Economic Quarterly—Volume 100, Number 4—Fourth Quarter 2014—Pages 279–295

The Dropout Option in a Simple Model of College Education

Ali Ozdagli and Nicholas Trachter

It has been estimated that college graduates earn substantially more than those who do not pursue higher education (Kane and Rouse 1995; Heckman, Lochner, and Todd 2008). But the high rates of return associated with college graduation seem at odds with the low enrollment and graduation rates observed in the data (Athreya and Eberly 2015). Most of the literature addresses this apparent inconsistency in two alternative ways. The first way introduces credit constraints that prevent some students from enrolling and some others from graduating (Keane and Wolpin 2001, Cameron and Taber 2004, Stinebrickner and Stinebrickner 2008, and Lochner and Monge-Naranjo 2011). The second way models college education as a process of learning about self-ability (Stinebrickner and Stinebrickner 2012, Strange 2012, Hendricks and Leukhina 2014, and Trachter 2015). The evidence in favor of the first way is mixed, and thus the second method gained attention in recent years. In this article we aim to provide a simple laboratory to provide some sense of the workings of the learning model.

We build a simple continuous-time model where high school graduates are uncertain about their innate ability to accumulate human capital in college.[1] They are endowed with a belief about their initial ability level. Those with pessimistic initial beliefs find it optimal not to pursue higher education and instead join the workforce, while

Ozdagli is a senior economist with the Boston Fed. Trachter is an economist with the Richmond Fed. The views expressed in this article are those of the authors and do not necessarily represent the views of the Federal Reserve Bank of Boston, the Federal Reserve Bank of Richmond, or the Federal Reserve System. E-mail: nicholas.trachter@rich.frb.org.

[1] The modeling approach relates to the ones in Miao and Wang (2007) and Ozdagli and Trachter (2011).

those with optimistic beliefs enroll in college. During college education students are confronted with random events, which we label as exams, that convey information about the student's true ability level. This information causes students to update their beliefs, making students either more or less optimistic about the expected wage increase upon obtaining a college degree. Those that become very pessimistic about it find it optimal to drop out and join the workforce without a college degree.

We make several assumptions to keep the model simple and tractable. Although some of the assumptions look stark, they allow not only for an extremely tractable model, but also allow us to solve the model in closed form for the college enrollment rate, dropout rate, and wage premium upon graduation. These three objects are at the core of the puzzling nature of college education described above. With this in mind, it is natural to calibrate the model to match these three figures. The model, for its given extreme simplicity, does a good job in matching these objects.

We use the model to gauge the importance of the dropout option. We do so by calculating the fraction of the returns to college enrollment that comes from the availability of this option relative to an alternative scenario where students are not allowed to drop out. We find that a large fraction of measured returns are associated with the fact that students are allowed to drop out.

The rest of the article is organized as follows. Section 1 presents the framework. Section 2 characterizes the problem of a worker, while Section 3 characterizes the problem of a student. Section 4 pertains to the calibration exercise, while Section 5 presents a simple measurement of the value attached to the dropout option. Finally, Section 6 concludes.

1. FRAMEWORK

We build a simple, continuous-time model of postsecondary education. Upon high school graduation, an individual has to decide between joining the workforce or pursuing a degree in a (four-year) college.[2] Agents are endowed with an initial wealth level a_0, and they differ in their ability to accumulate human capital in college. For simplicity, we assume that ability can take two levels: low and high. Let $\mu \in \{0, 1\}$ denote the ability level, with $\mu = 0$ denoting low ability. The ability level is

[2] Although relevant, for simplicity, we abstract from the decision to enroll or transfer to/from two-year colleges.

not observable, and the high school graduate is endowed with a belief $p_0 = \Pr(\mu = 1)$.

At any point in time, an agent can be working or attending college. We assume that working is an absorbing state, and that a college student graduates with intensity (the analogous of probability in a discrete time setup) ϕ, so that the expected time until obtaining a degree is $1/\phi$. Upon college graduation, the true ability level of the agent is revealed. We denote with τ college tuition, and by $h(\mu, GS)$ the wage, in any given period, of an agent joining the workforce with true ability level μ, and graduation status $GS \in \{0, 1\}$, where 0 implies that the agent does not hold a college degree, and 1 does. In particular, we let $h(0,0) = h(1,0) \equiv \underline{h}$, and $h(1,1) \equiv \bar{h} > h(0,1) = \underline{h}$. That is, there is a graduation premium only if the ability level of the agent is high. Notice that it follows that the process for wealth accumulation can be written as $da_t/dt = ra_t + h(\mu, GS) - c_t$ if the agent is working, and by $da_t/dt = ra_t - \tau - c_t$ if the agent is enrolled in college, where c_t denotes the consumption level at time t.

While attending college, students are faced with exams every period. The exams have two alternative outcomes: a passing grade or a good grade. We assume that only high-ability students are able to obtain a good grade, something that occurs with intensity λ.[3] Because the grade obtained in a particular exam is correlated with the ability level of the student, grades convey information that the students use to update their beliefs: Upon receiving a good grade, the belief of a student "jumps" from p to 1, because only high-ability students can obtain a good grade, while, upon receiving a passing grade, students update their beliefs by Bayes' rule. Suppose that the period length is given by Δ. Bayes' rule implies that the belief at period $t + \Delta$ is given by

$$p_{t+\Delta} = \frac{(1 - \lambda\Delta)p_t}{p_t(1 - \lambda\Delta) + (1 - p_t)} \ .$$

The denominator, $p_t(1-\lambda\Delta)+(1-p_t)$, accounts for the probability that we observe belief $p_{t+\Delta}$ at time $t+\Delta$, while the numerator accounts for the probability that we observe the high ability level $\mu = 1$ (which is conditional on the student being of high ability) times the belief that this event could occur, p_t. Subtracting p_t on both sides and dividing

[3] It is easy to add more grades and to correlate them with the ability level of students. However, this would add more parameters to the model, increasing the complexity of the calibration exercise.

by Δ provides

$$\frac{p_{t+\Delta} - p_t}{\Delta} = -p_t(1-p_t)\frac{\lambda}{p_t(1-\lambda\Delta) + (1-p_t)} \ .$$

Notice that $dp_t/dt = \lim_{\Delta \downarrow 0} \frac{p_{t+\Delta}-p_t}{\Delta}$. Then, taking the limit as Δ approaches zero provides that

$$dp_t/dt = -\lambda p_t(1-p_t) \ .$$

The learning equation shows that, conditional on the student not receiving a good grade, the beliefs of a student fall through time. This follows because as time passes and no fully revealing signal is received, the student updates her beliefs toward believing that her ability level is low. Further, notice that the speed of learning depends on (i) the intensity parameter λ regulating the probability of receiving a fully revealing signal, and (ii) the current level of beliefs p_t. A higher λ implies that signals arrive at a faster rate and thus not receiving the signal conveys more information too (the student updates her beliefs faster toward being of low ability). The process depends on p_t, as this regulates how uncertain is the student. In particular, the highest uncertainty level is $p_t = 1/2$, which maximizes the value $p_t(1-p_t)$. The higher the degree of uncertainty, the more relevant the arrival of information becomes.

