j \in occ_i} y^{BB}_j. \tag{17}$$

Post-graduate salaries (y^{BB}_j) and cell-sizes ($N^{BB}_{occ_i}$) are taken from the 2008/2012 Baccalaureate and Beyond (B&B) study, which we match to BPS occupation elicitations (occ_i) using three-digit occupation codes.[37] We then run an IV regression of z_i on $log(y_i)$, where we instrument $log(y_i)$ with z'_i, the log of average salaries in i's reported expected occupation. Table A4 reports the results for γ. We find an estimate of γ =0.69 (s.e. 0.16).

A potential concern with the empirical strategy is that the occupation-based salaries may be correlated with individuals subjective elicitation conditional on beliefs about future salary, so that the instrument would violate the exclusion restriction. We test for this potential concern by exploring the robustness of γ to alternative subjective information used as instruments. For example, in

[37] The B&B data include survey responses for a representative sample of four-year college graduates in the spring of 2008, followed up in 2011-2012. More information can be found at `https://nces.ed.gov/surveys/b&b/`.

Appendix Table A5, we the log of the self-reported salary individual i would expect had they not gone to college as the instrumental variable. This yields a similar value of γ =0.77 (s.e. 0.10).

With estimates of γ in hand, equations (15) and (16) form a latent factor model with two measurements (y_i and z_i) and three factors ($\tilde{\mu}_{\theta(i)}$, ϵ_i, and ν_i) (Aigner et al., 1984). The distributions of these factors can be estimated with minimal parametric assumptions using deconvolution methods (Bonhomme and Robin, 2010). Note, however, that a deconvolution of $log(y)$ and z would recover the *unconditional* distribution of beliefs, as equations (15) and (16) do not explicitly condition on observables, X. To construct estimates conditional on X, we residualize the left-hand sides of equations (15) and (16) using random-forest predictions of $E[y|X]$ from Section 5, using controls for institutional and academic characteristics.[38]

Using residualized values of $log(y)$ and z, we apply the deconvolution estimator from Bonhomme and Robin (2010) to recover the distributions of $\tilde{\mu}_{\theta(i)}$, ϵ_i, and ν_i conditional on public information. Appendix C provides details of how we implement the estimation procedure from Bonhomme and Robin (2010) for our estimation. Appendix Figure A2 plots the density estimates for $\tilde{\mu}_{\theta(i)}$ and ϵ_i. Importantly, our estimation method places no parametric assumptions on either distribution, which allows us to flexibly estimate the AV and WTA curves in the following section.

Having estimated the distributions of $\tilde{\mu}_{\theta(i)}$ and ϵ_i, we recover the distribution of $E[y|\theta,y>0]$ using the equation

$$E[y|\theta,y>0] = \int \exp(\tilde{\mu}_\theta + \epsilon) dF_\epsilon. \tag{18}$$

Appendix Figure A2C plots the resulting density for the beliefs about salary conditional on $y > 0$.

Beliefs about Employment Next we turn to the estimation of beliefs about employment. Let $e = 1\{y > 0\}$ denote an indicator for being employed and let $\omega_\theta \equiv \Pr\{y > 0|\theta\} = E[e|\theta]$ denote a type θ's expected probability of being employed. Because employment is a binary outcome, the identification result and deconvolution estimator in Bonhomme and Robin (2010) cannot be applied (a deconvolution of the distribution of a binary outcome into a continuous distribution

[38]We choose these controls because Tables 3-6 show that including additional controls does not generate a statistically significantly different coefficient on Z in any of our specifications. Moreover, existing attempts at income-contingent college financing have generally used a subset of these variables to set prices (Palacios, 2004).

of beliefs would violate the rank condition). However, we show here that one can use a flexible maximum likelihood estimator that is motivated by the non-parametric identification results in Hu and Schennach (2008). We use this approach below not just for the distribution of beliefs about employment but also for the outcomes corresponding to our other three state-contingent debt contracts.

To estimate the distribution of beliefs about employment, we take z_i to be the log of the self-reported salary individual i would expect had they not gone to college. We take the instrumental variable, z_i', to be the average employment rate among 2008 college graduates (j) in individual i's expected occupation, constructed analogously to the conditional salary case in equation (17):

$$z_i' = \frac{1}{N_{occ_i}^{BB}} \sum_{j \in occ_i} e_j^{BB}. \tag{19}$$

As in the salary case, we assume that we can write the elicitations as a linear function of the true beliefs of each individual,

$$z_i = \alpha + \gamma \omega_{\theta(i)} + \nu_i, \quad \nu_i \sim N\left(0,\sigma^2\right) \tag{20}$$

for some unknown variance, σ^2. We estimate γ as in the continuous case above: we regress z_i on an indicator for employment, $1\{y_i > 0\}$, and instrument for employment using z'. As above, the key identification assumption is that the measurement error is independent, so that z_i' is correlated with z_i only through its correlation with beliefs. Table A4 reports the results for γ. We find an estimate of γ =1.47 (s.e. 0.76).[39]

Next, consider the joint distribution of elicitations, z, and binary employment outcome e, $f_{e,z}(e,z)$. We can expand the observed density of e and z, $f_{e,z}(e,z)$, by conditioning on ω_θ:

$$f_{e,z}(e,z) = \int \omega_\theta^e (1-\omega_\theta)^{1-e} f_{z|\omega_\theta}(z|\omega_\theta)\, g(\omega_\theta)\, d\omega_\theta, \tag{21}$$

[39] Appendix Table A5 shows that these estimates are similar using an individual's self-reported likelihood of finding a job in their chosen occupation as an alternative instrument.

where $f_{y|\omega_\theta} = \omega_\theta^e (1-\omega_\theta)^{1-e}$ is the p.m.f. of e given ω_θ, $f_{z|\omega_\theta}$ is the distribution of the elicitations given ω_θ, and $g(\omega_\theta)$ is the distribution of beliefs about the likelihood of employment, $\omega_\theta = E[e|\theta] = \Pr[y > 0|\theta]$. Our estimates for α and γ and σ in equation (20) provide an estimate of $f_{z|\omega_\theta}$. The distribution of beliefs about employment, $g(\omega_\theta)$, can then be inferred from the joint distribution of y and z.[40] We flexibly specify the belief distribution, $g(\omega_\theta)$, as a grid of discrete point masses, so that it's c.d.f., $G(\omega_\theta)$, is given by

$$G(\omega_\theta) = \sum_j \xi_j \mathbf{1}\{\omega_\theta \leq a_j\}, \tag{22}$$

where $\{a_j\}$ is a set of twenty-five evenly-spaced point masses in [0,1]. Combining the flexible density function in (22) with the elicitation error distribution given by (20), we estimate $g(\omega_\theta)$ from the joint distribution of z and e by maximizing the likelihood given by equation 21.[41] Figure 4, panel C reports the distribution of the estimates ω_θ.

Constructing the Expected Salary Distribution Lastly, we combine our estimates for the distribution of expected salary conditional on employment and the distribution of beliefs about employment. To do so, we make a single index assumption that those with higher beliefs about employment also have higher expected salaries. This assumption is also consistent with the empirical literature suggesting that those with higher salaries also have stronger labor force attachment. This means that the α-quantile of the distribution of $E[y|\theta]$, $Q_\alpha(E[y|\theta])$, is given by the product of the two quantiles:

$$Q_\alpha(E[y|\theta]) = Q_\alpha(E[y|y>0,\theta])\, Q_\alpha(\Pr[y>0|\theta]) \tag{24}$$

[40] Hu and Schennach (2008) show that a sufficient set of requirements for $g(\omega_\theta)$ to be non-parametrically identified is that the linear mapping from $g(\circ)$ to $\int \theta^y (1-\theta)^{1-y} f_{Z|\theta}(z|\theta)\, g(\omega_\theta)\, d\theta$ is injective and that the distribution of z given θ has a known mapping, $E[m(z)|\theta] = \theta$. In our setting, when the elicitations are uncorrelated, γ_j is identified through an IV regression of the elicitation on the outcome, which corresponds to the required mapping. Because the elicitations are discrete, we are formally identified to some extent from the functional form choice of g and $f_{Z|\theta}$.

[41] In order to condition on observable characteristics, X, we augment equation (21) to allow for an additional point mass that varies with $E[y|X]$:

$$G(\omega_\theta) = w\mathbf{1}\{\omega_\theta \leq E[y|X] - a\} + (1-w)\sum_i \xi_i \mathbf{1}\{\omega_\theta \leq a_i\}. \tag{23}$$

Equation (24) enables the construction of the distribution of beliefs.

Results Figure 4, Panel A presents the results for the distribution of expected future salary. The solid red line presents the point estimates for the p.d.f. of the distribution of $E[y|\theta]$. The dashed red lines present a 95% confidence interval, constructed using a standard bootstrap procedure with 1,000 replications. We find a wide dispersion. The standard deviation of the distribution of beliefs is \$20,059 unconditional on X, and \$18,600 conditional on X, which contrasts with the total standard deviation of the realized y of \$25,376. This means that slightly more than half of the variance in income is known to individuals at the time they enroll in college . This means that there is significant private information, but there is considerable risk: much of the variation in salaries is uncertain at the time of enrollment.

6.2 Distribution of Beliefs about Binary Outcomes

For our three state-contingent debt markets in which y is a binary outcome, we use a procedure to estimate the distribution of beliefs that is analogous to the employment case above. Indeed, the we use the beliefs about employment, ω_θ, that are constructed above as the relevant belief distribution for the employment-based repayment market. For the the college completion outcome, we use individuals' reported likelihood of completing their degree on time for z and reported level of parental encouragement as z'. For on-time repayment outcome, we use reported level of parental encouragement as z and expected parental financial support as z'. The resulting estimates of γ are presented in Appendix Table A4.[42] The remaining estimation procedure is identical to the employment case above.

Figure 4, panels B through D present the results for the distribution of beliefs μ_θ in each of the three state-contingent debt markets. Panel B shows the distribution of $\mu_\theta = E[y|\theta]$ for the case when y is an indicator for college completion. The plot reveals a general grouping of mass between 0.5 and 0.8, but also a thick "lower tail" of mass near $\mu_\theta = 0$, which, as we will see in Section 6.3, imposes significant adverse selection problems. Panel C shows a similar pattern for beliefs about

[42] Appendix Table A5 shows that these estimates are similar using alternative elicitations as instruments.

employment; a distinct, though comparatively smaller, mass point lies near the leftmost point of the support, suggesting a small portion of the population is nearly certain they will not find employment after graduation. Panel D presents the distribution of beliefs concerning timely loan repayment. Here, there is mass ranging from 0 to 0.6, with a left-tail that again presents a potential threat of adverse selection.

6.3 Calculating the Average Value Curve

Figure 5 translates the distribution of beliefs for each setting into the average value curve, $AV(\theta)$, for each hypothetical market. The solid blue line displays the quantiles of $E[y|\theta]$, which reflects the actuarially fair price borrowers would be able to obtain for their future y. The solid green line reflects the average value curve, which equals the average value of y from those whose expectations are at or below those of a type θ. The shaded areas represent 95% confidence intervals.

In panel A, we see that the average value curve for the salary distribution falls well below the actuarially fair price line. For example, the median individual expects to earn $20,310 in 2017. On average, the people who expect to earn at or below $20,310 have an average salary of $12,373. Unless that median individual is willing to accept a 39% discount on the value of their future earnings, the financier could not obtain a profit. In other words, they must be willing to repay $1.64 plus interest for every dollar of upfront financing.

Panel B presents the results for the state contingent debt market that requires repayment only in the event of completion. We estimate an average value curve that again falls well below the actuarially fair curve. The median individual has an 60% chance of completing college. But, the average probability of those who believe they have less than an 60% chance of completing college is 34. Unless the median individual is willing to accept a contract in which they expect to repay $1.76 plus interest for every dollar of financing, the financier cannot profit by selling them a contract.

Panel C presents the results for the state contingent market that requires repayment only in the event of employment. The median individual has an 72% chance of being employed. But, the average probability of employment for those who believe they have less than an 72% chance of being employed is 61%. Unless the median individual is willing to accept a contract in which they expect

to repay $1.19 plus interest for every dollar of financing, the financier cannot profit by selling them a contract.

Finally, Panel D presents the results for the state contingent market that requires repayment only in the event the individual is not in financial distress and delinquent on their student loans. The median individual has an 30% chance of not being delinquent; but the average repayment rate of those who believe they have a more than 70% chance of falling delinquent is 16%. Unless the median individual is willing to accept a contract in which they expect to repay $1.87 plus interest for every dollar of financing, the financier cannot profit by selling them a contract.

In summary, the average value curves show that individuals would need to accept significant discounts in order for a financier to be able to profitably offer an equity or state-contingent financing contract. The next section explores whether college-goers would be willing to accept these discounts.

7 Willingness to Accept Curves

7.1 Estimation Approach

What reservation price are individuals willing to accept in order to obtain an additional dollar of college financing? To measure this willingness to accept (WTA), we follow an approach similar to the literature on optimal social insurance. We assume that the marginal utility function is given by a constant relative risk aversion, so that the marginal utility of consumption is equal to $c^{-\sigma}$. This means that $WTA(\theta)$ in equation (3) can be written as

$$WTA(\theta) = E[y|\theta] + cov\left(y, \frac{c(y)^{-\sigma}}{E\left[c(y)^{-\sigma}|\theta\right]}|\theta\right). \tag{25}$$

We then parameterize consumption levels for each of our four outcomes. For salary, we assume a consumption function of the form $c(y) = (1-\rho)E[y] + \rho(y - E[y])$, where ρ is the impact of variation in income on consumption. Drawing from a range of possible values found in the literature, we calibrate our baseline value of relative risk aversion to be $\sigma = 2$ but assess robustness

to $\sigma = 1$ and $\sigma = 3$.[43] We draw our baseline estimate of ρ from Ganong et al. (2020), who find that $\rho = \frac{dc}{dy} = 0.23$ for common sources of income shocks (e.g. changes in firm pay). We then use the distribution of y given θ to construct the covariance term and measure $WTA(\theta)$ in equation (25).

For binary contracts, we calibrate Δc, the proportional difference in consumption between good ($y = 1$) and bad ($y = 0$) states of the world, for each outcome.[44] For the completion-contingent loan contract, we approximate the increased consumption arising from degree completion using estimates from Zimmerman (2014). Relative to a base of non-enrollee incomes, Zimmerman (2014) estimates a 90% earnings increase from earning a BA degree, compared to a 22% increase from attendance alone. This implies a difference in earnings for those who complete versus do not of 68%. We translate this into the consumption difference by multiplying by $\rho = 0.23$ to obtain a consumption effect of $\Delta c = .16$.

For the employment-contingent loan contract, we approximate the increased consumption arising from employment using estimates from Hendren (2017). Hendren (2017) estimates a causal effect of unemployment on consumption ranging from 7% to 9%. To be conservative, we let $\Delta c = .09$.

Finally, for the non-dischargeable loan contract, we approximate the increased consumption arising from non-delinquency as follows.[45] We run a two-stage least-squares regression of realized salary against delinquency status and the "Expected Salary" elicitation, instrumenting for "Expected Salary" using the log of average earnings by occupation as in Section 6.1. Assuming independent measurement error of the elicitations, the instrumented elicitation controls for the portion of salary that is ex-ante known to the borrower, so that the residual correlation between delinquency and salary captures a causal effect of one on the other. This procedure yields an estimated earnings increase of 20%, which we multiply by $\rho = 0.23$ to obtain a consumption effect of $\Delta c = .05$.

[43] Empirical estimates of relative risk aversion often fall in the range of 0.5 to 4 (Chetty, 2006; Gandelman et al., 2015; Gourinchas and Parker, 2002; Pålsson, 1996), and calibrating σ to 2 is standard practice in many consumption-savings models (Jeanne and Rancière, 2006). Note that because our population of interest is relatively young, individuals may be less risk averse than the general population (Pålsson, 1996).

[44] Note that for binary contracts, equation (25) reduces to

$$WTA(\theta) = \left(1 + \frac{1 - E[y|\theta]}{E[y|\theta]}\left(1 + \Delta c^{\rho}\right)\right)^{-1}. \tag{26}$$

where Δc is the percentage difference in consumption if $y = 1$ versus $y = 0$.

[45] To our knowledge, there does not exist existing estimates of the income or consumption difference between those who have and have not defaulted on their student debt.

7.2 Baseline Results

Figure 6 reports the results. The solid red line presents the willingness to accept curve, $WTA(\theta)$. The solid green line presents the average value curve, $AV(\theta)$. For reference, we also present the actuarially fair expected outcome, $E[y|\theta]$, in the solid blue line. The shaded regions present 95% confidence intervals constructed via a bootstrap.