A high school graduate, endowed with p_0 and a_0, chooses her consumption stream c_t, whether to enroll and/or remain in college at every point in time, in order to maximize her time-separable expected discounted lifetime utility, with period utility function $u(c_t) = e^{\gamma c_t}/-\gamma$, where γ is the coefficient of constant absolute risk aversion. We further assume that the rate of discount equals the interest rate, and we denote both by r.

Let $V(p, a)$ denote the value for a college student with current belief p and wealth level a, and let $W(h(\mu, GS), a)$ denote the value for a worker with wage $h(\mu, GS)$ and wealth a. Given the structure of the problem, we start by solving the problem of workers; we then tackle the problem of students, and we finalize by obtaining the optimal policy of high school graduates.

2. THE PROBLEM OF A WORKER

We start by describing the problem of a worker. The value function $W(h(\mu, GS), a)$ solves

$$rW(h(\mu, GS), a) = \max_c \frac{e^{-\gamma c}}{-\gamma} + W_a(h(\mu, GS), a)(ra + h(\mu, GS) - c) \ .$$

The equation states that the flow value of being a worker with wage $h(\mu, GS)$ and wealth a, is equal to the instantaneous utility derived from consumption, in addition to the change in value accrued to the change in the wealth level of the worker. Notice that this is a standard savings problem, with solution given by $W(h(\mu, GS), a) = -\frac{e^{-\gamma(ra+h(\mu,GS))}}{\gamma r}$.

3. THE PROBLEM OF A STUDENT

The value function for a student with current belief p and wealth level a solves

$$\begin{aligned} rV(p,a) &= \max_c \frac{e^{-\gamma c}}{-\gamma} + p\lambda[V(1,a) - V(p,a)] \\ &\quad +\phi[pW(\bar{h},a) + (1-p)W(\underline{h},a) - V(p,a)] \\ &\quad -V_p(p,a)\lambda p(1-p) + V_a(p,a)(ra - \tau - c) \ . \end{aligned}$$

Let us provide some intuition to this Bellman equation. The left-hand side of the equation measures the flow value of being a student with current belief p and wealth a. This value has to be equal to the sum of five different terms. The first term is the instantaneous utility derived from consumption. The second term accounts for the change in value if the student is revealed to be of high ability, which, conditional on the student being of high ability (which occurs with probability p), happens with intensity λ. The third term accounts for the change in value if the student graduates (remember that upon graduation ability is revealed). Finally, the fourth and fifth terms account for the change in value accrued for the change in beliefs and wealth, if the student remains enrolled in school and no fully revealing signal had arrived.

Solving for the value function of students of unknown type requires us first to obtain an expression for the value function when the student is of high ability, $V(1,a)$. In Appendix A we solve for this value function. We obtain that $V(1,a) = -b\frac{e^{-\gamma(ra+\bar{h})}}{\gamma r}$, where $b > 1$ is the unique solution to $\phi/b = \phi - \gamma r(\bar{h} + \tau) + r \ln b$. We will restrict our attention to cases where $V(1,a) > W(\underline{h}, a)$, so that high-ability students always find it worthwhile to enroll in college and remain until graduation.

Using the expressions for the value function for a worker and for a student with $p = 1$, we can now solve for the value function of a student of unknown type. Details on the derivation can be found in Appendix B. We obtain that $V(p,a) = -\frac{e^{-\gamma(ra+f(p))}}{\gamma r}$, where $f(p)$ solves

$$(\phi + p\lambda - \gamma f'^{-\gamma f(p)} = p\lambda b e^{-\gamma\bar{h}} + \phi\left[pe^{-\gamma\bar{h}} + (1-p)e^{-\gamma\underline{h}}\right] \ .$$

In Appendix C we show that $f(p)$ is increasing. Although the proof is somewhat involved, the result is intuitive. The idea is that high-ability students are the ones who obtain a wage boost after graduation, and also are the ones likely to receive positive signals. As a result, a higher belief p implies a higher expected gain from college enrollment.

The decision to enroll in college involves comparing the value accrued if an agent enrolls, $V(p, a)$, with the value if she drops out, $W(\underline{h}, a)$. Notice that these two functions are multiplicative in $e^{-\gamma ra}$, and thus the optimal enrollment policy should not depend on the wealth level a. This is a natural implication of the exponential utility function used in the article that implies no income effects. As a result, the enrollment decision depends solely on the belief of the student. Because the value of college enrollment is increasing in the student's belief p—because $f(p)$ is increasing in p, we conjecture that there exists a threshold p^* such that those students with beliefs $p < p^*$ do not attend college, while those with $p > p^*$ pursue higher education. Furthermore, because there is no direct cost to be paid after enrollment or dropping out from college, the threshold p^* also regulates who drops out. That is, the decision about whether to remain enrolled at college with belief p is identical to the decision of whether to start college with a belief of p. Because the value function $V(p, a)$ is continuous in p, it has to be the case that a student with prior p^* and wealth level a is indifferent between working and enrolling in college. That is, $V(p^*, a) = W(\underline{h}, a)$, which is known as the value matching condition. Also, at the indifference point, it has to be the case that there is no extra value of staying in school for the marginal student. That is, $V_p(p^*, a) = 0$, a condition known as the smooth pasting condition. Using the expressions we found for $V(p, a)$ and $W(\underline{h}, a)$, we can reduce these two conditions to $f(p^*) = \underline{h}$, and $f'^*) = 0$. We can combine these two conditions with the expression that $f(p)$ has to satisfy to solve for p^*,

$$p^* = \frac{\gamma r(\tau + \underline{h})}{\phi\left(1 - e^{-\gamma(\bar{h} - \underline{h})}\right) + \lambda\left(1 - be^{-\gamma(\bar{h} - \underline{h})}\right)} \ .$$

Before concluding this section, it is useful to define the value of college attendance. Let $\sigma(p)$ denote the maximum amount of wealth a high school graduate with belief p is willing to forgo to have the option to attend college. Notice that $\sigma(p)$ solves $V(p, a - \sigma(p)) = W(\underline{h}, a)$. Immediate calculations provide that $\sigma(p) = (f(p) - \underline{h})/r$. Notice that the value of attending college for the student at the margin is zero. That is, $\sigma(p^*) = 0$.

4. CALIBRATION: A PARAMETRIC EXAMPLE

In this section we aim to provide a simple calibration of the model in order to be able to gauge the option value provided to students by having the opportunity to drop out from school. We will propose a simple way for constructing initial beliefs, i.e., p_0, and we will parameterize the model by a combination of direct imputation and calibration exercise.

Imputed Parameters

We set the interest rate r to be 0.02, which implies a yearly discount factor of 0.98. We standardize $\underline{h} \equiv 1$, we set the risk aversion parameter $\gamma = 8$, and, following Trachter (2015), we set $\tau = 0.32$.[4]

Calibrated Parameters

The remaining parameters, λ (the learning parameter), $\bar{h}$ (the wage premium parameter), and s (the share of high-ability students), are estimated by a method of moments. With this in mind, we build three moments that depend on these three parameters: the college enrollment rate (C^e), the college dropout rate (C^d), and the average wage premium of college graduates (W^p).