In all four market settings, we find the willingness to accept curve falls everywhere below the average value curve, $WTA(\theta) < AV(\theta) \ \forall\theta$. For the equity contract, we estimate the median individual is willing to accept a valuation of \$15,824, which is a 22% discount relative to their expected salary of \$20,310. But, it lies above the expected incomes of those who are willing to accept less than \$15,824, which is $AV(0.5)$ =\$12,373. The p-value for the test that there exists a value of θ such that $AV(\theta) \geq WTA(\theta)$ is less than 0.001.

In the completion contract market, we estimate that the median individual is willing to accept a valuation of $WTA(0.5)$ =\$0.53 per dollar lent, which is a 43 discount relative to their expected likelihood of completion of 60%. But, the mean completion likelihood amongst those willing to accept a valuation of up to \$0.53 per dollar is 34%, which means the firm cannot make a profit by setting $\lambda = \$0.53$. The p-value for the test that there exists a value of θ such that $AV(\theta) \geq WTA(\theta)$ is again less than 0.001.

For the state-contingent market requiring repayment in the event of employment, the WTA curve again lies above the AV curve. We estimate that the median individual is willing to accept a valuation of $WTA(0.5)$ =\$0.69 per dollar lent, which is a 16% discount relative to their expected likelihood of employment of 72%. But, the employment likelihood amongst those willing to accept a valuation of up to \$0.69 is 61%, which means the firm cannot make a profit by setting $\lambda = \$0.69$. The p-value for the test that there exists a value of θ such that $AV(\theta) \geq WTA(\theta)$ is less than 0.001.

For the state-contingent market requiring repayment in the event of non-delinquency on traditional student loans, we again find that the WTA curve lies above the AV curve. We estimate that the median individual is willing to accept a valuation of $WTA(0.5)$ =\$0.28 per dollar lent, which

is a 46% discount relative to their expected likelihood of not being delinquent on student debt of 30%. But, the likelihood amongst those willing to accept a valuation of up to \$0.28 is 16%, which means the firm cannot make a profit by setting$\lambda = \$0.28$. The p-value for the test that there exists a value of θ such that $AV(\theta) \geq WTA(\theta)$ is less than 0.001.

In sum, in all four market settings we find evidence consistent with market unraveling: the $WTA(\theta)$ curve lies everywhere above the $AV(\theta)$ curve.

7.3 Robustness

We assess the robustness of these estimates to three sets of alternative assumptions: (1) changes in the parameters σ used to construct the WTA curve and (2) differences in the relative risk-free interest rates faced by college-goers versus financiers, and (3) a relaxation of our assumption that willingness to accept varies only with the expected outcome (Assumption 2).

Alternative Risk Aversion Figure 7 presents the WTA curves using alternative values of the coefficient of relative risk aversion of 1 and 3. As expected, a higher coefficient of relative risk aversion leads to a lower WTA curve, but it still lies above the AV curve.

Interest Rates Our baseline specification assumes that individuals face the same interest rate on (non-dischargeable) student debt as financiers face. Currently the government-mandated interest rate on federal student loans for undergraduates is 3.7%, though the effective rate is lower after accounting for grace periods and deferments that pause the accrual of interest.[46] If this interest rate faced by college-goers is less than the rate faced by financiers due to the implicit subsidy, this would only reinforce our results: borrowers would only be willing to accept especially high-value contracts to compete with the favorable terms offered by a subsidized loan. The key concern is the other direction: one worries that college-goers face a higher interest rate on additional student debt relative to the interest rate faced by financiers. To assess the implications of this, the long-dash

[46] Congress has set rates on student loans since 1965, though it automated the process in 2013 with the Bipartisan Student Loan Certainty Act, which sets interest rates equal to the 10-year Treasury bond rate plus 205 basis points (360 bps for graduate students). Interest rates are fixed throughout the life of a loan and accrue as simple daily interest on principal only. Most interest does not accrue while borrowers are enrolled in college or graduate school, and all interest accrual has been temporarily paused in response to the COVID-19 pandemic.

line in Figure 7 reports the WTA curve under an assumption that college-goers face a 10pp higher interest rate than firms. This leads to a lower WTA curve, but it still falls above the AV curve. This suggests that even if college goers are significantly credit constrained on the margin, their WTA is not be low enough to make the state-contingent financing markets profitable.

Preference Heterogeneity Assumption 2 prevents those with the same expected outcome, $\mu_\theta = E[y|\theta]$, from having different willingness to accept, $WTA(\theta)$. This is potentially worrisome because a large body of work that studies adverse (or advantageous) selection finds evidence of multi-dimensional preference heterogeneity. Here, we assess how the presence of preference heterogeneity would affect our core conclusions.

Recall that when Assumption 2 does not hold, the AV curve defined in equation (6) no longer summarizes the impact of private information on the market. Instead, we need to use the more general average value curve in equation (5) that depends on the joint distribution of $E[y|\theta]$ and $WTA(\theta)$.

We calibrate two alternative specifications that allow for preference heterogeneity. In the first specification, we assume that the coefficient of relative risk aversion, σ, is drawn from a uniform distribution between 1 and 3, $\sigma \sim Unif[1,3]$. In the second specification, we assume σ is drawn from a uniform distribution between 0 and 4, $\sigma \sim Unif[0,4]$. Both specifications retain the mean risk aversion coefficient at 2, but introduce different degrees of heterogeneity.[47] We then compute AV curves under these alternative specifications using the more general form in equation (5).

Figure 8 plots the resulting AV curves that incorporate these two preference heterogeneity specifications.[48] We find the average value curve is virtually unchanged in both specifications and still falls below the willingness-to-accept curve—the market would continue to unravel even under

[47] Our simulation assumes that preference heterogeneity is not correlated with the level of the expected outcome. We view this as a natural benchmark. There is no robust reduced form evidence of correlated preference heterogeneity in other settings. In health contexts, several earlier studies have argued that there is 'advantageous selection' generated by the 'worried well', however Section 8.4 in Hendren (2013) argues that these correlations in earlier literature are likely driven by insurance companies choosing not to insure observably sick applicants as opposed to sick applicants having less preference for insurance.

[48] Note that we evaluate average value (equation (5)) at the same values of θ_λ as in the baseline case with homogenous risk preferences. As a result, the horizontal axis in Figure 8 corresponds to quantiles of WTA among individuals with $\sigma = 2$, rather than quantiles of the entire WTA distribution under preference heterogeneity. Appendix Figure A3 provides an alternative plot in which the horizontal axis follows a quantile-WTA interpretation for the alternative specification with highest risk-aversion, $\sigma \sim Unif[0,4]$.

these heterogeneous risk preferences.

Biased Beliefs Our baseline model and estimation approach relies on an assumption of unbiased beliefs (Assumption 1). However, previous literature has argued that college-goers may be over-optimistic about their future earnings prospects. When beliefs about average future income, μ_θ, exceed the average realized income of individuals holding those beliefs, $E[y|\theta]$, the willingness to accept in equation (3) is increased by a bias term, $bias = \mu_\theta - E[y|\theta]$.[49] Intuitively, overconfident college-goers would need to be offered a larger amount of money today to give up a fraction of their future earnings. In this case, the WTA curve would shift upward, strengthening our unraveling result.[50]

8 Welfare Impacts of Government Subsidies

If the private market can't profit from providing these contracts, should the government subsidize the provision of risk-mitigating contracts to help open up these markets? Should we have free college financed by general taxation? Or should we ask people who go to college to pay back an extra fraction of their future earnings if they go to college relative to those who don't go to college? How about government subsidies for other state-contingent debt contracts that would offer greater insurance value than the current Perkins and Stafford loan programs? These questions have obtained considerable theoretical interest in the economics literature (e.g. Jacobs and van Wijnbergen (2007); Stantcheva (2017)) and in recent consideration in political debates about student debt burdens (Warren, 2020; Harrison, 2021). In this section, we use our empirical estimates along with the framework from Section 2 to measure the welfare impact of subsidizing the four types of risk-mitigating contracts we study above.

We do so by constructing the marginal value of public funds (MVPF) of subsidies to these contracts. The MVPF measures the dollar value of the policy provided to its beneficiaries per

[49] Heterogeneity in the degree of bias would introduce variation in beliefs that is orthogonal to one's average outcome – as shown in Figure 8, this type of variation is unlikely to affect our conclusions.

[50] In the perhaps less relevant case where college-goers were uniformly pessimistic about their future outcomes, the unraveling condition in equation (8) quantifies how large the negative bias would need to be in order to prevent market unraveling.

dollar of net cost the government incurs from subsidizing these contracts. Comparisons of MVPFs across policies correspond to statements about the welfare impact of hypothetical budget neutral policies (Hendren and Sprung-Keyser, 2020). As a result, the MVPFs we construct here can be compared to the broader library of MVPFs for government expenditure policies constructed in Hendren and Sprung-Keyser (2020), Finkelstein and Hendren (2020), and others.

8.1 Equity Contracts

We begin with the MVPF of government subsidies for an equity contract for human capital.[51] We imagine that the government provides $\$\eta\lambda$ of equity financing for college in exchange for a η-share of earnings six years later, where we take $\eta \approx 0$ and consider the MVPF for a small increase in the amount of financing individuals can obtain, η. We compute the MVPF of an equity contract that prices a claim on an individual's earnings at mean realized salary for the population, $\lambda = E[y] = \$24{,}032$. This contract provides individuals with $\frac{\lambda}{100} = \$240$ in financing today in exchange for 1% of their post-college income in 2017.

An individual's willingness to pay for the subsidy is the difference between the price they accept, λ, and the amount they would have been willing to accept, $\lambda - wta(\theta)$. The average willingness to pay per person who takes up the contract is $\lambda - E[y|\theta \leq \theta_\lambda]$, or $1 - \frac{1}{\lambda}E[y|\theta \leq \theta_\lambda]$ per unit of the price, λ. If these were the only costs, the MVPF of offering subsidies at valuation λ would be given by $\frac{\lambda - E[wta(\theta)|\theta \leq \theta_\lambda]}{\lambda - E[y|\theta \leq \theta_\lambda]}$. The policy would have an MVPF that exceeds one to the extent to which individuals' willingness to accept is less than their expected future incomes. However, from the government perspective, there is an additional fiscal externality from any behavioral responses to the subsidy. While the first dollar of the contract does not affect the profitability of a financier, it does have first-order effects on tax revenue because of pre-existing positive tax rates. In Appendix E,

[51] We assume for simplicity that financiers earn zero profits, so that the MVPF for subsidies of these contracts in a private market is equivalent to the MVPF of direct government provision of the contracts.

we show that the MVPF can be written as:

$$MVPF = \frac{1 - \frac{1}{\lambda} E\left[WTA\left(\theta\right) | \theta \leq \theta_\lambda\right]}{1 - \frac{1}{\lambda} E\left[y | \theta \leq \theta_\lambda\right] + \underbrace{\frac{1}{\lambda} \frac{\tau}{1-\tau} E\left[y | \theta \leq \theta_\lambda\right] \epsilon_{y,1-\tau}}_{\text{Tax Distortion}} - \underbrace{\tau \frac{dE\left[y^L\right]}{dg} \frac{1}{\Pr\left\{\theta \leq \theta_\lambda\right\}}}_{\text{Impact on Earnings via Grant}}}, \tag{27}$$

where the two latter terms in the denominator quantify the impact of the behavioral response to the equity contract on government tax revenue. The parameter τ is the marginal tax rate; $\epsilon_{y,1-\tau} = \frac{1-\tau}{E[y|\theta \leq \theta_\lambda]} \frac{dE[y|\theta \leq \theta_\lambda]}{d(1-\tau)}$ is the elasticity of taxable income in 2017 with respect to taxes in 2017; and $\tau \frac{dE\left[y^L\right]}{dg} \frac{1}{\Pr\{\theta \leq \theta_\lambda\}}$ is the impact of a \$1 grant for college financing on lifetime tax payments.[52] The MVPF exceeds 1 to the extent to which individuals' economic surplus from the contract, which is given by the difference between their expected incomes and their willingness to accept, $E\left[y|\theta\right] - wta\left(\theta\right)$, exceeds the fiscal externality that their financing imposes on the government, which is given by the sum of the tax distortion and earnings impact of a grant.

While we do not directly observe this fiscal externality from variation in our data, equation (27) shows that it depends on two parameters commonly estimated in previous literature: the elasticity of taxable income, $\epsilon_{y,1-\tau}$, and the impact of a grant for college on lifetime earnings. For $\epsilon_{y,1-\tau}$, we draw upon estimates from the large literature on the taxable income elasticity, surveyed in Saez et al. (2012). This literature finds a midpoint elasticity of taxable income with respect to the keep rate of $\epsilon_{y,1-\tau} = 0.3$.[53]

For the impact of college funding on future tax payments per dollar, $\tau \frac{dE\left[y^L\right]}{dg} \frac{1}{\Pr\{\theta \leq \theta_\lambda\}}$, we use estimates of loan financing from Gervais and Ziebarth (2019), who use discontinuities in federal student debt limits to estimate the impact of loan financing on earnings.[54] They find that \$1,000 in loan financing increases earnings by 0.028 ten years after graduation (Gervais and Ziebarth (2019)

[52] We use $dE\left[y^L\right]/dg$ to denote the impact of a grant on average lifetime income in the population so that normalizing by the take-up rate $\Pr\left\{\theta \leq \theta_\lambda\right\}$ corresponds to the treatment on the treated effect of the grant. This term is scaled by λ to capture the fact that the marginal η contract provides $\eta\lambda$ units of a grant.

[53] Saez et al. (2012) document a wide variation of estimates, and we therefore consider a value of 1 (and 0) for a point of comparison. We also discuss estimates from Britton and Gruber (2019) who exploit discontinuities in the slope of the marginal tax rate for income-contingent loans in the UK. They find no behavioral response from these "kinks" in the tax schedule, which would imply a compensated elasticity of $\epsilon_{y,1-\tau} = 0$.

[54] Note that because our elasticity estimates are compensated, we use the impact of a loan as opposed to grant on outcomes.

Table 4, Column 1). Applying this growth rate to the distribution of expected incomes in our sample yields an average implied income change of \$0.78 per dollar of financing.[55] We assume throughout a tax rate of $\tau = 20\%$, so that the average change in tax revenue is 0.156 per dollar of spending.

Table 8 outlines our estimates of each component of the MVPF in equation (27). The contract would break even if everyone were to accept the contract. But, our results suggest that 79% of the population would take up the contract. Their expected outcomes would be equal to $E[y|\theta \leq \theta_\lambda] =$ \$12,209, which is \$11,823 less than the average earnings in the population. However, the government recoups 5% of each dollar financed from the impact on earnings in adulthood (9% from the reduction in earnings due to higher tax rates and -4% from the increase in earnings due to the college financing). On average, amongst the 79% who take up the contract, they are on average willing to accept \$15,490, which means they enjoy an average surplus of \$3,281. This means that the consumption smoothing benefits from the equity contract (\$0.17 per dollar of financing) are four times larger than the costs of distortions from higher implicit tax rates (\$0.04 per dollar of financing when $\epsilon = 0.3$).[56]

Totaling these components, we estimate an MVPF of the policy of 1.86.[57] Every dollar of government spending would generate \$1.86 of benefits to college-goers. Comparing to the library of estimates in Hendren and Sprung-Keyser (2020), this MVPF is higher than most other social programs in the US, aside from targeted direct investments in low-income children.

[55] It is also possible that the availability of equity finance increases college enrollment, not just completion. Any increased earnings through such enrollment effects would increase tax revenue and further lower the net cost of the policy, increasing its MVPF. We omit these effects in our calculations, so that our estimates should be seen as conservative.

[56] Even if $\epsilon = 1$, the distortion is still less than the consumption smoothing benefits.

[57] If the elasticity of taxable income is zero instead of 0.3, the MVPF would be 2.51; if the elasticity of taxable income is 1, the MVPF would be 1.45.

8.2 Completion-Based Repayment, Employment-Based Repayment, and Dischargeable Loans

We also construct the MVPF of state-contingent debt contracts. Appendix E shows that the MVPF for binary outcomes has a very similar formula to equation (27):

$$MVPF^{binary} \approx \frac{1-\frac{1}{\lambda}E\left[WTA\left(\theta\right)|\theta\leq\theta_{\lambda}\right]}{-\frac{1}{\lambda}E\left[y|\theta\leq\theta_{\lambda}\right]+\frac{\tau}{\lambda}\frac{1}{\Pr\{\theta\leq\theta_{\kappa}\}}\frac{dE[y^{L}]}{dD}-\tau\frac{dE[y^{L}]}{dg}\frac{1}{\Pr\{\theta\leq\theta_{\lambda}\}}}$$

The distinction relative to equation (27) is that the behavioral response due to higher implicit tax rates from the equity contract is replaced with the causal effect of the debt repayment incentives on earnings, $\frac{\tau}{\lambda}\frac{1}{\Pr\{\theta\leq\theta_{\kappa}\}}\frac{dE\left[y^{L}\right]}{dD}$. For the employment-contingent loan contract, we draw upon the literature on UI that shows behavioral responses to UI mean that every \$1 of UI spending actually costs the government around \$1.50 (Schmieder and Von Wachter (2016)). Since the odds of being unemployed are 12.3%, this implies every \$1 of financing that only requires repayment in the event of employment has an additional cost of $\frac{0.123}{1-0.123}0.5 = .07$ to the government. To the best of our knowledge, there does not exist empirical evidence on the impact of dischargeable loans and completion-based repayment contracts on taxable income. We therefore assume for simplicity that this fiscal externality per person taking up the contract is equal to the fiscal externality from the earning-based repayment disincentive.[58] The calculation in Table 8 breaks out the MVPF calculation to show clearly how future empirical evidence on these fiscal externalities could affect the results.

Completion-Based Repayment The second row of Table 8 presents the results for the completion-contingent loan. We estimate that roughly 52% of individuals would take-up this contract at a price that would reflect population-wide actuarially fair prices ($\lambda = E\left[y\right]$). Every \$1 of financing would cost the government \$0.31 to cover the cost of adverse selection. This would deliver a consumption smoothing benefit of \$0.10, for a total WTP of \$0.41. The net cost to the government includes the upfront cost of adverse selection of \$0.31, minus the benefits from increased earnings from

[58] In principle, the fiscal externalities reflect not only any earnings effects, but also any effects on loan repayments that lead the government to not fully recoup their existing base of student loan spending.

the financing of \$0.09, plus the cost of distortions arising from the disincentives in the contract of \$−0.13, for a total net cost of \$0.35. This implies an MVPF of 1.16.

Employment-Based Repayment The third row of Table 8 presents the results for the employment-contingent loan contracts. We estimate that \$56 of the population would take up this contract at $\lambda = E[y]$. The cost from adverse selection per dollar of financing would be \$0.11. The consumption smoothing benefit is \$0.05, for a total WTP of \$0.17. The net cost to the government includes the \$0.11 subsidy to cover the cost of adverse selection, minus the positive impact on future tax revenue from increased human capital of \$0.10, plus the impact of distortions on future earnings of \$−0.10. This implies a total net cost of \$0.12, and an implied MVPF of \$1.42.

Dischargeable Loans The fourth row of Table 8 presents the results for the dischargeable loans that embody greater limited liability protection. We estimate roughly 44% of the population would take up this contract at $\lambda = E[y]$. Every \$1 of financing would cost the government \$0.73 to cover the cost of adverse selection. These contracts would deliver \$0.02 of consumption smoothing benefits to the beneficiaries, for a net WTP of \$0.75. The cost to the government includes the \$0.73 cost of adverse selection minus the benefits of \$0.08 from increased tax revenue resulting from increased future earnings. This is counteracted by the incentive to lower incomes and default of \$0.30, for a net cost of \$0.94. Combining, this implies an MVPF of 0.79.

8.3 "Free College"

Instead of subsidizing equity contracts that ask college-goers to pay back some fraction of their future earnings, the government could alternatively subsidize college financing without requiring repayment. In the model, this corresponds to the limit as the price goes to infinity, $\lambda \to \infty$, so that

$$MVPF^{FreeCollege} = \frac{1}{1 - \tau \frac{dE[y^L]}{dg} \frac{1}{\Pr\{\theta \leq \theta_\lambda\}}}$$

The MVPF equals 1 if the financing has no effect on future earnings. But, the cost to the government is reduced to the extent to which providing free college increases future earnings. The estimates

from Gervais and Ziebarth (2019) suggest that tax revenue increases by 0.15, so that the net cost per dollar of spending is 0.85, which implies an MPVF of 1.17. The MVPF of an unconditional grant is lower than the MPVF of the equity contract, which was 1.86. This is because the equity contract offers an "efficient" method of asking individuals to repay: on average, the consumption-smoothing benefits of the equity contract (difference between $WTA(\theta)$ and $E[y|\theta]$) exceed the distortionary costs from higher post-college tax rates, $\frac{\tau}{1-\tau}E[y|\theta \leq \theta_\lambda]\,\epsilon_{y,1-\tau}$. While fewer people would choose to take up the equity contract, those who do are the ones that expect to have lower future incomes. In this sense, equity financing is not only more efficient but also its induced self-selection into the contract also provides positive equity benefits.

9 Conclusion

This paper develops a framework for quantifying the frictions imposed by private information in markets for financing human capital investment. Our results suggest that adverse selection prevents private markets from offering risk-mitigating financial contracts like the equity contracts envisioned by Friedman (1955). From a policy standpoint, our results relate to recent debates about the rising burden of college debt. Methods of financing beyond the dominant non-dischargeable debt contracts that reduce the debt burden for those who grow up to have low incomes could deliver significant welfare gains.

Our results contribute to a growing literature suggesting that the set of financial markets we observe is limited by the existence of private information. To that aim, our framework and empirical approach could be used beyond the education financing literature to other settings. For example, the Small Business Administration spends significant resources intervening in capital markets for firms. And, in times of crisis, the Federal Reserve has recently expanded its balance sheet to take equity stakes in private enterprise. Our framework and methods could be extended to these settings to understand the frictions preventing efficient capital markets and the welfare impacts of this type of government intervention. Our methods could also be used to investigate the role of private information elsewhere in the labor market. For example, adverse selection might help explain why

some industries fail to form unions, or why some occupations pay piece rates rather than flat wages. In the meantime, our results here suggests that adverse selection in the market for human capital financing may limit the economic opportunities available to many potential college-goers.

References

Abraham, Filiz-Ozbay, Ozbay, and Turner (2018) "Behavioral Effects of Student Loan Repayment Plan Options on Borrowers' Career Decisions: Theory and Experimental Evidence," July, National Bureau of Economic Research Working Paper 24804.

Abraham, Katharine G, Emel Filiz-Ozbay, Erkut Y Ozbay, and Lesley J Turner (2018) "Framing Effects, Earnings Expectations, and the Design of Student Loan Repayment Schemes," National Bureau of Economic Research Working Paper 24484.

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

Agarwal, Sumit, Souphala Chomsisengphet, and Chunlin Liu (2010) "The Importance of Adverse Selection in the Credit Card Market: Evidence from Randomized Trials of Credit Card Solicitations," *Journal of Money, Credit and Banking*, 42 (4), 743–754.

Aigner, Dennis J, Cheng Hsiao, Arie Kapteyn, and Tom Wansbeek (1984) "Latent variable models in econometrics," *Handbook of econometrics*, 2, 1321–1393.

Akerlof, George A (1970) "The market for lemons: Quality uncertainty and the market mechanism," *The quarterly journal of economics*, 488–500.

Arcidiacono, Peter, V Joseph Hotz, Arnaud Maurel, and Teresa Romano (2020) "Ex ante returns and occupational choice," *Journal of Political Economy*, 128 (12), 4475–4522.

Ausubel, Lawrence M. (1999) "Adverse Selection in the Credit Card Market," Working Paper.

Barr, Nicholas, Bruce Chapman, Lorraine Dearden, and Susan Dynarski (2017) "Getting student financing right in the US: lessons from Australia and England," Centre for Global Higher Education Working Paper.

Berman, Jillian (2017) "Students Get Tuition Aid for a Piece of Their Future," *The Wall Street Journal*.

Bonhomme, Stéphane and Jean-Marc Robin (2010) "Generalized non-parametric deconvolution with an application to earnings dynamics," *The Review of Economic Studies*, 77 (2), 491–533.

Britton, Jack W and Jonathan Gruber (2019) "Do Income Contingent Student Loan Programs Distort Earnings? Evidence from the UK," National Bureau of Economic Research Working Paper 25822.

Brunnermeier, Thomas Eisenbach, Markus K and Yuliy Sannikov (2013) "Macroeconomics with Financial Frictions: A Survey."

Chapman, Bruce (2006) "Income contingent loans for higher education: International reforms," *Handbook of the Economics of Education*, 2, 1435–1503.

Chapman, Bruce and Andrew Leigh (2009) "Do very high tax rates induce bunching? Implications for the design of income contingent loan schemes," *Economic Record*, 85 (270), 276–289.

Chetty, Raj (2006) "A new method of estimating risk aversion," *American Economic Review*, 96 (5), 1821–1834.

Conlon, John J, Laura Pilossoph, Matthew Wiswall, and Basit Zafar (2018) "Labor market search with imperfect information and learning,"Technical report, National Bureau of Economic Research.

Cowen, Tyler (2019) "How Much Is Your Education Worth? Depends How Much You Make," *Bloomberg News*.

Cox, James C, Daniel Kreisman, and Susan Dynarski (2018) "Designed to Fail: Effects of the Default Option and Information Complexity on Student Loan Repayment," National Bureau of Economic Research Working Paper 25258.

Delaigle, Aurore and Irène Gijbels (2004) "Practical bandwidth selection in deconvolution kernel density estimation," *Computational statistics & data analysis*, 45 (2), 249–267.

Department of Education (2021) "Federal Student Loan Portfolio," Website, `https://studentaid.gov/data-center/student/portfolio` Accessed: 2021-07-14.

Dobbie, Will and Paige Marta Skiba (2013) "Information Asymmetries in Consumer Credit Markets: Evidence from Payday Lending," *American Economic Journal: Applied Economics*, 5 (4), 256–282.

Einav, Liran, Amy Finkelstein, and Mark R Cullen (2010) "Estimating welfare in insurance markets using variation in prices," *The quarterly journal of economics*, 125 (3), 877–921.

Einav, Liran, Amy Finkelstein, and Neale Mahoney (2021) "The IO of Selection Markets," *Working Paper*.

Einav, Liran, Mark Jenkins, and Jonathan Levin (2012) "Contract Pricing in Consumer Credit Markets," *Econometrica*, 80 (4), 1387–1432.

Field, Erica (2009) "Educational debt burden and career choice: Evidence from a financial aid experiment at NYU Law School," *American Economic Journal: Applied Economics*, 1 (1), 1–21.

Finkelstein, Amy and Nathaniel Hendren (2020) "Welfare Analysis Meets Causal Inference," *Journal of Economic Perspectives*, 34 (4), 146–67.

Friedman, Milton (1955) "The Role of Government in Education."

Gandelman, Nestor, Ruben Hernandez-Murillo et al. (2015) "Risk Aversion at the Country Level," *Federal Reserve Bank of St. Louis Review*, 97 (1), 53–66.

Ganong, Peter, Damon Jones, Pascal Noel, Fiona Greig, Diana Farrell, and Chris Wheat (2020) "Wealth, race, and consumption smoothing of typical income shocks," *NBER Working Paper* (w27552).

Gervais, Martin and Nicolas L Ziebarth (2019) "Life after debt: Postgraduation consequences of federal student loans," *Economic Inquiry*, 57 (3), 1342–1366.

Gong, Yifan, Todd Stinebrickner, and Ralph Stinebrickner (2019) "Uncertainty about future income: Initial beliefs and resolution during college," *Quantitative Economics*, 10 (2), 607–641.

Gourinchas, Pierre-Olivier and Jonathan A Parker (2002) "Consumption over the life cycle," *Econometrica*, 70 (1), 47–89.

Gupta, Arpit and Christopher Hansman (2019) "Selection, Leverage, and Default in the Mortgage Market," Working Paper.

Guvenen, Fatih (2007) "Learning your earning: Are labor income shocks really very persistent?" *American Economic Review*, 97 (3), 687–712.

Harrison, David (2021) "Schumer Again Calls On Biden to Cancel $50,000 per Borrower in Federal Student Debt," *The Wall Street Journal*, `https://www.wsj.com/articles/schumer-again-calls-on-biden-to-cancel-50-000-per-borrower-in-federal-student-debt-11624911980` Accessed: 2021-07-14.

Hartley, Jon (2016) "Purdue's Income Sharing Agreement Solution To The Student Debt Crisis," *Forbes Magazine*.

Hendren, Nathaniel (2013) "Private information and insurance rejections," *Econometrica*, 81 (5), 1713–1762.

——— (2017) "Knowledge of future job loss and implications for unemployment insurance," *American Economic Review*, 107 (7), 1778–1823.

Hendren, Nathaniel and Ben Sprung-Keyser (2020) "A unified welfare analysis of government policies," *The Quarterly Journal of Economics*, 135 (3), 1209–1318.

Herbst, Daniel (2020) "Liquidity and Insurance in Student Loan Contracts: The Effects of Income-Driven Repayment on Default and Consumption," Working Paper.

Horowitz, Joel L and Marianthi Markatou (1996) "Semiparametric estimation of regression models for panel data," *The Review of Economic Studies*, 63 (1), 145–168.

Hu, Yingyao and Susanne M Schennach (2008) "Instrumental variable treatment of nonclassical measurement error models," *Econometrica*, 76 (1), 195–216.

Jacobs, Bas and Sweder JG van Wijnbergen (2007) "Capital-market failure, adverse selection, and equity financing of higher education," *FinanzArchiv/Public Finance Analysis*, 1–32.

Jeanne, Olivier and Romain G Rancière (2006) "The optimal level of international reserves for emerging market countries: formulas and applications."

Karamcheva, Nadia, Jeffrey Perry, Constantine Yannelis et al. (2020) "Income-Driven Repayment Plans for Student Loans: Working Paper 2020-02,"Technical report.

Karlan, Dean and Jonathan Zinman (2009) "Observing unobservables: Identifying information asymmetries with a consumer credit field experiment," *Econometrica*, 77 (6), 1993–2008.

Kreighbaum, Andrew (2019) "Lifting the Curtain on Income-Share Agreements," *Inside Higher Ed*, `https://www.insidehighered.com/news/2019/09/26/two-emerging-players-isas-deliberate-and-different-approaches` Accessed: 2021-02-12.

Ladine, Bret (2001) "'70s debt program finally ending," *Yale Daily News*, `https://yaledailynews.com/blog/2001/03/27/70s-debt-program-finally-ending/` Accessed: 2021-02-12.

Li, Tong and Quang Vuong (1998) "Nonparametric estimation of the measurement error model using multiple indicators," *Journal of Multivariate Analysis*, 65 (2), 139–165.

Looney, Adam and Constantine Yannelis (2015) "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.

Meinshausen, Nicolai (2006) "Quantile regression forests," *Journal of Machine Learning Research*, 7 (Jun), 983–999.

Milgrom, Paul and Ilya Segal (2002) "Envelope theorems for arbitrary choice sets," *Econometrica*, 70 (2), 583–601.

Mueller, Holger M and Constantine Yannelis (2019) "Reducing Barriers to Enrollment in Federal Student Loan Repayment Plans: Evidence from the Navient Field Experiment," Unpublished manuscript.https://faculty.chicagobooth.edu/constantine.yannelis/IBR.pdf.

Mumford, Kevin J (2020) "Student selection into an income share agreement," Working Paper.

Palacios, Miguel (2004) *Investing in human capital: a capital markets approach to student funding*: Cambridge University Press.

Pålsson, Anne-Marie (1996) "Does the degree of relative risk aversion vary with household characteristics?" *Journal of economic psychology*, 17 (6), 771–787.

Rudegeair, Peter (2016) "Online Lender's Secret to Success: Past Failure," *The Wall Street Journal*.

Saez, Emmanuel, Joel Slemrod, and Seth H Giertz (2012) "The elasticity of taxable income with respect to marginal tax rates: A critical review," *Journal of Economic Literature*, 50 (1), 3–50.

Schmieder, Johannes F and Till Von Wachter (2016) "The effects of unemployment insurance benefits: New evidence and interpretation," *Annual Review of Economics*, 8, 547–581.

Shavell, Steven et al. (1979) "On Moral Hazard and Insurance," *The Quarterly Journal of Economics*, 93 (4), 541–562.

Stantcheva, Stefanie (2017) "Optimal taxation and human capital policies over the life cycle," *Journal of Political Economy*, 125 (6), 1931–1990.

Stroebel, Johannes (2016) "Asymmetric Information about Collateral Values," *Journal of Finance*, 71 (3), 1071–1112.

University of Utah (2021) "Invest in U," Website, https://isa.utah.edu Accessed: 2021-02-12.

Warren, Elizabeth (2020) "My Plan to Cancel Student Loan Debt on Day One of my Presidency," Website, https://elizabethwarren.com/plans/student-loan-debt-day-one Accessed: 2021-07-14.