To calibrate the model, we need to specify the true ability level of agents and their initial beliefs. We do so by setting up density functions from where the ability levels are drawn, and we then compute, using Bayes' rule, the initial belief of each individual. There are alternative ways of producing initial beliefs. Some studies, such as Hendricks and Leukhina (2014) and Trachter (2015), estimate initial beliefs by the means of parametric models that link them to observable and unobservable measures of initial ability. A simpler approach that does not rely on microdata would be to specify a structure of initial beliefs and figure out the mass of the individual at each belief level so as to match moments in the data. The method we follow in this article follows this idea, but with the further assumption that we impose structure on how we relate individual types, initial beliefs, and the measure of agents at each belief level.

We normalize the total mass of high school graduates to be one. We let s denote the share of high-ability students. Each agent receives a signal, q_0, drawn from a distribution that is type dependent, $\Gamma_\mu(q_0)$.

[4] The value for τ follows from obtaining the average expenditures in tuition (including room and board), and then normalizing by the wage of those not pursuing college education and dropouts. More details can be found in Trachter (2015).

For any given q_0, then, p_0 is just the result of Bayes' rule, and the gamma distributions, transformed by Bayes' rule, provide the distribution of initial beliefs p_0. For simplicity, we assume that $\Gamma_1(q_0) = 2q_0$, and $\Gamma_0(q_0) = 2 - 2q_0$, for $q_0 \in [0, 1]$, to capture the idea that a high q_0 (and thus a high p_0) is more likely to be drawn from the distribution for high-ability students. Bayes' rule provides that

$$p_0 \equiv \Pr(\mu = 1 | q_0) = \frac{\Pr(q_0 | \mu = 1)\Pr(\mu = 1)}{\Pr(q_0)} = \frac{q_0 s}{sq_0 + (1-s)(1-q_0)} . \quad (1)$$

We now compute the three required moments. We start by providing an expression for the college enrollment rate, which is given by

$$\begin{aligned} C^e &= \int_{q(p^*)}^{1} (h_0(q_0)(1-s) + h_1(q_0)s) dq_0 \\ &= 1 - 2\left[(1-s)q(p^*) + (2s-1)\frac{q(p^*)^2}{2}\right] , \end{aligned}$$

where $q(p^*)$ follows from inverting the expression in (1).

The college dropout rate is given by

$$\begin{aligned} C^d &= s\int_{q(p^*)}^{1} Q(p(q_0), 1)\frac{h_1(q_0)}{1 - H_1(q^*)} dq_0 \\ &\quad + (1-s)\int_{q(p^*)}^{1} Q(p(q_0), 0)\frac{h_0(q_0)}{1 - H_0(q^*)} dq_0 , \end{aligned}$$

where $Q(p, \mu)$ denotes the dropout probability of a college student of type μ and belief p (see Appendix D for a derivation of the expressions for $Q(p, \mu)$). Some algebra provides that

$$\begin{aligned} C^d &= s\frac{2}{1 - q(p^*)^2}\left(\frac{p^*}{1-p^*}\frac{1-s}{s}\right)^{\frac{\phi+\lambda}{\lambda}} Beta\left(q(p^*), 2 + \frac{\phi}{\lambda}, 1 - \frac{\phi}{\lambda}\right) \\ &\quad + (1-s)\frac{2}{1 - 2q(p^*) + q(p^*)^2}\left(\frac{p^*}{1-p^*}\frac{1-s}{s}\right)^{\frac{\phi}{\lambda}} \\ &\quad Beta\left(q(p^*), 2 + \frac{\phi}{\lambda}, 1 - \frac{\phi}{\lambda}\right) , \end{aligned}$$

where details of the derivation can be obtained in Appendix E, and where $Beta(\cdot, \cdot, \cdot)$ denotes an incomplete Beta function, and where it is required that $\lambda > \phi$.

Finally, the wage premium upon graduation is given by

$$W^p = \bar{h}\frac{sN_1}{sN_1 + (1-s)N_0} + \underline{h}\frac{(1-s)N_0}{sN_1 + (1-s)N_0} ,$$

Table 1 Moments in Model and Data

	Model	**Data**
Enrollment rate	0.34	0.25
Dropout rate	0.38	0.4
Wage premium	0.34	0.21

where the amount sN_1 $((1-s)N_0)$ accounts for the mass of high- (low-) ability students who graduate, where $\int_{q(p^*)}^{1} (1 - Q(p(q_0), 1))\, h_1(q_0) dq_0$ and $N_0 = \int_{q(p^*)}^{1} (1 - Q(p(q_0), 0))\, h_0(q_0) dq_0$.

Following a similar method as the one used for the dropout rate, we can find expressions for N_1 and N_0,

$$N_1 = 1 - q(p^*)^2 - 2\left(\frac{p^*}{1-p^*}\frac{1-s}{s}\right)^{\frac{\phi+\lambda}{\lambda}} Beta\left(q(p^*), 2 + \frac{\phi}{\lambda}, 1 - \frac{\phi}{\lambda}\right) ,$$

$$N_0 = 1 + q(p^*)^2 - 2q(p^*) - 2\left(\frac{p^*}{1-p^*}\frac{1-s}{s}\right)^{\frac{\phi}{\lambda}} Beta\left(q(p^*), 2 + \frac{\phi}{\lambda}, 1 - \frac{\phi}{\lambda}\right) .$$

We aim to calibrate $\bar{h}$, λ, and s by looking for the values of these three parameters that make the model implied enrollment rate, dropout rate, and average wage premium to be as close as possible to its empirical counterpart. We again borrow the empirical figures from Trachter (2015), which computed these figures from the *National Longitudinal Study of the High School Class of 1972* (NLS-72). Table 1 presents the data and model-implied moments. The model, for its given simplicity, does a good job in matching the data.[5] It overly predicts college enrollment and wage premium of graduates, while the college dropout rate is very close to its empirical counterpart. This exercise provided us with estimates for the remaining parameters: $\bar{h} = 1.36$, $\lambda = 0.241$, and $s = 0.6$.

5. THE VALUE OF THE DROPOUT OPTION

In this section, we provide a simple measurement on the valued added by the dropout option. We do so by comparing the gains from college enrollment that a student gets if the option is available, $\sigma(p)$, with those

[5] The fit of the model can be improved in several ways, as done in Trachter (2015). For example, by using a less restrictive belief construction or learning process.

Figure 1 Decomposition of Returns

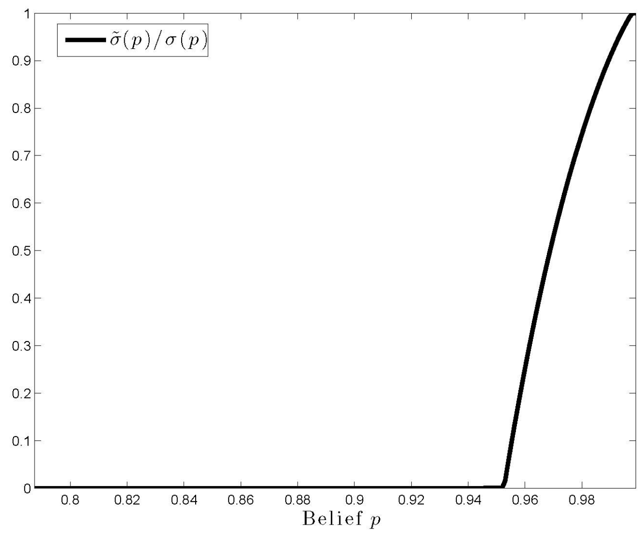

returns if the option is removed, which we define by $\tilde{\sigma}(p)$.[6] Notice that $\tilde{\sigma}(p)/\sigma(p)$ will measure the fraction of the gains that follow from the fact that students are allowed to enroll, while $1 - \tilde{\sigma}(p)/\sigma(p)$ measures the fraction of the gains that follow from the fact that students are also allowed to drop out.