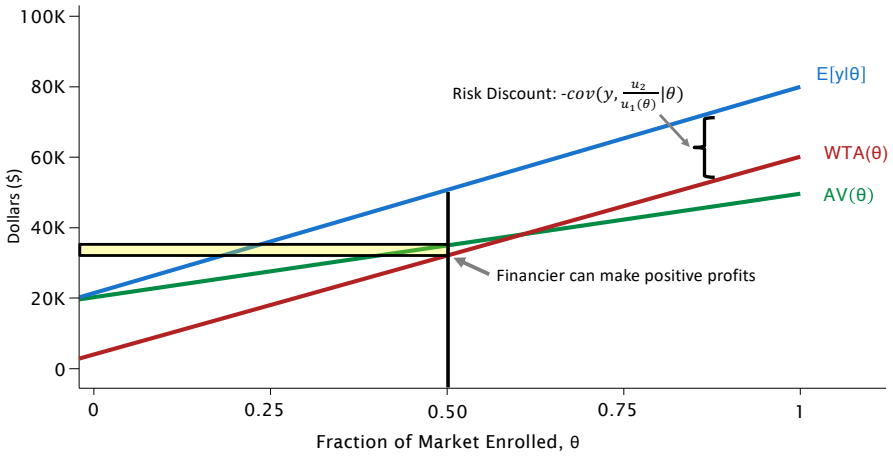
Wiswall, Matthew and Basit Zafar (2021) "Human capital investments and expectations about career and family," *Journal of Political Economy*, 129 (5), 1361–1424.

Zimmerman, Seth D (2014) "The returns to college admission for academically marginal students," *Journal of Labor Economics*, 32 (4), 711–754.

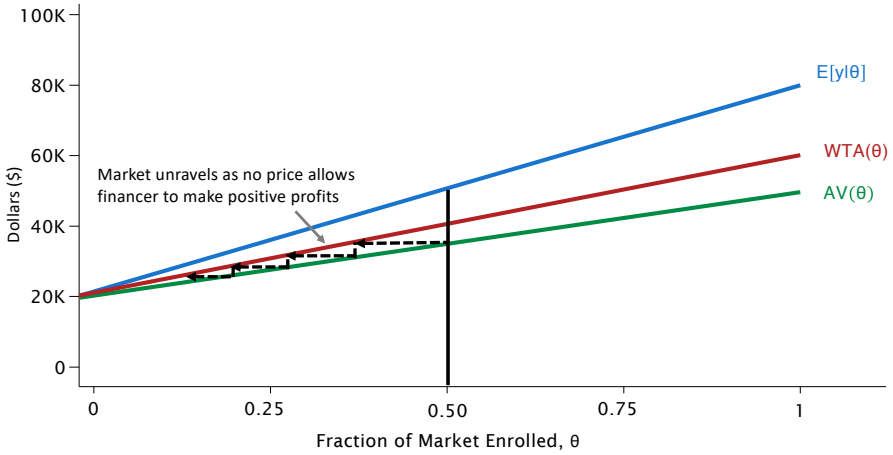
Zingales, Luigi (2012) "The College Graduate as Collateral," *The New York Times*, `https://www.nytimes.com/2012/06/14/opinion/the-college-graduate-as-collateral.html` Accessed: 2021-02-12.

Tables and Figures

Figure 1: Model of Market Unraveling: $AV(\theta)$ and $WTA(\theta)$ Curves



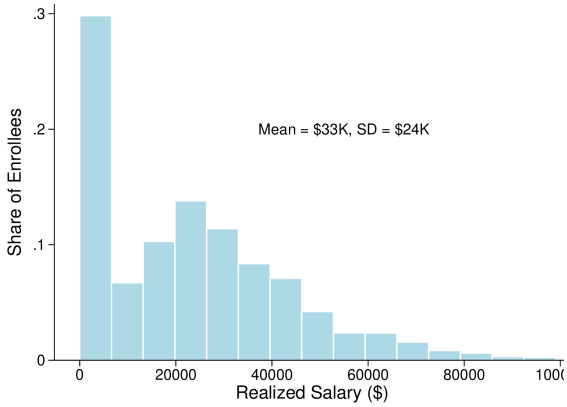
(A) Firms Can Make Profits: $WTA(\theta) < AV(\theta)$ **For Some** θ



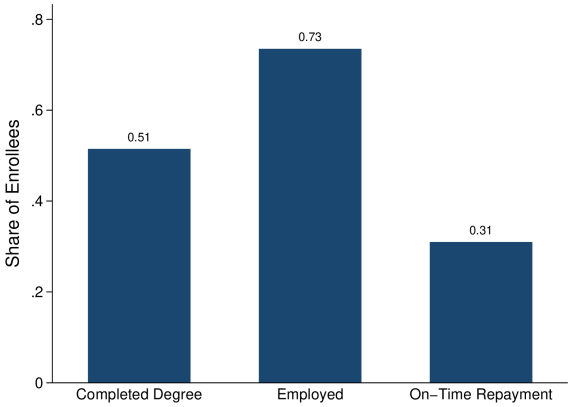
(B) Market Fully Unravels, $AV(\theta) < WTA(\theta)$ **For All** θ

Note: This figure provides a graphical representation of market unraveling for an equity contract. The blue line plots uniformly-distributed quantiles of privately-expected salaries, $\mu_\theta = E[y|\theta]$, between \$20,000 and \$80,000. The red line plots the willingness-to-accept curve, $WTA(\theta)$. The green line plots the average value curve, $AV(\theta)$, which corresponds to the average expected salary among those with who expect incomes below the corresponding point on the $E[y|\theta]$ line. On the horizontal axis, types θ are enumerated in ascending order based on their willingness to accept, $WTA(\theta)$. Panel A depicts a scenario where the financier can make a profit, in which individuals are willing to accept less than the \$35,000 necessary for a market to be profitable when $\theta = 0.5$. Panel B depicts a scenario where the market unravels, in which no one is willing to accept the average value of expected incomes lower than their own.

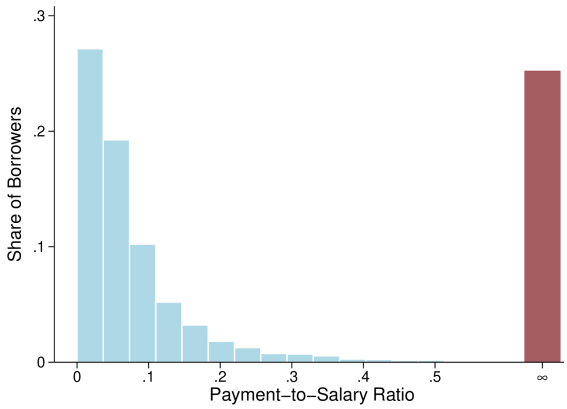
Figure 2: Summary Statistics for Contract Outcomes



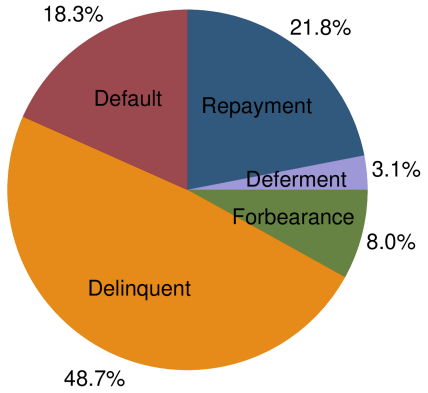
(A) Histogram of Realized Salary



(B) Mean Binary Outcomes



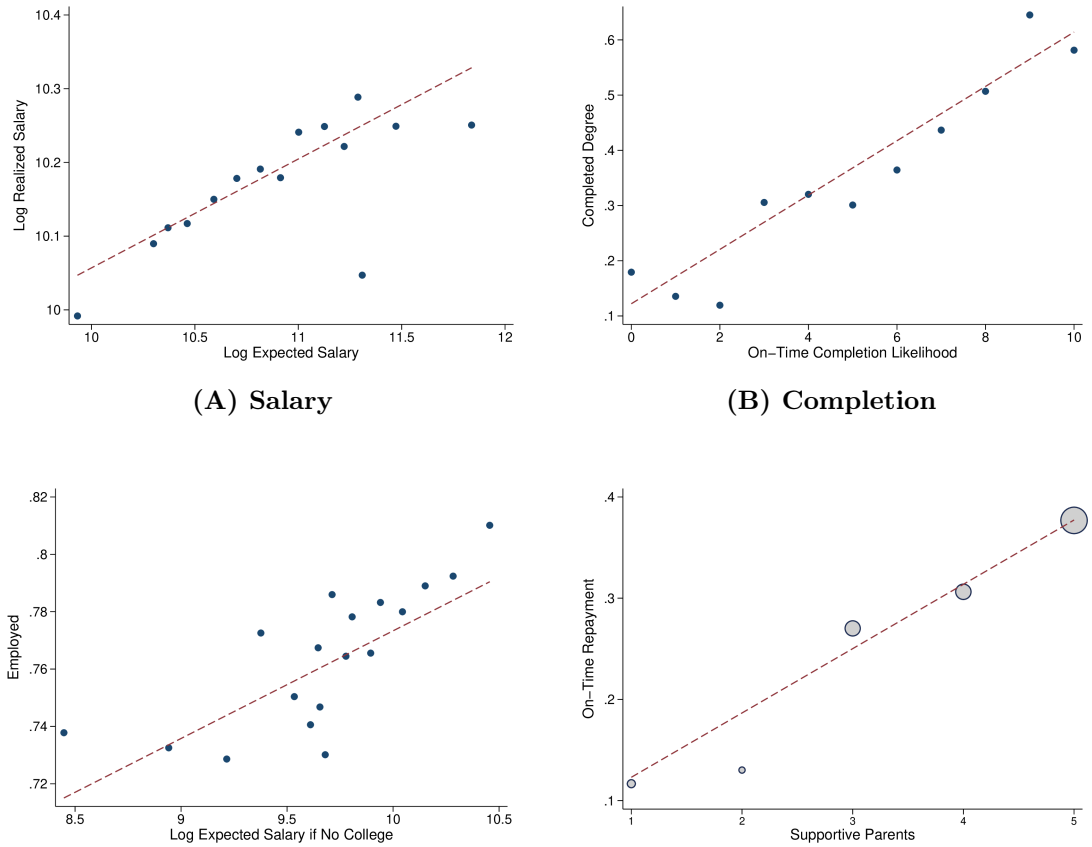
(C) Debt-Payment-to-Salary Ratio



(D) Loan Repayment Status

Note: This figure reports employment and financial outcomes among student borrowers in the 2012 cohort as of 2017. Panel A reports realized salaries, including zeros for those who are unemployed or not in the labor force. Panel B reports mean degree completion and employment for all students in our sample, as well as the share of borrowers in our sample with no delinquencies. Panel C reports a histogram of monthly loan-payment-to-salary ratios among student borrowers who have begun the repayment period on their federal student loans. The "∞" bar represents the portion of borrowers who report not having employment in 2017. Panel D reports a pie chart of loan status among borrowers in repayment. Each portion of the pie represents the share of borrowers whose most severe non-repayment event since leaving college corresponds to the labeled status. For example, those who are in default are delinquent but are counted as "Default" in the chart above. Sample and variable definitions are provided in Table 1. Source: U.S. Department of Education, National Center for Education Statistics, 2012/17 Beginning Postsecondary Students (BPS) study, authors' calculations (September 2020).

Figure 3: Realizations Versus Elicitations

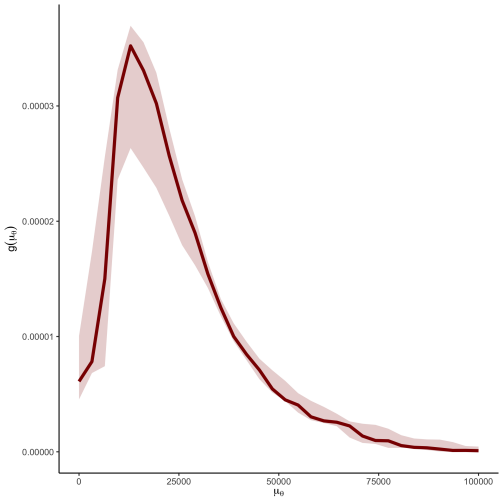


(A) Salary **(B) Completion**

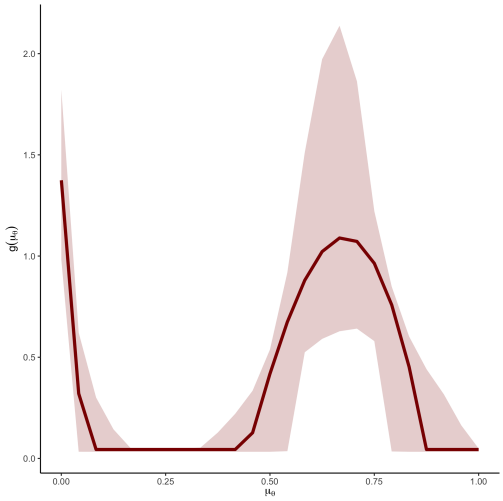
(C) Employment **(D) On-Time Repayment**

Note: This figure plots realized outcomes against subjective elicitations asked in the 2012 survey. Panels A through C report binned scatter plots. Panel A reports the relationship between log salary in 2017 against an individual's log of expected salary. Panel B reports reports the likelihood of completing college against the elicited likelihood of on-time completion. Panel C reports the likelihood of being employed against the log salary the respondent would expect if they were not enrolled in college. Panel D reports average loan repayment by respondents' responses when asked whether they agree with the statement, "My parents encourage me to stay in college." Responses are coded as (1) "Strongly disagree," (2) "Somewhat disagree," (3) "Neither disagree nor agree," (4) "Somewhat agree," and (5)"Strongly agree." Grey bubbles reflect relative number of individuals reporting each response. In all four panels, dotted lines denote linear OLS predictions. Source: U.S. Department of Education, National Center for Education Statistics, 2012/17 Beginning Postsecondary Students (BPS) study, authors' calculations (September 2020).

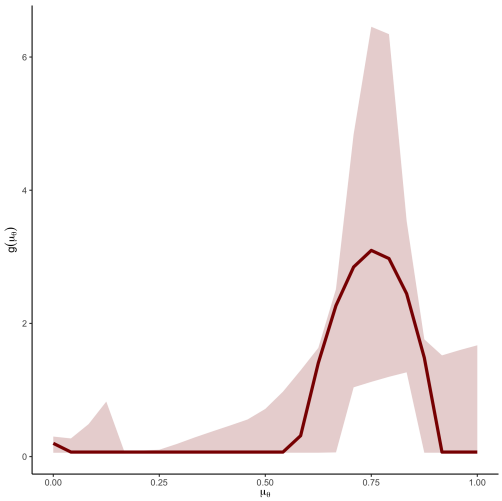
Figure 4: Estimates of Belief Distributions



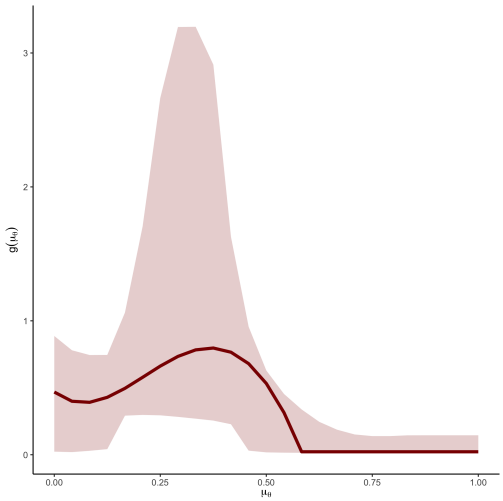
(A) Earnings Equity



(B) Completion-Contingent Loan



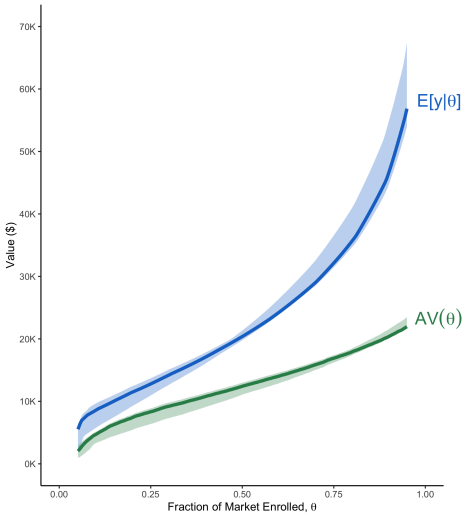
(C) Employment-Contingent Loan



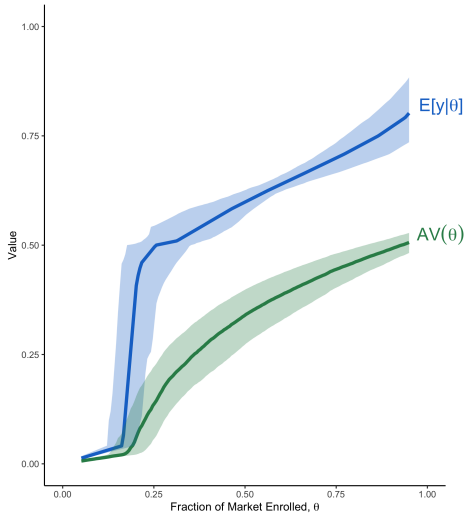
(D) Dischargeable Loan

Note: This figure plots the distribution of $E[y|\theta]$ conditional on public information for each of the outcomes corresponding to the four markets of interest. Public information controls include all variables from the institutional and academic categories defined in Appendix Table A1. Panel A presents the results for the equity market, Panel B presents the results for state-contingent debt market with repayment only in the event of college completion, Panel C presents the results for the state-contingent debt market with repayment only in the event of employment, and Panel D presents the results for the dischargeable loan market requiring repayment only if not delinquent on traditional student loans. The shaded region presents 95% confidence intervals constructed via bootstrap resampling. Source: U.S. Department of Education, National Center for Education Statistics, 2012/17 Beginning Postsecondary Students (BPS) study, authors' calculations (September 2020).

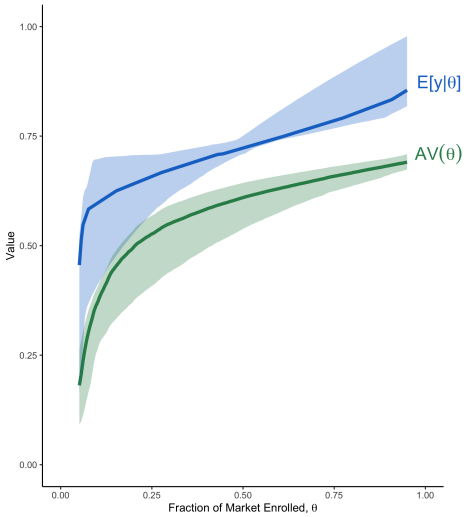
Figure 5: Estimates of Average Value



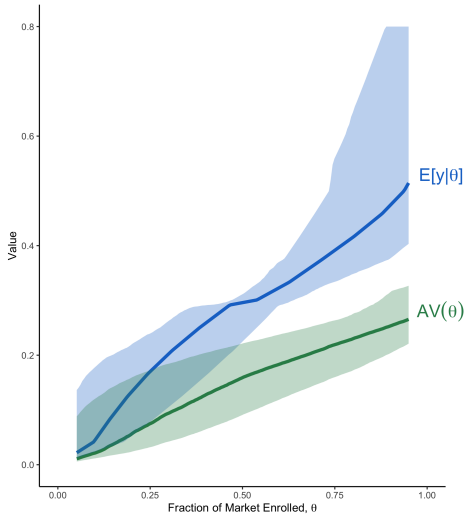
(A) Earnings Equity



(B) Completion-Contingent Loan



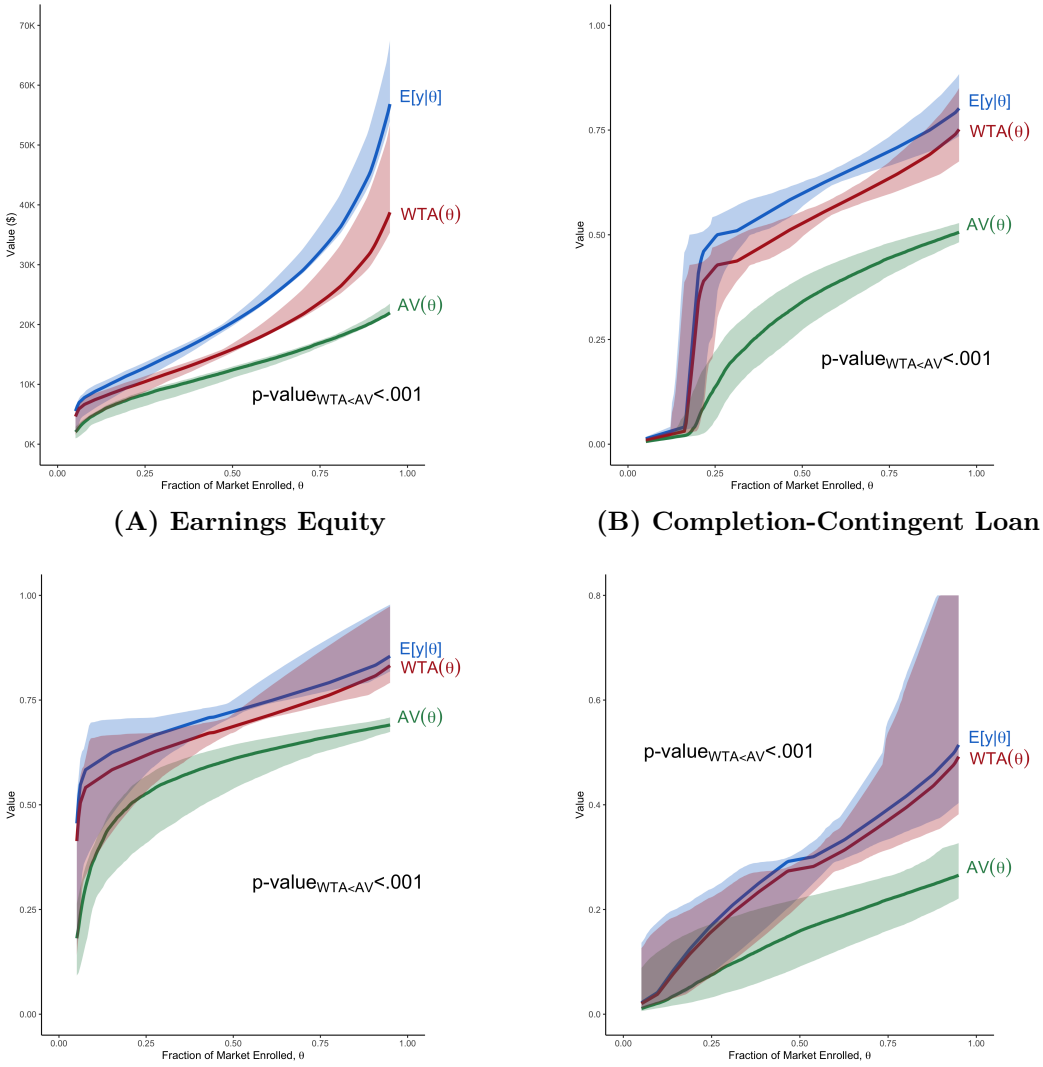
(C) Employment-Contingent Loan



(D) Dischargeable Loan

Note: This figure plots the average value curve, $AV(\theta)$, and expected outcomes, $E[y|\theta]$, for each of the four markets of interest. We plot each curve against the fraction of the market insured, θ, on the horizontal axis. The blue line plots percentiles of privately-expected salaries, $E[y|\theta]$. The green line presents the average value curve, $AV(\theta)$. Panel A presents the results for the equity market, Panel B presents the results for state-contingent debt market with repayment only in the event of college completion, Panel C presents the results for the state-contingent debt market with repayment only in the event of employment, and Panel D presents the results for the dischargeable loan market requiring repayment only if not delinquent on traditional student loans. Results are conditional on institutional and academic categories of public information, as defined in Appendix Table A1. The shaded region presents 95% confidence intervals constructed via bootstrap resampling. Source: U.S. Department of Education, National Center for Education Statistics, 2012/17 Beginning Postsecondary Students (BPS) study, authors' calculations (September 2020).

Figure 6: Estimates of Average Value and Willingness-to-Accept



(A) Earnings Equity

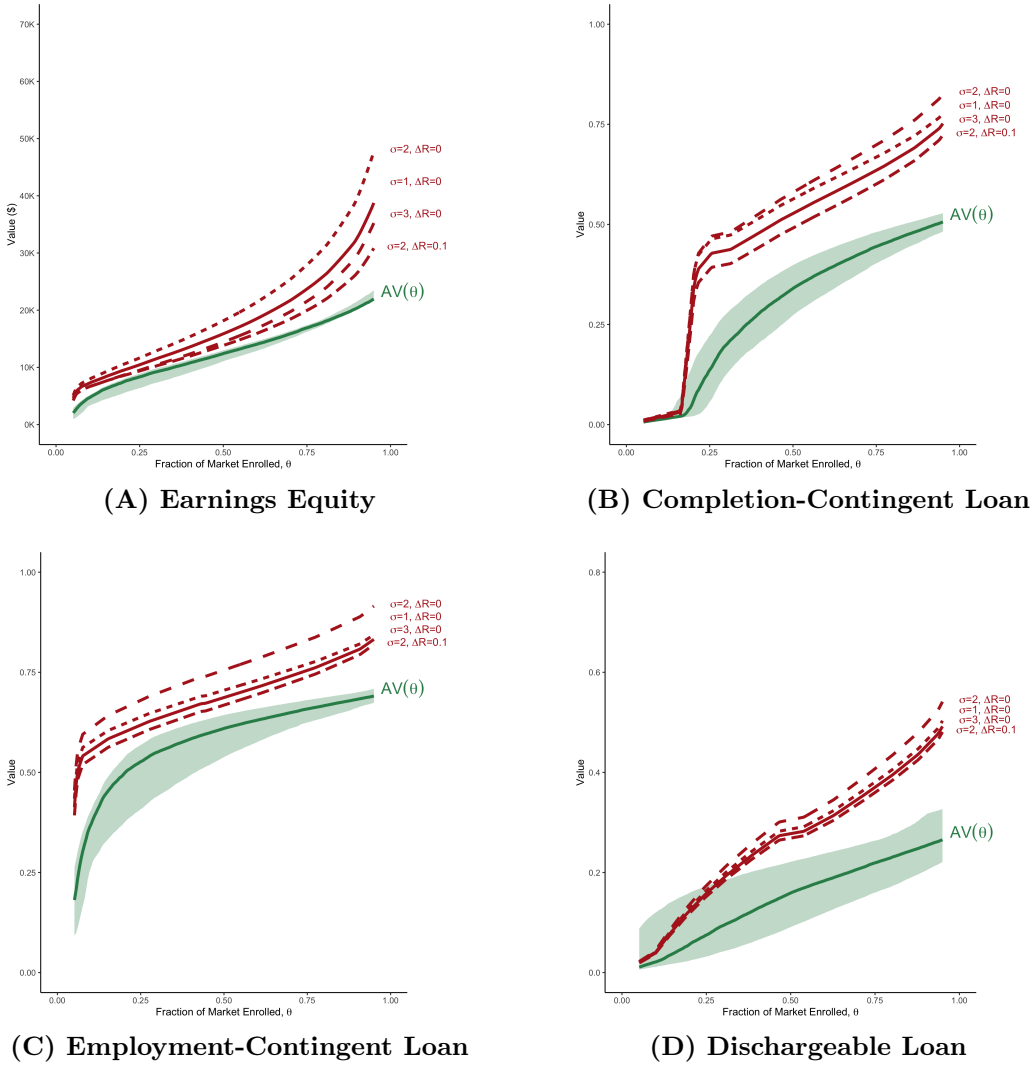
(B) Completion-Contingent Loan

(C) Employment-Contingent Loan

(D) Dischargeable Loan

Note: This figure plots the average value curve, $AV(\theta)$, and the willingness-to-accept curve, $WTA(\theta)$, for each of the four markets of interest. We plot each curve against the fraction of the market insured, θ, on the horizontal axis. The blue line plots percentiles of privately-expected salaries, $E[y|\theta]$. The green line presents the average value curve, $AV(\theta)$. The red line presents the willingness-to-accept curve, $WTA(\theta)$. Panel A presents the results for the equity market, Panel B presents the results for state-contingent debt market with repayment only in the event of college completion, Panel C presents the results for the state-contingent debt market with repayment only in the event of employment, and Panel D presents the results for the dischargeable loan market requiring repayment only if not delinquent on traditional student loans. Results are conditional on institutional and academic categories of public information, as defined in Appendix Table A1. The shaded region presents 95% confidence intervals constructed via bootstrap resampling. We also present the p-value for a test of the market unraveling condition in equation (8), which is given by the fraction of bootstrap draws for which there exists a value of θ such that $WTA(\theta) < AV(\theta)$. Note that this p-value accounts for correlated sampling error between the $WTA(\theta)$ and $AV(\theta)$ curves. Source: U.S. Department of Education, National Center for Education Statistics, 2012/17 Beginning Postsecondary Students (BPS) study, authors' calculations (September 2020).

Figure 7: Estimates of Average Value and Alternative Willingness-to-Accept Specifications

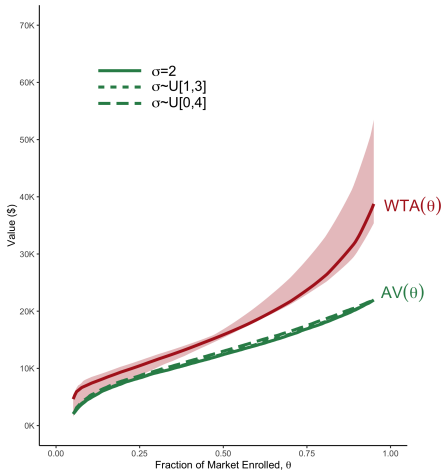


(A) Earnings Equity **(B) Completion-Contingent Loan**

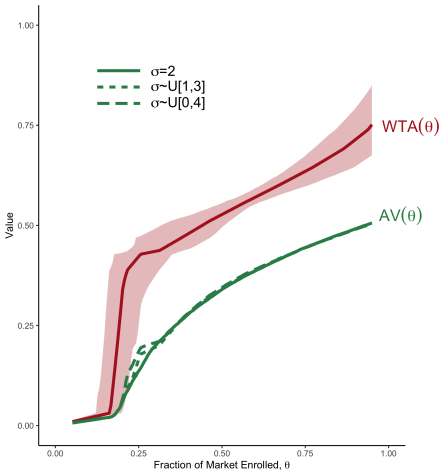
(C) Employment-Contingent Loan **(D) Dischargeable Loan**

Note: This figure plots alternative specifications for the willingness-to-accept curve, $WTA(\theta)$, for different values of the coefficient of relative risk aversion, σ, and assumptions about the difference between the interest rate faced by financiers and the implicit interest rate rationalizing the Euler equation of college-goers (ΔR). We plot each curve against the fraction of the market insured, θ, on the horizontal axis. For reference, the green line presents the average value curve, $AV(\theta)$, from the baseline specification. The solid red line presents the willingness-to-accept curve, $WTA(\theta)$, from the baseline specification. The three dashed red lines present alternative specifications for $WTA(\theta)$ using $\sigma = 1$ and $\sigma = 3$, and an alternative specification assuming college-goers face a 10pp higher implicit interest rate than financiers, $\Delta R = 0.10$. Source: U.S. Department of Education, National Center for Education Statistics, 2012/17 Beginning Postsecondary Students (BPS) study, authors' calculations (September 2020).

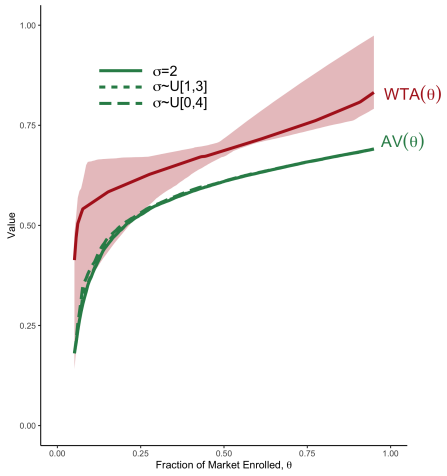
Figure 8: Estimates of Average Value and Willingness-to-Accept under Preference Heterogeneity



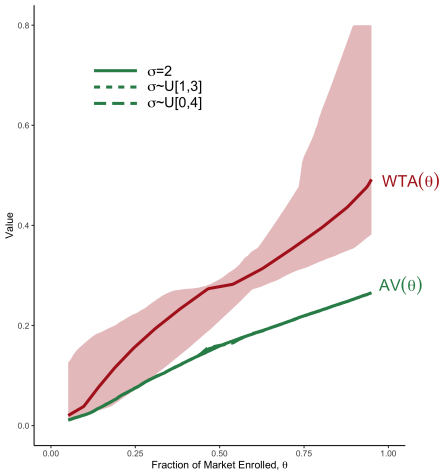
(A) Earnings Equity



(B) Completion-Contingent Loan



(C) Employment-Contingent Loan



(D) Dischargeable Loan

Note: This figure compares average value and willingness-to-accept under alternative specifications that allow for heterogeneity in risk aversion, σ. The red line presents the quantiles of the willingness to accept from the baseline specification. The solid, dotted, and dashed green lines present average value curves, $AV(\theta)$, under each alternative specification. The $AV(\theta)$ curves using equation (5) as the average value of y for those who have a lower willingness to accept than the plotted value of the willingness to accept curve. For ease of comparison, the figure holds the levels of the $WTA(\theta)$ curve fixed from the baseline specification when computing the AV curve. This allows the figure to illustrate the no trade condition relative to a single standardized $WTA(\theta)$ curve, but the fraction of the market taking up the contract differs slightly from θ across specifications. For comparison, Appendix Figure A3 presents the WTA and AV curves using the exact quantiles of the willingness to accept curve in the alternative specification. Source: U.S. Department of Education, National Center for Education Statistics, 2012/17 Beginning Postsecondary Students (BPS) study, authors' calculations (September 2020).