Figure 1 presents the decomposition of returns to enrollment for all the beliefs that imply college enrollment (i.e., $p > p^* = 0.7875$). For a large fraction of beliefs (up to $p \approx 0.95$), the dropout option accounts for all of the returns to college enrollment. This follows because students with beliefs in this range do not find it very likely to graduate and be of high ability, so they value highly the dropout opportunity. As the belief keeps rising toward one, the enrollment option gains value as students are more likely to be of high ability and so value less the dropout option. Overall, the figure shows that the dropout option accounts for a large part of the returns to enrollment. In fact, using the densities for the initial types, we can compute the average value added of the dropout option: The dropout option accounts for 85 percent of the value of college enrollment.

[6] We do not include this second model in the article, but it can be easily computed following a similar procedure to the one used for the model with the dropout option.

6. CONCLUDING REMARKS

This article proposes a simple and highly tractable model of postsecondary education. In the model, students are allowed to drop out at any point in time as they update their beliefs about the gains accrued from college education. For its apparent simplicity, the parameterized version of the model is able to generate patterns consistent with the data. As a result, this article showcases the potential of learning models to explain the patterns observed in postsecondary education.

The simplified nature of the model implied that we abstracted from several noteworthy elements that any exercise attempting to gauge the returns to education might need to address. First, the model does not provide any wage premium for graduates of low ability. Lifting this assumption is very simple and would allow us to gauge more precisely the importance of the dropout option as, without the wage premium for low-ability students, the model overestimates the importance of the dropout option. Although, as shown in Trachter (2015), the qualitative properties of this article would survive. Second, the model does not allow for any wealth effects, something undesirable if we want to explore the connection between wealth differences and education outcomes. In Ozdagli and Trachter (2011), we evaluate a setup similar to the one developed here but where wealth differences affect the risk preferences of students. We show that, consistent with the data, wealthier students are less likely to drop out and, conditional on this event, they drop out later than poorer students.

The goal of this article is to introduce the reader to a particular class of models being used to analyze postsecondary patterns of education. These models are highly tractable, are easy to handle, and are able to fit several salient features of the data (for example, see Hendricks and Leukhina [2014] and Trachter [2015]). The ongoing success of these models calls for further research in the area.

APPENDIX A: A STUDENT OF HIGH ABILITY

Evaluating the problem of a student at $p = 1$ provides

$$rV(1,a) = \max_{c} \frac{e^{-\gamma c}}{-\gamma} + \phi[W(\bar{h},a) - V(1,a)] + V_a(1,a)(ra - \tau - c) \ .$$

Notice that standard techniques imply that the solution to this problem is unique. Plugging in the first-order condition provides

$$rV(1,a) = \frac{V_a(1,a)}{-\gamma} + \phi[W(\bar{h},a) - V(1,a)] + V_a(p,a)(ra - \tau + \gamma^{-1}\ln V_a(1,a)) .$$

We look for the solution to this problem by a guess and verify method, which readily provides the expression provided in the main text.

APPENDIX B: A STUDENT OF UNKNOWN TYPE

The problem of the student can be rewritten as

$$\begin{aligned}(r+\phi+p\lambda)V(p,a) &= \max_c \frac{e^{-\gamma c}}{-\gamma} + p\lambda V(1,a) \\ &\quad +\phi[pW(\bar{h},a) + (1-p)W(\underline{h},a)] \\ &\quad -V_p(p,a)\lambda p(1-p) + V_a(p,a)(ra-\tau-c) .\end{aligned}$$

As with the problem of a worker, it is straightforward to argue that this problem has a unique solution. The first-order condition provides $e^{-\gamma c} = V_a(p,a)$. Solving for c, we get that $c = -\gamma^{-1}\ln V_a(p,a)$. Using these two expressions, we obtain that

$$\begin{aligned}(r+\phi+p\lambda)V(p,a) &= \frac{V_a(p,a)}{-\gamma} + p\lambda V(1,a) + \phi[pW(\bar{h},a) \\ &\quad +(1-p)W(\underline{h},a)] - V_p(p,a)\lambda p(1-p) \\ &\quad +V_a(p,a)(ra-\tau+\gamma^{-1}\ln V_a(p,a)) .\end{aligned}$$

Notice that this last expression implies that the value function $V(p,a)$ has to satisfy a partial differential equation. We solve for $V(p,a)$ by a guess and verify method. We guess that $V(p,a) = -\frac{e^{-\gamma(ra+f(p))}}{\gamma r}$, where $f(p)$ is a function to be determined. Under this guess, the previous expression reduces to

$$(\phi + p\lambda - \gamma f'^{-\gamma f(p)} = p\lambda b e^{-\gamma\bar{h}} + \phi\left[pe^{-\gamma\bar{h}} + (1-p)e^{-\gamma\underline{h}}\right] ,$$

where we also used the solution to the worker problem and the problem of a student with belief $p = 1$.

APPENDIX C: DETAILS ON $f(p)$

Here we show that $f(p)$ is increasing in p. To do so, notice that we can rewrite the expression defining $f(p)$ as $-\gamma f'(p)\lambda p(1-p) = \theta(p, f(p))$, where

$$\begin{aligned}\theta(p, f(p)) &\equiv p(\lambda b + \phi)e^{-\gamma(\bar{h}-f(p))} + \phi(1-p)e^{-\gamma(\underline{h}-f(p))} \\ &\quad -(\phi + p\lambda - \gamma r(\tau + f(p))) \ .\end{aligned}$$

Notice that

$$\frac{\partial\theta(p, f(p))}{\partial f(p)} = \gamma p(\lambda b + \phi)e^{-\gamma(\bar{h}-f(p))} + \gamma\phi(1-p)e^{-\gamma(\underline{h}-f(p))} + \gamma r > 0 \ ,$$

and

$$\begin{aligned}\frac{\partial\theta(p, f(p))}{\partial p} &= (\lambda b + \phi)be^{-\gamma(\bar{h}-f(p))} - \phi e^{-\gamma(\underline{h}-f(p))} - \lambda \\ &= \phi e^{\gamma f(p)}\left(e^{-\gamma\bar{h}} - e^{-\gamma\underline{h}}\right) + \lambda\left(be^{-\gamma(\bar{h}-f(p))} - 1\right) \ .\end{aligned}$$

To sign this last expression we need to notice first that the highest attainable value of college enrollment is $V(a,1)$, so that $V(a,p) \leq V(a,1)$ for all p. This condition reduces to $be^{-\gamma(\bar{h}-f(p))} - 1$. Then, we obtain that $\frac{\partial\theta(p,f(p))}{\partial p} < 0$. We use the sign of these two derivatives to show that $f(p)$ increases with p.