Table 1: Summary Statistics: Elicitations and Realizations

Category	Variable	Mean	SD
Panel A: Ex-Ante Elicitations	Ever Completion Likelihood	9.314	1.838
	On-Time Completion Likelihood	8.413	2.103
	Expected Completion Year	2014.3	1.091
	Employment Likelihood	8.154	1.734
	Exp. Occ. Employed	0.847	0.0937
	Expected Salary	64064.2	44800.8
	Highest Expected Salary	117110.8	142762.8
	Lowest Expected Salary	43923.5	26926.0
	Expected Salary if No College	17332.5	7823.6
	Exp. Occ. Salary	30073.1	8503.5
	Elicited Discount Factor	0.370	0.321
	Supportive Friends	4.372	0.971
	Supportive Classmates	4.226	1.071
	Supportive Parents	4.227	1.072
	Parent Financial Support	6463.8	9512.1
Panel B: Ex-Post Outcomes	Completed Degree	0.515	0.500
	On-Time Repayment	0.310	0.462
	Delinquent	0.620	0.485
	Default	0.165	0.371
	Employed	0.735	0.441
	Unemployed	0.123	0.328
	Realized Salary	32701.5	24345.6
	Number of Credit Cards	1.051	0.816
	Credit Card Balance	1234.9	3171.3
	Paid Credit Card Balance	0.604	0.489

Note: This table provides summary statistics for the complete set of outcomes and elicitations used in our non-parametric deconvolution, and maximum-likelihood exercises. Data are taken from the 2012-2017 Beginning Post-secondary Students (BPS) study. Elicitations are measured in winter and spring of 2012. Outcomes are measured in the spring of 2017. "Completed Degree" indicates whether the respondent had completed their intended degree as of June 2017. "Non-Repayment" indicates whether the respondent reported being in default, delinquency, or forbearance on their student loans at least once since beginning repayment. "Employed" indicates whether the respondent reported holding a job at some point between February and June of 2017. "Unemployed" indicates whether the respondent was not employed and looking for work for one or more months since leaving college. "Realized Salary" is the respondent's reported salary for their most recently held job since February 2017, excluding those without jobs. "Number of Credit Cards" and "Credit Card Balance" provides the self-reported total number and monthly balance on credit cards among respondents who held credit cards in 2017. "Paid Credit Card Balance" indicates credit-card holders said they do not usually carry a balance month to month. Elicitations are defined in Appendix B. Sample size is 22,530 individuals, rounded to the nearest ten. Source: U.S. Department of Education, National Center for Education Statistics, 2012/17 Beginning Postsecondary Students (BPS) study, authors' calculations (September 2020).

Table 2: Summary Statistics: Public Information

Category	Variable	Mean	SD
Institution	Four-Year	0.540	0.498
	Private	0.299	0.458
	For-Profit	0.128	0.334
	Enrollment	18218.3	34962.9
	Tuition	9620.2	10939.2
	Share Female	0.573	0.123
	Share Black	0.138	0.163
	Admissions Rate	0.633	0.199
	Completion Rate	0.411	0.245
	Avg. SAT Score	1102.0	137.5
	Md. Parent Income	32142.5	20580.4
	Md. 6-Yr Earnings	29530.3	8106.7
Academic	Age	20.54	5.948
	BA Program	0.472	0.499
	STEM Major	0.182	0.386
Performance	High School GPA	3.058	0.613
	SAT Score	1008.7	203.3
Demographics	Female	0.565	0.496
	Black	0.176	0.381
	US Citizen	0.945	0.228
	Married	0.0585	0.235
	Children	0.121	0.326
Parental	Parent has BA	0.386	0.487
	Parents Married	0.661	0.473
	Dependent	0.783	0.412
	Parental Income	77816.3	73684.7
	EFC	10245.3	16865.8
Protected Classes	Black	0.176	0.381
	Female	0.565	0.496

Note: This table provides summary statistics for the for key public-information and demographic variables used in our non-parametric deconvolution, and maximum-likelihood exercises. All variables in this table are classified as public information in our various control specifications with the exception of gender and race (these are protected classes and cannot be used in pricing or screening for financial products). Sample size is 22,530 individuals, rounded to the nearest ten. Source: U.S. Department of Education, National Center for Education Statistics, 2012/17 Beginning Postsecondary Students (BPS) study, authors' calculations (September 2020).

Table 3: Presence of Private Information about Future Salary

	(1) Log Salary	(2) Log Salary	(3) Log Salary	(4) Log Salary	(5) Log Salary	(6) Log Salary
Log Expected Salary	0.113***	0.0602***	0.0446***	0.0432***	0.0327**	0.0314**
	(0.0159)	(0.0159)	(0.0161)	(0.0160)	(0.0158)	(0.0158)
Institution		X	X	X	X	X
Academic			X	X	X	X
Performance				X	X	X
Demographics					X	X
Parental						X
Partial R-Squared	0.009	0.067	0.101	0.104	0.119	0.123
R-squared	0.009	0.003	0.002	0.001	0.001	0.001
N	12580	12580	12580	12580	12580	12580

Note: This table reports coefficients and associated standard errors from a multivariate regression of log salary in 2017 on the log of 2012 elicited salary. Columns (1)–(6) include an increasing set of controls for observable information that are classified in Appendix Table A1. Column (1) includes no additional controls; Column (2) adds controls for the type of institution attended, Column (3) adds controls for academic information, Column (4) adds controls for high school performance, Column (5) adds controls for demographic information, and Column (6) adds controls for parental information. Number of observations are rounded to the nearest ten. Source: U.S. Department of Education, National Center for Education Statistics, 2012/17 Beginning Postsecondary Students (BPS) study, authors' calculations (September 2020).

Table 4: Presence of Private Information about Degree Completion

	(1) Degree Completion	(2) Degree Completion	(3) Degree Completion	(4) Degree Completion	(5) Degree Completion	(6) Degree Completion
On-Time Completion Likelihood	0.0492***	0.0365***	0.0364***	0.0345***	0.0343***	0.0332***
	(0.00223)	(0.00223)	(0.00224)	(0.00225)	(0.00221)	(0.00220)
Institution		X	X	X	X	X
Academic			X	X	X	X
Performance				X	X	X
Demographics					X	X
Parental						X
Partial R-Squared	0.045	0.215	0.222	0.239	0.249	0.264
R-squared	0.045	0.029	0.028	0.028	0.028	0.026
N	22340	22340	22340	22340	22340	22340

Note: This table reports coefficients and associated standard errors from a multivariate regression of degree completion by 2017 on the 2012 elicited likelihood of on-time completion. Columns (1)-(6) include an increasing set of controls for observable information that are classified in Appendix Table A1. Column (1) includes no additional controls; Column (2) adds controls for the type of institution attended, Column (3) adds controls for academic information, Column (4) adds controls for high school performance, Column (5) adds controls for demographic information, and Column (6) adds controls for parental information. Number of observations are rounded to the nearest ten. Source: U.S. Department of Education, National Center for Education Statistics, 2012/17 Beginning Postsecondary Students (BPS) study, authors' calculations (September 2020).

Table 5: Presence of Private Information about Future Employment

	(1) Employed	(2) Employed	(3) Employed	(4) Employed	(5) Employed	(6) Employed
Log Expected Salary if No College	0.0313***	0.0243**	0.0212**	0.0199*	0.0175	0.0169
	(0.0107)	(0.0109)	(0.0108)	(0.0107)	(0.0106)	(0.0106)
Institution		X	X	X	X	X
Academic			X	X	X	X
Performance				X	X	X
Demographics					X	X
Parental						X
Partial R-Squared	0.012	0.026	0.035	0.038	0.042	0.046
R-squared	0.012	0.008	0.007	0.007	0.006	0.006
N	17480	17480	17480	17480	17480	17480

Note: This table reports coefficients and associated standard errors from a multivariate regression of employment status in 2017 on the 2012 log elicited salary the respondent would expect if they were not employed. Columns (1)-(6) include an increasing set of controls for observable information that are classified in Appendix Table A1. Column (1) includes no additional controls; Column (2) adds controls for the type of institution attended, Column (3) adds controls for academic information, Column (4) adds controls for high school performance, Column (5) adds controls for demographic information, and Column (6) adds controls for parental information. Number of observations are rounded to the nearest ten. Source: U.S. Department of Education, National Center for Education Statistics, 2012/17 Beginning Postsecondary Students (BPS) study, authors' calculations (September 2020).

Table 6: Presence of Private Information about On-Time Loan Repayment

	(1) On-Time Repayment	(2) On-Time Repayment	(3) On-Time Repayment	(4) On-Time Repayment	(5) On-Time Repayment	(6) On-Time Repayment
Supportive Parents	0.0635***	0.0349***	0.0336***	0.0305***	0.0301***	0.0285***
	(0.00505)	(0.00502)	(0.00497)	(0.00491)	(0.00488)	(0.00483)
Institution		X	X	X	X	X
Academic			X	X	X	X
Performance				X	X	X
Demographics					X	X
Parental						X
Partial R-Squared	0.030	0.114	0.123	0.136	0.144	0.155
R-squared	0.030	0.014	0.014	0.015	0.015	0.014
N	15520	15520	15520	15520	15520	15520

Note: This table reports coefficients and associated standard errors from a multivariate regression of the incidence of on-time repayment as of 2017 on the 2012 elicited level of parents' emotional support. We restrict the sample to those who utilize student loans for college financing. Columns (1)-(6) include an increasing set of controls for observable information that are classified in Appendix Table A1. Column (1) includes no additional controls; Column (2) adds controls for the type of institution attended, Column (3) adds controls for academic information, Column (4) adds controls for high school performance, Column (5) adds controls for demographic information, and Column (6) adds controls for parental information. Number of observations are rounded to the nearest ten. Source: U.S. Department of Education, National Center for Education Statistics, 2012/17 Beginning Postsecondary Students (BPS) study, authors' calculations (September 2020).

Table 7: Lower-Bound on the Magnitude of Private Information

	Category				
	(1) *No Public Info*	(2) *Institution + Academic*	(3) *Institution + Academic + Performance + Demographics*	(4) *Institution + Academic + Performance + Demographics + Parental*	(5) *Institution + Academic + Performance + Demographics + Parental + Protected*
Earnings Equity	5765	5314	3797	2907	2381
Completion-Contingent Loan	0.20	0.16	0.13	0.11	0.11
Employment-Contingent Loan	0.09	0.11	0.07	0.05	0.04
Dischargeable Loan	0.13	0.13	0.07	0.05	0.04

Note: This table provides estimates of $E\left[r\left(X_i,Z_i\right)\right]$, the lower-bound on the average difference between the average value curve, $AV\left(\theta\right)$, and the expected outcome, $E\left[y|\theta\right]$, for each of our four contracts. Values are calculated from equation 12 using random-forest estimates of $E\left[y|X_i,Z_i\right]$ and $E\left[y|X_i\right]$. X_i includes the set of publicly known variables corresponding to each column label. Column (1) includes no controls for observable variables. Column (2) adds controls for the type of institution attended and academic information. Column (3) adds controls for high school performance and demographic information. Column (4) adds controls for parental information. Column (5) adds information on race and gender. Z_i includes all private elicitations in Table 1, as well as any observable variables not included in the specified set of public information, X. These categories are defined in Table 2. Source: U.S. Department of Education, National Center for Education Statistics, 2012/17 Beginning Postsecondary Students (BPS) study, authors' calculations (September 2020).

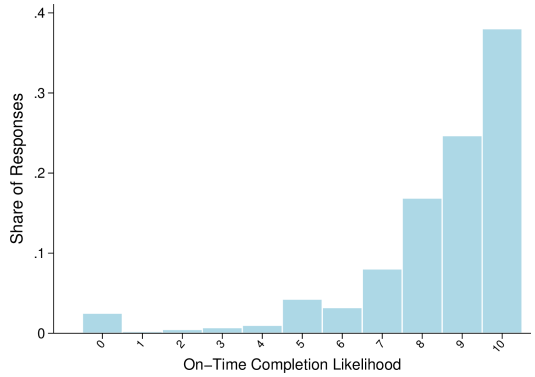
Table 8: MVPF Components

	(1) Take-up	(2) Transfer	(3) Smoothing	(4) WTP	(5) FE Grant	(6) FE Tax Distortion	(7) Cost	(8) MVPF
Earnings Equity	0.79	0.30	0.17	0.47	0.09	-0.04	0.25	1.86
	(0.02)	(0.04)	(0.01)	(0.03)	(0.00)	(0.00)	(0.04)	(0.15)
Completion-Contingent Loan	0.52	0.31	0.10	0.41	0.09	-0.13	0.35	1.16
	(0.01)	(0.03)	(0.00)	(0.03)	(0.00)	(0.00)	(0.03)	(0.03)
Employment-Contingent Loan	0.56	0.11	0.05	0.17	0.10	-0.10	0.12	1.42
	(0.03)	(0.05)	(0.00)	(0.05)	(0.00)	(0.00)	(0.05)	(0.11)
Dischargeable Loan	0.44	0.73	0.02	0.75	0.08	-0.30	0.94	0.79
	(0.09)	(0.13)	(0.01)	(0.12)	(0.01)	(0.01)	(0.13)	(0.02)
Grant	1.00	1.00	0.00	1.00	0.15	-0.00	0.85	1.17
	–	–	–	–	–	–	–	–

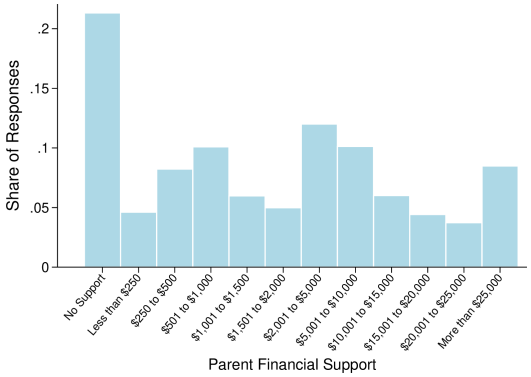
Note: This table reports components of the marginal value of public funds (MVPF), defined in Section 8. Components are reported for each of four hypothetical contracts: salary-based equity contract (row 1), and state-contingent debt contracts that are dischargeable in the event of dropout (row 2), non-employment (row 3), and non-repayment (row 4). For each contract, the MVPF is calculated for the actuarially fair price under no private information, $\lambda = \frac{1}{E[y]}$ and $\kappa = \frac{1}{E[y]}$, so that the government would break even if there was no differential selection into the contract. Column (1) reports the "Take-up", which denotes the share of individuals who would accept the contract, column (2) reports the size of the "Transfer", which equals the average expected surplus contractees would receive (i.e., expected negative profits the financier would incur). Column (3) reports the consumption "Smoothing" benefits individuals derive from the contract. Column (4) reports the willingness to pay by those who choose to take up the contract, which is the sum of the size of the transfer and consumption smoothing benefits. Columns (5)–(6) turn to the components of costs that arise from fiscal externalities from behavioral responses to the financing. Column (5) reports the size of the fiscal externality resulting from the provision of the education finance, "FE Grant". Column (6) reports the fiscal externality from the distortion associated with the implicit tax on earnings associated with the risk-mitigating contracts. Column (7) measures total cost, which equals the size of the transfer minus the two fiscal externality terms. Column (8) reports the MVPF, which is the ratio of WTP in Column (4) to net government Cost in Column 7. Source: U.S. Department of Education, National Center for Education Statistics, 2012/17 Beginning Postsecondary Students (BPS) study, authors' calculations (September 2020).