We prove that $f'(p) > 0$ by a contradiction argument. Suppose there exists a belief p_1, $p_1 > p^*$, such that $f'(p_1) \leq 0$. Then, it has to be the case that $\theta(p_1, f(p_1)) \geq 0$. Pick a belief p_2 in the neighborhood of p_1 that satisfies $p_2 < p_1$. Because $f'(p_1) \leq 0$, we have that $f(p_2) \geq f(p_1)$. Then, because $\theta(\cdot)$ increases in p and decreases in $f(p)$, we have that $\theta(p_2, f(p_2)) > \theta(p_1, f(p_1)) \geq 0$. This provides that $f'(p_2) < 0$. Iterating on this process provides that $f'^{*}) < 0$, which contradicts the fact that, by construction, $f'^{*}) = 0$. Then, $f'(p) > 0$ for all p.

APPENDIX D: DROPOUT PROBABILITY

Let $Q(p, \mu)$ denote the dropout probability of a student with current belief p and true ability level μ. Notice that we are not including the wealth level a as a state, as we already concluded that the dropout threshold is independent of a. Also, notice that, for a given belief p, a high-ability student is less likely to drop out than a low-ability student given that $\lambda > 0$, $V(1,a) > W(\underline{h},a)$.

For a high-ability student, we have that

$$Q(p,1) = \phi dt \cdot 0 + (1-\phi dt)\left[\lambda dt Q(1,1) + (1-\lambda dt) Q(p - \lambda p(1-p)dt, 1)\right],$$

where $dp/dt = -\lambda p(1-p)$. Given that, $V(1,a) > W(\underline{h},a)$ implies that $Q(1,1) = 0$. Also, noticing that $Q(p - \lambda p(1-p)dt, 1) \approx Q(p,1) - Q_p(p,1)\lambda p(1-p)dt$, we can rewrite the previous expression as

$$\begin{aligned}
Q(p,1) &= (1-\phi dt)(1-\lambda dt)\left[Q(p,1) - Q_p(p,1)\lambda p(1-p)dt\right] \ , \\
Q(p,1) &= (1-(\phi+\lambda)dt + \phi\lambda dt^2)Q(p,1) \\
&\quad -(1-\phi dt)(1-\lambda dt)Q_p(p,1)\lambda p(1-p)dt \ , \\
Q(p,1) &= -\frac{(1-\phi dt)(1-\lambda dt)Q_p(p,1)\lambda p(1-p)}{((\phi+\lambda) - \phi\lambda dt)} \ .
\end{aligned}$$

Taking the limit as dt approaches zero provides that $Q(p,1)$ satisfies a first-order linear ordinary differential equation, $\frac{Q_p(p,1)}{Q(p,1)} = -\frac{\phi+\lambda}{\lambda p(1-p)}$ with boundary condition $Q(p^*,1) = 1$. The general solution to the differential equation is $Q(p,1) = C_1\left(\frac{1-p}{p}\right)^{\frac{\phi+\lambda}{\lambda}}$, where C_1 is such that $Q(p^*,1) = 1$. It is immediate to obtain that $C_1 = \left(\frac{p^*}{1-p^*}\right)^{\frac{\phi+\lambda}{\lambda}}$, so that $Q(p,1) = \left(\frac{p^*}{p}\frac{1-p}{1-p^*}\right)^{\frac{\phi+\lambda}{\lambda}}$.

We now derive $Q(p,0)$. This value solves

$$Q(p,0) = \phi dt \cdot 0 + (1-\phi dt)Q(p - \lambda p(1-p)dt, 0) \ .$$

As we did above, from this expression we obtain that $Q(p,0)$ satisfies $\frac{Q_p(p,0)}{Q(p,0)} = -\frac{\phi}{\lambda p(1-p)}$, with boundary condition $Q(p^*,0) = 1$. It follows that $Q(p,0) = \left(\frac{p^*}{p}\frac{1-p}{1-p^*}\right)^{\frac{\phi}{\lambda}}$.

APPENDIX E: COMPUTING DROPOUT RATES

We begin by computing $\int_{q(p^*)}^{1} Q(p(q_0),1)\frac{h_1(q_0)}{1-H_1(q^*)}dq_0$. Using the expression for the density $h_1(q_0)$, together with the expression for $Q(p,1)$

derived in Appendix D, simple algebra provides that

$$\begin{aligned}&\int_{q(p^*)}^{1} Q(p(q_0),1)\frac{h_1(q_0)}{1-H_1(q(p^*))}dq_0 \\ =\ & \frac{2}{1-q(p^*)^2}\left(\frac{p^*}{1-p^*}\frac{1-s}{s}\right)^{\frac{\phi+\lambda}{\lambda}}\int_{q(p^*)}^{1} q_0^{1-\frac{\phi+\lambda}{\lambda}}(1-q_0)^{\frac{\phi+\lambda}{\lambda}}dq_0\ ,\end{aligned}$$

where, from (1), we can obtain that $\frac{1-p}{p}=\frac{1-q_0}{q_0}\frac{1-s}{s}$. Define $y\equiv 1-q_0$. This change of variables allows us to rewrite the previous expression as

$$\begin{aligned}&\int_{q(p^*)}^{1} Q(p(q_0),1)\frac{h_1(q_0)}{1-H_1(q(p^*))}dq_0 \\ =\ & \frac{2}{1-q(p^*)^2}\left(\frac{p^*}{1-p^*}\frac{1-s}{s}\right)^{\frac{\phi+\lambda}{\lambda}}\int_{0}^{1-y(p^*)} y^{\frac{\phi+\lambda}{\lambda}}(1-y)^{1-\frac{\phi+\lambda}{\lambda}}dy\ .\end{aligned}$$

Notice that the expression $\int_0^{1-y(p^*)} y^{\frac{\phi+\lambda}{\lambda}}(1-y)^{1-\frac{\phi+\lambda}{\lambda}}dy$ accounts for an incomplete Beta function $Beta(q(p^*),1+\frac{\phi+\lambda}{\lambda},2-\frac{\phi+\lambda}{\lambda})$. Then,

$$\begin{aligned}&\int_{q(p^*)}^{1} \frac{Q(p(q_0),1)h_1(q_0)}{1-H_1(q(p^*))}dq_0 \\ =\ & \frac{2}{1-q(p^*)^2}\left(\frac{p^*}{1-p^*}\frac{1-s}{s}\right)^{\frac{\phi+\lambda}{\lambda}} Beta\left(q(p^*),1+\frac{\phi+\lambda}{\lambda},2-\frac{\phi+\lambda}{\lambda}\right)\ .\end{aligned}$$

We now turn to compute the term $\int_{q(p^*)}^{1} Q(p(q_0),0)\frac{h_0(q_0)}{1-H_0(q^*)}dq_0$. Again, we use the expression for $h_0(q_0)$ and the expression for $Q(p,0)$ derived in Appendix D to obtain

$$\begin{aligned}&\int_{q(p^*)}^{1} Q(p(q_0),0)\frac{h_0(q_0)}{1-H_0(q^*)}dq_0 \\ =\ & \frac{2}{1-2q(p^*)+q(p^*)^2}\left(\frac{p^*}{1-p^*}\frac{1-s}{s}\right)^{\frac{\phi}{\lambda}}\int_{q(p^*)}^{1}\left(\frac{1-q_0}{q_0}\right)^{\frac{\phi}{\lambda}}(1-q_0)dq_0\ .\end{aligned}$$