Appendix A Additional Figures and Tables

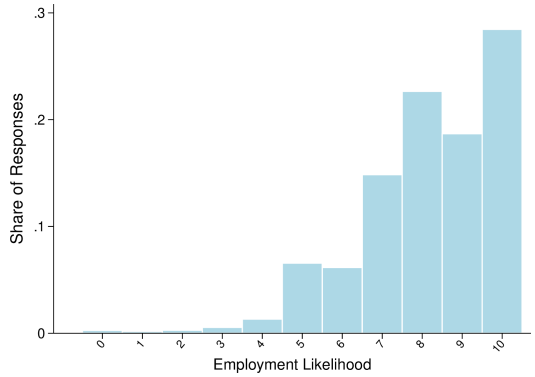
Figure A1: Elicitations



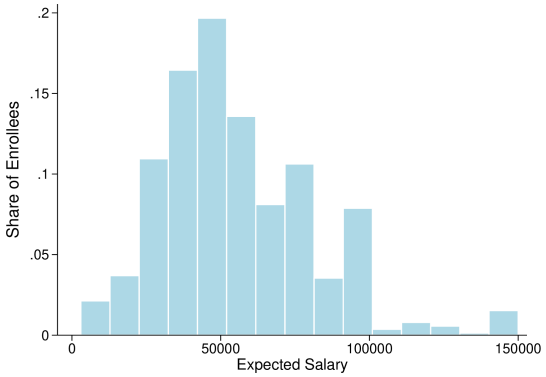
(A) Completion Likelihood



(B) Parents' Financial Support



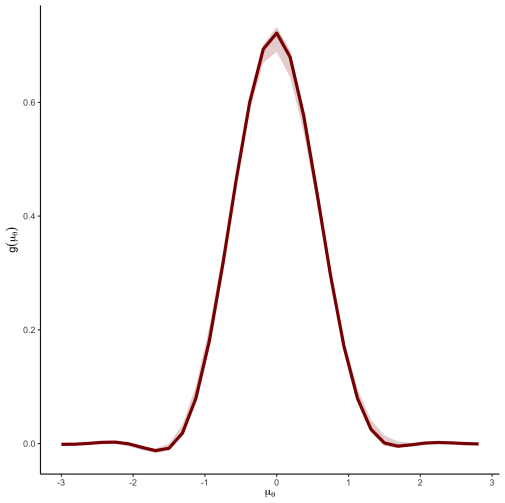
(C) Employment Likelihood



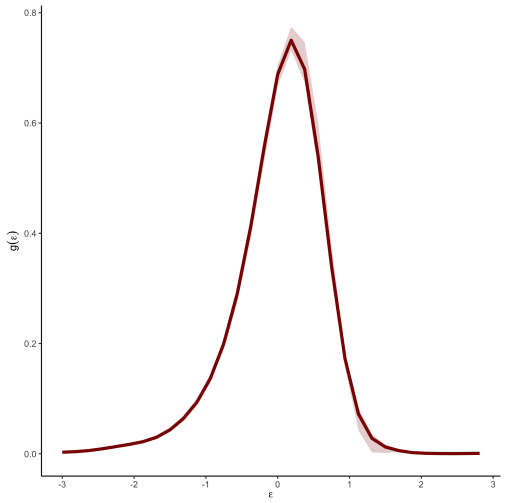
(D) Expected Salary

Note: This figure plots histograms for four elicitation variables: completion likelihood, parents' financial support, employment likelihood, and expected salary. Elicitations are defined in Appendix B. Sample details are provided in Table 1. Source: U.S. Department of Education, National Center for Education Statistics, 2012/17 Beginning Postsecondary Students (BPS) study, authors' calculations (September 2020).

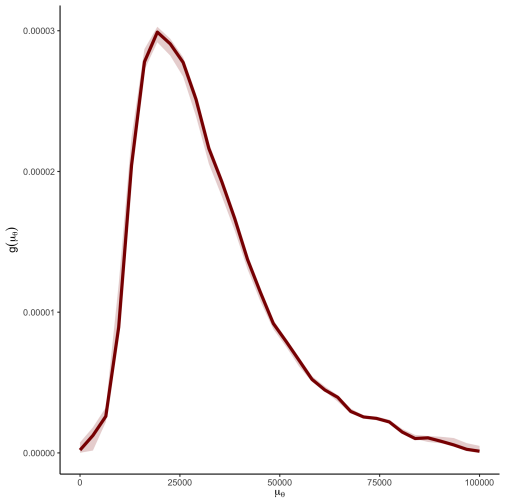
Figure A2: Components of the Distribution of Expected Salary



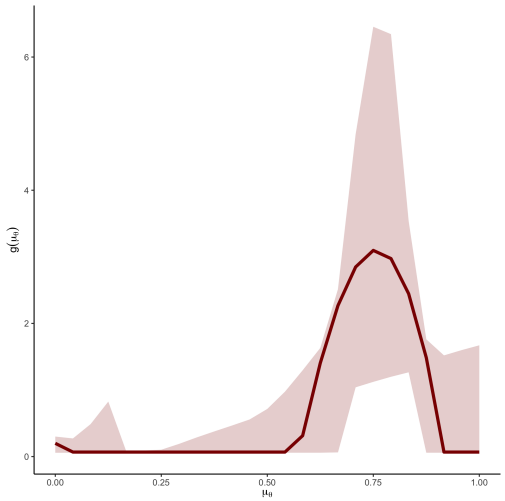
(A) Expected Log Salary if Employed, $\tilde{\theta}$



(B) Log Income Uncertainty Realization, ϵ



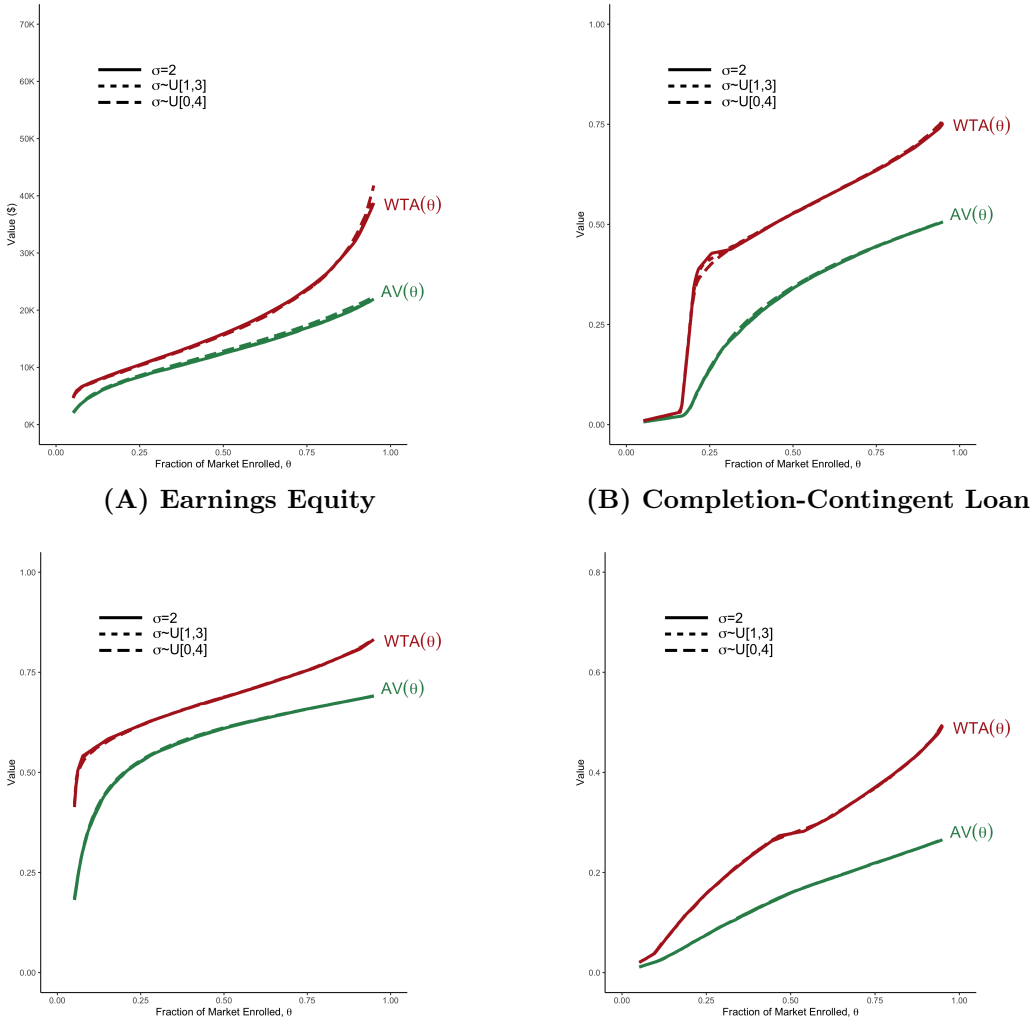
(C) Expected Salary if Employed, $E[y|y > 0,\theta]$



(D) Probability of Employment, $Pr\left(y > 0|\theta\right)$

Note: This figure plots the components of the distribution of expected salary, $E\left[y|\theta\right]$, after residualizing on public information that includes all variables from the institutional and academic categories defined in Appendix Table A1. Panel A presents the p.d.f. of the distribution of each type, θ, expected log salary. Panel B presents the distribution of log income uncertainty, ϵ, that equals the unknown component of future income. Panel C presents the distribution of expected salary if employed. Panel D presents the distribution of the probability of employment. The shaded region presents 95% confidence intervals constructed via bootstrap resampling. Source: U.S. Department of Education, National Center for Education Statistics, 2012/17 Beginning Postsecondary Students (BPS) study, authors' calculations (September 2020).

Figure A3: Estimates of Average Value and Willingness-to-Accept under Preference Heterogeneity



(A) Earnings Equity

(B) Completion-Contingent Loan

(C) Employment-Contingent Loan

(D) Dischargeable Loan

Note: This figure compares average value and willingness-to-accept under alternative specifications that allow for heterogeneity in risk aversion, σ. We plot each curve against the fraction of the market insured. The solid, dotted, and dashed lines present average value curves, $AV(\theta)$ in green, and willingness-to-accept curves, $WTA(\theta)$ in red, under each alternative specification. Source: U.S. Department of Education, National Center for Education Statistics, 2012/17 Beginning Postsecondary Students (BPS) study, authors' calculations (September 2020).

Table A1: Categorization and Description of Publicly Observable Information, X

Category	Variable
Institutional Characteristics	Four-Year College
	Private/Public Status
	For-Profit
	Region (8 Categories)
	Enrollment Size
	Share Black
	Share Female
	Admissions Rate
	Completion Rate
	Average SAT Score
	Median Parental Income
	Median 6-Year Salary
Academic Characteristics	Age at Enrollment
	Type of Degree (BA, AA)
	Field of Study (14 Categories)
Performance	High School GPA
	SAT Score
Demographics	Citizenship Status
	Marital Status
	Number of Dependents
Parental Characteristics	Parents' Highest Education
	Parents' Marital Status
	Students' Dependency Status
	Parents' Income
	Expected Family Contribution (FAFSA)
Protected Classes	Race
	Gender

Note: This table lists names and categories for all variables used as observable characteristics in our analysis. The right column provides the variable name. The left column provides category names for each group of variables. More detailed variable definitions can be found at the National Center for Education Statistics, 2012/17 Beginning Postsecondary Students (BPS) study website: `https://nces.ed.gov/surveys/bps/`.

Table A2: Predictive Performance With and Without Elicitations

Outcome	Statistic	Category (1) *Institution + Academic*	(2) *Institution + Academic + Performance + Demographics*	(3) *Institution + Academic + Performance + Demographics + Parental*	(4) *Institution + Academic + Performance + Demographics + Parental + Protected*	(5) *All Public + Elicitations*
Panel A: Log Salary	R^2	0.068	0.073	0.078	0.092	0.108
		(0.009)	(0.010)	(0.010)	(0.010)	(0.011)
	RMSE	0.641	0.638	0.636	0.631	0.626
		(0.012)	(0.013)	(0.012)	(0.012)	(0.012)
	MAE	0.464	0.461	0.460	0.455	0.453
		(0.007)	(0.007)	(0.007)	(0.007)	(0.007)
Panel B: Dropout	Pseudo R^2	0.096	0.157	0.166	0.170	0.231
		(0.013)	(0.008)	(0.007)	(0.007)	(0.007)
	ROC	0.742	0.761	0.768	0.770	0.813
		(0.006)	(0.006)	(0.006)	(0.006)	(0.005)
	Accuracy	0.684	0.697	0.701	0.704	0.741
		(0.006)	(0.006)	(0.006)	(0.006)	(0.005)
Panel C: On-Time Repayment	Pseudo R^2	0.060	0.133	0.155	0.158	0.170
		(0.014)	(0.011)	(0.010)	(0.010)	(0.009)
	ROC	0.723	0.758	0.773	0.775	0.785
		(0.008)	(0.008)	(0.008)	(0.008)	(0.007)
	Accuracy	0.755	0.763	0.761	0.763	0.766
		(0.006)	(0.006)	(0.006)	(0.006)	(0.006)
Panel D: Employment	Pseudo R^2	-0.110	0.002	0.021	0.027	0.042
		(0.022)	(0.007)	(0.006)	(0.006)	(0.005)
	ROC	0.565	0.596	0.610	0.621	0.640
		(0.009)	(0.009)	(0.009)	(0.009)	(0.009)
	Accuracy	0.700	0.719	0.719	0.721	0.723
		(0.006)	(0.006)	(0.006)	(0.006)	(0.006)

Note: This table reports out-of-sample prediction performance statistics for each outcome. Each column corresponds to an increasing set of predictor variables that are included in a random forest model trained on a 70% sample. Column (1) includes institution and academic variables. Column (2) adds performance and demographics. Column (3) adds parental characteristics. Column (4) adds information on race and gender. Each of these categories is defined in Appendix Table A1. Finally, column (5) adds in the elicitations. Numbers in parentheses denote standard deviations of prediction statistics calculated over 1000 bootstrap samples of the 30% holdout sample. Pseudo-R^2 is calculated as $1 - \frac{\ln L_M}{\ln L_0}$, where L_M and L_0 denote the likelihood of observed outcomes given predictions from the random forest model and sample mean, respectively. Source: U.S. Department of Education, National Center for Education Statistics, 2012/17 Beginning Postsecondary Students (BPS) study, authors' calculations (September 2020).

Table A3: Lower-Bound on the Magnitude of Private Information, Excluding Observables from Private Information

	Category				
	(1) *No Public Info*	(2) *Institution + Academic*	(3) *Institution + Academic + Performance + Demographics*	(4) *Institution + Academic + Performance + Demographics + Parental*	(5) *Institution + Academic + Performance + Demographics + Parental + Protected*
Earnings Equity	4881	4163	2903	2502	2216
Completion-Contingent Loan	0.22	0.15	0.12	0.11	0.11
Employment-Contingent Loan	0.12	0.09	0.05	0.05	0.04
Dischargeable Loan	0.12	0.11	0.06	0.04	0.04

Note: This table provides lower-bound estimates, $E\left[r\left(X_i,Z_i\right)\right]$, under the assumption that private information, Z_i, includes all private elicitations in Table 1 and no other observable characteristics. Values are calculated from equation 12 using random-forest estimates of $E\left[y|X_i,Z_i\right]$ and $E\left[y|X_i\right]$. X_i includes the set of publicly known variables corresponding to each column label. Column (1) includes no controls for observable variables. Column (2) adds controls for the type of institution attended and academic information. Column (3) adds controls for high school performance and demographic information. Column (4) adds controls for parental information. Column (5) adds information on race and gender. These categories are defined in Table 2. Source: U.S. Department of Education, National Center for Education Statistics, 2012/17 Beginning Postsecondary Students (BPS) study, authors' calculations (September 2020).

Table A4: Elicitation Details and γ-Estimates

(1) Outcome	(2) Elicitation	(3) Instrument	(4) γ-Estimate
Salary	*Log Expected Salary*	*Log Avg. Salary Expected Occ.*	0.69
			(0.16)
Completion	*On-Time Completion Likelihood*	*Supportive Parents*	3.20
			(0.23)
Employment	*Log Expected Salary if No College*	*Avg. Employment Expected Occ.*	0.59
			(0.29)
On-Time Repayment	*Supportive Parents*	*Parents' Financial Support*	1.47
			(0.76)

Note: This table summarizes the specifications used for each outcome in our IV estimation of the elicitation-belief relationship, γ, in equations (16) and (20) of the text. Column (1) lists the names of the outcome variables, y. Column (2) lists the names of the focal elicitations, z, used as dependent variables. Column (3) lists the names of instrumental variables, z', used to instrument for z in each regression. Column (4) reports point estimates of γ for each outcome-elicitation pair. Standard errors are reported in parentheses. Full elicitation descriptions are provided in Appendix B. Source: U.S. Department of Education, National Center for Education Statistics, 2012/17 Beginning Postsecondary Students (BPS) study, authors' calculations (September 2020).