The same change of variables we used above provides

$$\begin{aligned}&\int_{q(p^*)}^{1} Q(p(q_0),0)\frac{h_0(q_0)}{1-H_0(q^*)}dq_0 \\ =\ & \frac{2}{1-2q(p^*)+q(p^*)^2}\left(\frac{p^*}{1-p^*}\frac{1-s}{s}\right)^{\frac{\phi}{\lambda}}\int_{0}^{1-y(p^*)} y^{1+\frac{\phi}{\lambda}}(1-y)^{-\frac{\phi}{\lambda}}dy\ ,\end{aligned}$$

from where it follows that

$$\begin{aligned}&\int_{q(p^*)}^{1} Q(p(q_0),0)\frac{h_0(q_0)}{1-H_0(q^*)}dq_0 \\ =\ & \frac{2}{1-2q(p^*)+q(p^*)^2}\left(\frac{p^*}{1-p^*}\frac{1-s}{s}\right)^{\frac{\phi}{\lambda}} Beta\left(q(p^*),2+\frac{\phi}{\lambda},1-\frac{\phi}{\lambda}\right)\ .\end{aligned}$$

REFERENCES

Athreya, Kartik, and Janice Eberly. 2015. "The College Premium, College Noncompletion, and Human Capital Investment." Federal Reserve Bank of Richmond Working Paper 13-02R (February).

Cameron, Stephen V., and Christopher Taber. 2004. "Borrowing Constraints and the Returns to Schooling." *Journal of Political Economy* 112 (February): 132–82.

Heckman, James J., Lance J. Lochner, and Petra E. Todd. 2008. "Earnings Functions and Rates of Return." *Journal of Human Capital* 2 (1): 1–31.

Hendricks, Lutz, and Oksana Leukhina. 2014. "The Return to College: Selection and Dropout Risk." Center for Financial Studies Working Paper 487 (November).

Kane, Thomas J., and Cecilia Elena Rouse. 1995. "Labor-market Returns to Two- and Four-year College." *American Economic Review* 85 (June): 600–14.

Keane, Michael P., and Kenneth I. Wolpin. 2001. "The Effect of Parental Transfers and Borrowing Constraints on Educational Attainment." *International Economic Review* 42 (November): 1,051–103.

Lochner, Lance J., and Alexander Monge-Naranjo. 2011. "The Nature of Credit Constraints and Human Capital." *American Economic Review* 101 (October): 2,487–529.

Miao, Jianjun, and Neng Wang. 2007. "Experimentation under Uninsurable Idiosyncratic Risk: An Application to Entrepreneurial Survival." Unpublished manuscript.

Ozdagli, Ali K., and Nicholas Trachter. 2011. "On the Distribution of College Dropouts: Household Wealth and Uninsurable Idiosyncratic Risk." Federal Reserve Bank of Boston Working Paper 11-8 (July).

Stange, Kevin M. 2012. "An Empirical Examination of the Option Value of College Enrollment." *American Economic Journal: Applied Economics* 4 (January): 49–84.

Stinebrickner, Ralph, and Todd Stinebrickner. 2008. "The Effect of Credit Constraints on the College Drop-out Decision: A Direct Approach using a New Panel Study." *American Economic Review* 98 (December): 2,163–84.

Stinebrickner, Todd, and Ralph Stinebrickner. 2012. “Learning about Academic Ability and the College Dropout Decision.” *Journal of Labor Economics* 30 (October): 707–48.

Trachter, Nicholas. 2015. “Stepping Stone and Option Value in a Model of Postsecondary Education.” *Quantitative Economics* 6 (March): 223–56.

Economic Quarterly—Volume 100, Number 4—Fourth Quarter 2014—Pages 297–352

Loan Guarantees for Consumer Credit Markets

Kartik B. Athreya, Xuan S. Tam, and Eric R. Young

Two specific subsets of the U.S. population—the young and those with temporarily low income (and wealth)—have long been identified as pervasively facing liquidity constraints. Empirical work has long measured the fraction of constrained households at close to 20 percent of the U.S. population (see Zeldes 1989; Jappelli 1990; and Hubbard, Skinner, and Zeldes 1994). More recent work again places importance on the inability to cheaply access unsecured credit when needed (see, e.g., Gross and Souleles 2002 and Telyukova 2013). While the preceding work takes substantial care to arrive at estimates, even the simplest summary measures in U.S. data suggest a lack of access to consumer credit for the groups mentioned above. For example, among families with heads of household who have income between $25,000 to $50,000, 34 percent have *no* credit card at all. Moreover, 60 percent roll over debt and pay interest rates of approximately 15 percent per year, a clear indication of their inability to access cheaper alternatives. Among younger families, those with heads of household of age 35–44, similar patterns emerge: 32 percent have no credit card, while 32 percent roll over debt and also pay the same (high) interest rates. Lastly, poorer households fare badly: 55 percent of those with income between $10,000 and $25,000 have no credit card, while 55 percent of renters

We thank Larry Ausubel, Dean Corbae, Martin Gervais, Kevin Reffett, and participants in seminars at Arizona State, Georgia, Iowa, and the Consumer Credit and Bankruptcy Conference at the University of Cambridge for helpful comments and discussions. We also thank Brian Gaines for his editorial assistance. We especially thank Borys Grochulski, Robert Hetzel, and Miki Doan for detailed comments, and Edward Prescott for his guidance as editor. The opinions expressed here are not necessarily those of the Federal Reserve Bank of Richmond or the Federal Reserve System. E-mail: kartik.athreya@rich.frb.org.

don't have a credit card.[1] The lack of wealth among both young and low-income households also precludes the use of far less expensive secured credit, such as home equity loans or lines of credit. Thus, despite the apparent ubiquity of consumer credit, the young and the poor both appear to face tight restrictions on access to the principal source of credit available to them.

The populations most routinely identified as credit constrained are also precisely those groups who are generally most lacking in wealth that could be pledged as collateral. For example, the young and poor do not possess sufficient collateral. In the student loan market, the private sector's inability to attach human capital in the event of default has been viewed as a basis for credit policy since at least Becker (1967). More recently, and very specifically to our inquiry, quantitative work suggests that the market for unsecured consumer credit is significantly hindered by the availability of low-cost personal bankruptcy (see, e.g., Chatterjee et al. 2007; Livshits, MacGee, and Tertilt 2007; or Athreya 2008) and by the presence of private information about borrower default risk (e.g., Sánchez 2009; Athreya, Tam, and Young 2012b). The absence of collateral is important: Full collateralization will, by definition, make both limited-commitment and private-information frictions irrelevant to lending decisions. In other words, it is the *unsecured* credit market whose functioning is likely to be most important for the populations identified above.

Given that the unsecured credit market is the one most central to the consumption-smoothing objectives of a significant share of U.S. households, the question is then: What, if anything, can be done in this market to improve outcomes? A first answer might be to make bankruptcy law tougher: If limited commitment is the problem, why not directly address the issue by making debt harder to repudiate? The problem is that while formal bankruptcy is a currently important source of credit losses, informal default remains, in practice, a clear option. Recent work of Athreya, Tam, and Young (2012b) suggests that this option seriously limits the power of bankruptcy policy to diminish incentives for debt repudiation, and, hence, to mitigate limited commitment as an impediment to unsecured consumer lending. That is, once calibrated to match the salient facts on consumption and borrowing, informal default remains a viable option that borrowers choose in the face of even modest increases in the cost of formal bankruptcy. This is seen in the data in terms of the high rate of bankruptcy and

[1] Source: www.census.gov/compendia/statab/2012/tables/12s1189.pdf. The unit is the family in the 2007 Survey of Consumer Finances, and the sample is *before either* the deleveraging or Great Recession. All dollars are 2007.

delinquency in unsecured credit in U.S. data. However, even if incentives to *default* are difficult to alter, an alternative already employed in a wide array of settings—but not yet for unsecured consumer credit—seems promising: public loan guarantees.