Table A5: Elicitation Details and γ-Estimates: Alternative Specification

(1) Outcome	(2) Elicitation	(3) Alternative Instrument	(4) γ-Estimate
Salary	*Log Expected Salary*	*Log Expected Salary if No College*	0.77
			(0.10)
Completion	*On-Time Completion Likelihood*	*Parents' Financial Support*	2.04
			(0.23)
Employment	*Log Expected Salary if No College*	*Likelihood Employed in Expected Occ.*	1.34
			(0.62)
On-Time Repayment	*Supportive Parents*	*Avg. Employment Expected Occ.*	1.86
			(0.70)

Note: This table summarizes the alternative specifications used for each outcome in our secondary IV estimation of the elicitation-belief relationship, γ. Column (1) lists the names of the outcome variables, y. Column (2) lists the names of the focal elicitations, z, used as dependent variables. Column (3) lists the names of instrumental variables used to instrument for z in each regression. Column (4) reports point estimates of γ for each outcome-elicitation pair. Standard errors are reported in parentheses. Full elicitation descriptions are provided in Appendix B. Source: U.S. Department of Education, National Center for Education Statistics, 2012/17 Beginning Postsecondary Students (BPS) study, authors' calculations (September 2020).

Appendix B Descriptions of Elicitation Variables

The elicitations variables we use are the recorded responses to first-wave survey questions from the 2012/17 Beginning Postsecondary Students (BPS) study. The question text corresponding to each elicitation is provided below.

- *Expected Occupation*: "What is the title of the job you want to have after you complete your education?" [Response options correspond to 2010-13 Occupational Information Network-Standard Occupational Classification (O*NET-SOC) codes.]
- *Expected Salary*: "We have some questions about the range of salary you expect to make once you finish your education. What is...your expected yearly salary?"
- *Likelihood Employed in Expected Occupation*: "On a scale from 0-10, how likely do you think it is that, five years from now you will hold your intended occupation?"
- *On-Time Completion Likelihood*: "On a scale from 0-10, how likely is it you will finish your degree by [EXPECTED DATE]?"
- *Supportive Parents*: "On a scale of 1-5, how much do you agree with the following statement: 'My parents encourage me to stay in college.'?"
- *Expected Salary if No College*: "How much do you think you would have earned at all your jobs together if you had not attended college in the 2011-2012 school year?"
- *Parents' Financial Support*: "Through the end of the 2011-2012 school year, about how much will your parents (or guardians) have helped you pay for any of your education and living expenses while you are enrolled in school?"

More information on the survey design and implementation can be found at `https://nces.ed.gov/surveys/bps/`.

In addition to the elicitations above, we construct two additional Z-variables—*Log Average Salary in Expected Occupation* and *Average Employment in Expected Occupation*—using responses

to the *Expected Occupation* question. Specifically, for each individual i, we take averages of outcomes among college graduates (j) who had worked in individual i's expected occupation (occ_i) as of the BPS 2012 survey:

$$\textit{Log Avg. Salary Expected Occ.} = \log \frac{1}{N^{BB}_{occ_i}} \sum_{j \in occ_i} y^{BB}_j \tag{28}$$

$$\textit{Avg. Employment Expected Occ.} = \frac{1}{N^{BB}_{occ_i}} \sum_{j \in occ_i} e^{BB}_j. \tag{29}$$

Post-graduate salaries and employment (y^{BB}_j and e^{BB}_j), and cell-sizes ($N^{BB}_{occ_i}$) are taken from the 2008/2012 Baccalaureate and Beyond (B&B) study, which we match to BPS occupation elicitations (occ_i) using three-digit occupation codes. The B&B data include survey responses for a representative sample of four-year college graduates in the spring of 2008, followed up in 2011-2012. More information can be found at `https://nces.ed.gov/surveys/b&b/`.

Appendix C Deconvolution Details

Bonhomme and Robin (2010) deconvolve linear independent multi-factor models of the form $Y = AX$, where Y is a vector of observed measurements, X is a vector of latent variables, and A is a matrix of factor loadings, assumed to be known. In our context, we let

$$\boldsymbol{Y} = \begin{bmatrix} y \\ Z_{sal} \end{bmatrix}, \boldsymbol{A} = \begin{bmatrix} 1 & 1 & 0 \\ \gamma & 0 & 1 \end{bmatrix}, \text{ and } \boldsymbol{X} = \begin{bmatrix} \mu_\theta \\ \epsilon_i \\ \nu_i \end{bmatrix},$$

where y and Z_{sal} denote log realized salary and log elicited salary expectation, respectively. The belief-elicitation relationship, γ, is estimated prior to the deconvolution following the instrumental-variables procedure in Section 6.1.

Since beliefs (μ_θ), expectational error (ϵ), and elicitation error (ν) are mutually independent, we can use the Bonhomme-Robin framework to non-parametrically estimate both density of types across individuals, $g(\mu_\theta)$, and density of expected log salary within type, $f_{y|\theta}(y) = f_\epsilon(y - \mu_\theta)$.

The procedure uses empirical characteristic functions of observed measurements to uncover the empirical characteristic functions of unobserved latent factors. These characteristic functions are then transformed into density functions through inverse Fourier transformation. This transformation requires kernel and bandwidth choice to facilitate smoothing. We use the second-order kernel specified in Bonhomme and Robin (2010). To select bandwidth, we use the recommended bandwidth selector from Delaigle and Gijbels (2004).

Appendix D Machine Learning Details

To assess the predictive power of private information and form non-parametric lower bounds on the magnitude of private information, we use machine learning techniques to form predictions of y using public and private information ($E[y|X,Z]$), and using public information only ($E[y|X]$). For each binary outcome, we train a ten-fold cross-validated random forest model with five-thousand

trees on a 70% sample of our data and measure its predictive performance using the 30% holdout sample. We repeat this procedure for each subset of predictor variables given by the categories listed at the top of Table A2, using the first three subsets to estimate $E[y|X]$ under alternative definitions of X, and using the final subset, "All Public + Elicitations", to estimate $E[y|X,Z]$.

For log salary, we follow the same procedure as we do for binary outcomes, but adapt the random forest algorithm to predict not just the conditional mean of y, $E[y|X]$, but also its conditional quantile function, $F^{-1}(\alpha|X)$ for all $\alpha \in [0,1]$, a technique known as quantile regression forests (Meinshausen, 2006). We use these estimated quantile functions to form predicted *level* salary conditional on employment, $E\left[\widehat{e^{\log(y^S)}|y>0},X,Z\right]$, which we then combine with employment predictions, $Pr(\widehat{y>0}|X,Z)$ to form predicted *unconditional* level salary:

$$E\left[\widehat{y^S|X},Z\right] = Pr(\widehat{y>0}|X,Z) * E\left[\widehat{e^{\log(y^S)}|y>0},X,Z\right]. \tag{30}$$

Appendix E MVPF Derivation

Continuous Contracts

We begin the construction of the MVPF with the costs. Let $C(\eta,\lambda)$ denote the net cost to the government of offering a contract of size η at price, λ. The marginal cost, $\frac{\partial C(\eta,\lambda)}{\partial \eta}$, of providing the first dollar of equity financing at price at price λ is given by the sum of two terms. First, there is the marginal cost of subsidizing an adversely-selected contract. These (negative) profits are given by:

$$\Pi(\lambda) = \Pr[\theta \leq \theta_\lambda](E[y|\theta \leq \theta_\lambda] - \lambda). \tag{31}$$

Note that if $\lambda = E[y]$, the contract would break even in the absence of adverse selection. But, the fact that the no trade condition holds above implies that $\Pi(\lambda)$ is negative for all possible values of λ.

In contrast to a private financier, the government also incurs any fiscal externalities from changes

in individuals' (lifetime) earnings in response to the contract.[59] To capture these effects, we conceptualize the equity contract as the union of two components: an increase in college funding, g, given by $\frac{dg}{d\eta} = \lambda$, and an increase in future tax rates, τ, given by $\frac{d\tau}{d\eta} = 1$. The net effect of these components on lifetime earnings, denoted as y^L, can be positive or negative. On the one hand, the increased up-front funding might improve future earnings by relaxing liquidity constraints and increasing human capital investments (g may increase y^L). On the other hand, higher post-college tax rates may reduce earnings (τ may decrease y^L). We express the equity contract's net effect on earnings for each type θ as the sum of these two effects:

$$FE(\lambda) \equiv \tau \frac{dE\left[y^L\right]}{d\eta} = \underbrace{-\frac{\tau}{1-\tau} E\left[y|\theta \leq \theta_\lambda\right] \epsilon_{y,1-\tau}(\theta_\lambda)}_{\text{Tax Distortion}} + \underbrace{\tau \frac{dE\left[y^L\right]}{dg}}_{\text{Impact on Earnings via Grant}}, \tag{32}$$

where $\epsilon_{y,1-\tau} = \frac{1-\tau}{E[y|\theta \leq \theta_\lambda]} \frac{dE[y|\theta \leq \theta_\lambda]}{d(1-\tau)}$ is the elasticity of taxable income at year six, $\frac{dE[y|\theta \leq \theta_\lambda]}{d\tau}$ is the impact of higher implicit taxes from the equity repayment, and $\tau \frac{dE\left[y^L\right]}{dg}$ is the impact of a \$1 grant for college financing on lifetime tax payments.

Putting these terms together, the total marginal cost to the government is the sum of the negative profits from the contract and the fiscal externality on tax revenue,

$$\begin{aligned} \frac{dC(\eta,\lambda)}{d\eta}|_{\eta=0} &= -\Pi(\lambda) - FE(\lambda) \\ &= \Pr\{\theta \leq \theta_\lambda\} \left[\lambda - E\left[y|\theta \leq \theta_\lambda\right] - \tau \frac{dE\left[y^L\right]}{dg} \frac{1}{\Pr\{\theta \leq \theta_\lambda\}} + \frac{\tau}{1-\tau} E\left[y|\theta \leq \theta_\lambda\right] \epsilon_{y,1-\tau}\right] \end{aligned} \tag{33}$$

Next we turn to the aggregate willingness to pay among enrollees. The value of contract λ for an individual of type θ equals its impact on expected utility, $u_1(\theta) - E(\lambda y u_2|\theta)$, divided by the marginal utility of income at the time financing is received, $u_1(\theta)$.

[59] Hendren and Sprung-Keyser (2020) shows that these behavioral responses have only second order effects on financier profits. But, these effects are first order to the government because of pre-existing tax distortions.

Normalizing by expected utility, this individual's willingness to pay is given by

$$wtp\left(\theta\right)=\frac{\frac{dU}{d\eta}}{u_1\left(\theta\right)}=\lambda-\frac{E\left(yu_2|\theta\right)}{u_1\left(\theta\right)}$$
$$=\underbrace{\lambda-E\left[y|\theta\right]}_{\text{Transfer}}+\underbrace{E\left[y|\theta\right]-WTA\left(\theta\right)}_{\text{Consumption Smoothing}}, \tag{34}$$

where $WTA(\theta)$ is the willingness to accept for a type θ so that $\lambda - WTA\left(\theta\right)$ is the net surplus to the individual. Integrating over all types θ who choose to take up the contract, $\theta \leq \theta_\lambda$, and dividing by the government's net marginal cost, $\frac{dC(\eta,\lambda)}{d\eta}$, yields the MVPF:

$$\begin{aligned} MVPF\left(\lambda\right) &= \frac{\int_0^{\theta_\lambda} wtp(\theta)dG(\mu_\theta)}{-\Pi\left(\lambda\right)-FE\left(\lambda\right)} \\ &= \frac{\lambda-E\left[y|\theta\leq\theta_\lambda\right]+\left(E\left[y|\theta\leq\theta_\lambda\right]-E\left[WTA\left(\theta\right)|\theta\leq\theta_\lambda\right]\right)}{\lambda-E\left[y|\theta\leq\theta_\lambda\right]-\tau\lambda\frac{dE[y^L]}{dg}\frac{1}{\Pr\{\theta\leq\theta_\lambda\}}+\frac{\tau}{1-\tau}E\left[y|\theta\leq\theta_\lambda\right]\epsilon_{y,1-\tau}}. \end{aligned} \tag{35}$$

Binary Contracts

This section derives the MVPF for our binary contract in which $y = 1$ corresponds to repayment at price κ. To begin, note that the willingness to pay out of today's income for the contract by a type θ is given by:

$$wtp^d\left(\theta\right)=1-\kappa E\left[y|\theta\right]E\left[\frac{u_2}{u_1}|y=1,\theta\right]$$

The individual receives \$1 but must repay κ in the event of repayment, which occurs with probability $\Pr\{y=1|\theta\}=E\left[y|\theta\right]$. But, θ is not observed. A natural choice would then be that the government sets $\kappa=\frac{1}{E[y|\theta]}$, which would generate zero profits in the absence of adverse selection. Or, one could set it at $\kappa = 1$ so that the government subsidizes all costs of default. We let θ_κ denote the type that is indifferent to the loan at price κ, so that $wtp^d\left(\theta_\kappa\right)=0$. The aggregate willingness to pay is

then

$$\begin{aligned}
WTP^{debt}(\kappa) &= \Pr\{\theta \leq \theta_\kappa\}\left[1-\kappa E_\theta\left[E[y|\theta]\,E\left[\frac{u_2}{u_1(\theta)}|y=1,\theta\right]|\theta\leq\theta_\kappa\right]\right] \\
&\approx \Pr\{\theta \leq \theta_\kappa\}\left(1-\kappa E_\theta\left[E[y|\theta]\left(1+\sigma\left(1-E[y|\theta]\right)\frac{\left(E[c|y=1,\theta]-E[c|y=0,\theta]\right)}{\bar{c}}\right)|\theta\leq\theta_\kappa\right]\right) \\
&\approx \Pr\{\theta \leq \theta_\kappa\}\left(1-\kappa E\left[E[y|\theta]\,|\theta\leq\theta_k\right]+\right. \\
&\qquad \left.\kappa\sigma E_\theta\left[\left(1-E[y|\theta]\right)E[y|\theta]\,\frac{\left(E[c|y=0,\theta]-E[c|y=1,\theta]\right)}{\bar{c}}|\theta\leq\theta_\kappa\right]\right) \\
&\approx \Pr\{\theta \leq \theta_\kappa\}\left(1-\kappa E\left[y|\theta\leq\theta_k\right]+\kappa\sigma var\left(y|\theta\leq\theta_\kappa\right)\frac{\Delta c}{c}\right)
\end{aligned}$$

where the first line uses the Taylor expansion above to measure willingness to pay and the last line assumes the consumption difference is constant across θ (which is a common simplification often made in existing literature in optimal social insurance).

The marginal cost to the government follows a similar pattern to the equity contract. The lost profits to the financier is the cost of the \$1 provision minus the repayment κ for those that repay. In addition, we have fiscal externalities arising from two sources: the upfront grant of η and the repayment of $\eta\kappa$ in the event of non-default. Similar to the equity contract, the upfront grant increases tax revenue by $\tau\frac{dE[y^L]}{dg}$. The debt repayment in adulthood likely reduces tax revenue by $-\kappa\tau\frac{dE[y^L]}{dD}$, where $\tau\frac{dE[y^L]}{dD}$ is the impact of \$1 additional debt burden on tax revenue. Summing,

$$\frac{dC^{debt}(\kappa)}{d\eta}|_{\eta=0} = \Pr\{\theta\leq\theta_\kappa\}\left(1-\kappa E\left[y|\theta\leq\theta_\kappa\right]\right)-\tau\frac{dE\left[y^L\right]}{dg}-\kappa\tau\frac{dE\left[y^L\right]}{dD}$$

So, the MVPF of a debt contract that requires repayment of κ is

$$MVPF^{debt}(\kappa) \approx \frac{1-\kappa E\left[y|\theta\leq\theta_k\right]+\kappa\sigma var\left(y|\theta\leq\theta_\kappa\right)\frac{\Delta c}{c}}{1-\kappa E\left[y|\theta\leq\theta_\kappa\right]-\tau\frac{1}{\Pr\{\theta\leq\theta_\kappa\}}\frac{dE[y^L]}{dg}-\kappa\tau\frac{1}{\Pr\{\theta\leq\theta_\kappa\}}\frac{dE[y^L]}{dD}}$$