Such guarantees work by using public funds to defray private losses from default. In the United States, the most obvious loan guarantee programs for households are those that accompany home loans. For example, the Federal Housing Administration (FHA) and the Veterans Administration both offer loan guarantees to private lenders, and, in both 2009 and 2010, the FHA alone issued roughly two million guarantees and, as of 2010, insured nearly one-tenth of the stock of outstanding U.S. mortgage debt. Similarly, the U.S. Student Loan Administration (Sallie Mae) is active in arranging guaranteed loans, with recent flows on the order of $100 billion annually and a stock of approximately $500 billion. Loan guarantees also play a sizeable role in credit to households attempting self-employment, with the U.S. Small Business Administration's (SBA) 7a loan program guaranteeing roughly $100 billion in credit per decade since 1990.[2] However, despite their similarity to the programs we study in this article, the closest analogy might be instead to flood insurance. The reason for this will be made explicit below but stems from the fact that in our model, loan guarantees act in a manner that lowers the cost of moving consumption across both time and states-of-nature, in much the same manner as a subsidized form of insurance might.

The goal of this article is not to analyze a specific extant policy but rather to take a first step, *within a specific model class*, toward understanding the potential gains from extending loan guarantees to unsecured consumer credit markets. As such, and especially because they are currently not in use, some motivation for why one might study loan guarantees, as opposed to any number of other interventions in consumer credit markets, deserves discussion.

[2] In addition to these officially guaranteed loan programs, there is one that dwarfs them all, and this is the one operated by the two main government-sponsored enterprises, Fannie Mae and Freddie Mac. These entities issue securities to investors that come with a guarantee against default risk. The ultimate originators of mortgage credit taken by homebuyers thereby receive, in essence, a loan guarantee. While such guarantees have historically not been backed by the Treasury, they now clearly are: mortgage-backed securities investors receive Fannie and Freddie guarantees on loans with a face value of approximately $5 trillion, nearly half of the value of all household mortgage debt. See Li, (2002), Walter and Weinberg (2002), and Malysheva and Walter (2013) for more details. These articles show that the overall contingent-liabilities of the U.S. government have grown substantially over time. Lastly, beyond their sheer size, the *scope* of activities receiving guarantees is noteworthy. Endeavors ranging from nuclear power plant construction, trade credit, microenterprises, and support for female entrepreneurs all currently receive loan guarantees.

One important reason to view loan guarantees as potentially valuable in improving credit access is that under competitive conditions, loan guarantees decouple loan pricing from credit risk. This is relevant for two reasons. First, a growing body of work shows that in the absence of complete insurance markets, risk-averse households can benefit from the state contingency introduced by the option to default in bad states of the world (see, e.g., Zame 1993; Dubey, Geanakoplos, and Shubik 2005; Chatterjee et al. 2007; and Livshits, MacGee, and Tertilt 2007). What this means, intuitively, is that while nondefaultable debt requires the borrowers to always repay debt as promised, once default is allowed, matters are not so stark. Why? Because a borrower in dire straits can now invoke the option to not repay debt if doing so in the current period would expose them to severe hardship. This is what is meant by "state contingency." In a world where such an option is present, given the absence of other more explicit forms of financial contracts to help deal with risk, most notably insurance contracts against income loss, defaultable debt can be beneficial to borrowers.

Moreover, in existing work, consumers have been shown to benefit despite the presence of loan pricing that moves "against" the riskiest borrowers. However, these gains are not necessarily accessible in all *a priori* plausible environments. In recent quantitative work on the value of defaultable consumer debt, a variety of authors (such as Athreya, Tam, and Young 2009) have found that in many cases the ability of lenders to reprice loans at the same frequency as the arrival of new information on income risk undoes insurance benefits altogether. In other words, every time a consumer is hit by a persistent (but not permanent) bad shock, she will find her ability to commit to loan repayment eroded, and any borrowing she might attempt will become expensive. From the perspective of borrowers, if competitive lenders are made partially whole, they cannot "risk adjust" interest rates as much and so such loans will better assist households in consumption smoothing. Indeed, in the context of boosting aggregate consumption, researchers have recently started considering ways to direct unsecured credit to households at "favorable interest rates," with the public sector bearing default risk, exactly as would occur under the loan guarantees to consumers we study.[3]

Second, if information on borrowers has improved secularly over time, as has been suggested by many as having occurred in recent decades in U.S. consumer credit markets (see, e.g., Sánchez 2009; Athreya, Tam, and Young 2012b; and Narajabad 2012), then loans

[3] They are referred to as federal lines of credit. Details are here: www-personal.umich.edu/~mkimball/fiscal-bang-for-buck-29may12.pdf.

are now priced more accurately. This is likely to make relatively risky borrowers' access either worse or improve it by less than that of safer borrowers. Indeed, Sánchez (2009) suggests that this will be the case. Moreover, improvements in information will certainly bring the risk sensitivity of loan pricing closer to what we study below in the "full information" (FI) case. In these cases, as noted above, standard models suggest that unsecured credit will not work well as a smoothing device. Thus, policies that allow for default but break the link between credit risk and credit pricing are promising candidates to improve allocations—at least to borrowers.[4,5]

Despite their likely benefits, loan guarantees will create costs, particularly in two places. First, default rates are likely to rise, generating more deadweight loss (whether pecuniary or nonpecuniary in nature). The rise in default rates occurs for the very reason that loan guarantees "work": They lead to the systematic underpricing of loans by lenders, given their risk. Relatively larger loans will now attract relatively high-risk borrowers. As a result, the more effective any loan guarantee scheme is in spurring borrowing and consumption, the more prevalent that default and deadweight losses will be on the equilibrium path. In the context of loan guarantees for entrepreneurial ventures, the work of Lelarge, Sraer, and Thesmar (2010) documents precisely this type of response in a near-natural French experiment. As they note, "it [loan guarantee] significantly increases their probability of default, suggesting that risk-shifting may be a serious drawback for such loan guarantee programs." This inevitable tradeoff means that the real questions are: "By how much?" and "does risk-shifting happen, and if so, is it welfare-enhancing?"[6]

[4] In addition to decoupling risk and pricing, loan guarantees will also reduce *average* interest rates, all else equal. This is relevant for two reasons. First, concern with the consequences of frequent repricing of consumer debt has already led to policy changes. Most noticeably, the CARD Act of 2009 has responded by essentially *requiring* longer-term commitments from lenders in an attempt to deter frequent repricing. However, as studied by Tam (2009), such policies may carry serious side effects. In particular, average interest rates are predicted to rise substantially to offset the ability of a borrower to "dilute" his debt (much as in the sovereign debt literature). Second, average borrowing rates are likely important for welfare: Calem, Gordy, and Mester (2006) show that many U.S. households appear to use credit cards for relatively long-term financing, making the roughly 10-percentage-point cost differential between secured and unsecured interest rates quantitatively important.

[5] Andolfatto (2002) develops a simple model to illustrate how government policies (e.g., interest rate ceilings) may induce unintended outcomes (e.g., credit constraints) that generate calls for further policies to deal with these side effects (e.g., loan guarantees). A related point is that to the extent that public insurance simply crowds out familial or other forms of private insurance, the effects will be overstated. This possibility is not addressed in our article, and so should be kept in mind.

[6] With respect to the federal lines of credit noted earlier, and the subsidy that will allow the scheme to affect allocations (unlike the actuarially fair arrangement that

Additionally, tax revenue must be raised to finance transfers to lenders *ex post*. Under incomplete markets, the taxes used to finance these transfers have two *opposing* effects on welfare. First, if, as was the case in the in study of Lelarge, Sraer, and Thesmar (2010), a relatively large fraction of households faces a tax that a relatively small proportion benefit most significantly from, the introduction of a publicly funded loan guarantee program will reduce the mean level of income for many households. In particular, if it is a relatively small measure of households who run up substantial debts that, absent the guarantee, would demand high interest rates, they then receive a transfer from all other households. Second, nonregressive taxes reduce the variance of after-tax income, especially when one's *expected* lifetime income (as captured by *ex ante* uncertainty over one's eventual educational attainment) is uncertain.

While ours is not a policy evaluation article, the model class we study contains features that we believe will be essential to include in any empirically relevant policy related to consumer credit access. Specifically, our model contains a well-defined life cycle for household income that motivates credit use for intertemporal smoothing and uninsurable risk that motivates the use of credit to smooth across states. Importantly, our model features credit constraints that are endogenously derived in response to a limited commitment friction and, in other cases, to asymmetric information as well. Before proceeding, we also note that there is a distinction between what we study here and a more complicated alternative that in some ways may be more natural: We are interested in the implications of the *replacement* of the current nonguaranteed system with one in which guarantees necessarily apply to loans below a certain size threshold. Future work will, ideally, allow for the addition of guaranteed consumer lending as an option, with households self-selecting into programs with and without loan guarantees.

Our results can be summarized as follows. First, we find in our model that loan guarantees are powerful in influencing allocations: Even modest limits on qualifying loan size invite very large borrowing—as perhaps intended by proponents. However, these same limits also bring very large increases in default rates relative to a world without guarantees and, as a result, transfer resources in significant amounts from the *ex post* lucky to the *ex post* unlucky, in addition to transferring

we show is irrelevant), the idea's originator, professor Miles Kimball of the University of Michigan, argues as follows: "I am assuming the government will lose money doing this—just not as much as if they handed the money away as a tax rebate with no obligation of repayment. *The losing money part would stop private lenders cold* [emphasis ours]."

wealth across education types. Indeed, this is the key tradeoff that differentiates our work from existing research on unsecured credit, such as Chatterjee et al. (2007) or Livshits, MacGee, and Tertilt (2007). These articles find some gains from lax bankruptcy despite the fact that such rules make borrowing more expensive and, hence, tighten credit constraints, as such rules provide valuable insurance. By contrast, our article shows that loan guarantees can improve welfare despite greater ex-post deadweight loss due to bankruptcies, as they make borrowing cheaper and, hence, provide insurance as they *relax* borrowing constraints.

Second, we find that loan guarantees yield significant benefits as long as they are not too generous (whereby only small loans qualify). At low levels of the guarantee, this welfare gain is disproportionately experienced by low-skilled households who face flat paths for their average income over the life cycle and the risk of relatively large shocks. As loan guarantees are made more generous, however, higher-skilled types rapidly begin to experience welfare losses. This occurs because loan guarantees induce a transfer from skilled to unskilled households, which can be substantial, while the gains to skilled households from improved loan pricing as a result of guarantees are relatively small.

Third, we show that restricted guarantees clearly dominate unrestricted programs in terms of welfare. As noted, households in our model face risk, including that of large negative income shocks. It is plausible, therefore, that a more conditional loan guarantee program, available only to households hit by such shocks, might allow policymakers to provide the benefits from guarantees while limiting their cost. The model thus suggests that this intuition is likely to be valid.

In the case of asymmetric information, the size of this friction will be *endogenous* as well and will depend on how heterogeneous borrowers are, not only in terms of both exogenous shocks, but also in terms of endogenously determined and unobservable net asset positions. One of the earliest studies of loan guarantees is that of Smith and Stutzer (1989). These authors show in a stylized model with two types of borrowers (high risk and low risk) that the reduction in the sensitivity of loan interest rates to default that accompanies loan guarantees also reduces the high-risk types' incentives to reduce their borrowing inefficiently simply by mimicking the low-risk types. This incentive effect contributes positively to the welfare impact of loan guarantees. And while private information may not currently be a crucial problem in U.S. consumer credit markets, given extensive recordkeeping and information sharing via credit bureaus, it was likely present both in the United States in earlier periods (see Sánchez 2009 and Athreya, Tam, and Young 2012b), and plausibly remains an impediment in developing countries currently. A goal of this article is to measure the effects of loan guarantees under these more difficult circumstances. We find that under private information, the gains from guarantees are meaningfully larger than in the absence of private information, quantitatively consistent with the prediction of Smith and Stutzer (1989).

It is important to recognize that loan guarantees can *only* matter for allocations and welfare when debt is imperfectly collateralized. Even when a loan guarantee program is nominally targeted at a secured form of lending, such as mortgage loan guarantees, they can only alter allocations because there is a positive probability of the loan becoming at least partially unsecured *ex post*. This leads us to focus on the effects of introducing guarantees for uncollateralized consumer lending, which is the most prominent form of unsecured credit.[7] Nonetheless, our setting will clearly be informative for the effects of loan guarantees in any market in which there exists states of nature where repayment is less attractive than paying the costs of default.[8]

Our article is linked to three strands of research in public interventions in credit markets. First, our focus on consumer credit with default risk connects this article closely to recent work of Chatterjee et al. (2007); Livshits, MacGee, and Tertilt (2007); Athreya, Tam, and Young (2012b); and Narajabad (2012). In this line of work, however, guarantees are not studied, but both voluntary default and asymmetric information have been shown to matter for the allocation of consumer credit in the absence of guarantees (see, e.g., Sánchez 2009 and Athreya, Tam, and Young 2012b). As noted earlier, our research is novel in studying a distinct mechanism from this strand of work, whereby welfare can be improved by relaxing constraints, rather than tightening them through the promotion of debt forgiveness.

Second, our work is clearly connected to more recent research on *quantitative* analysis of the allocational consequences of loan guarantees. This work began, to our knowledge, with Gale (1990) and was followed by the rich, fully dynamic, and relatively tractable incomplete-market models developed in Li (1998) and Jeske, Krueger, and Mitman (2010). The last article is the first to focus centrally on credit markets in a consumption smoothing context. However, with respect to

[7] See, e.g., Federal Reserve Release G.19: www.federalreserve.gov/releases/g19/Current/g19.pdf.

[8] Sometimes, these costs are primarily those arising from the surrender of tangible collateral that, ex post, becomes less valuable than reneging on the repayment obligation, e.g., as recent house price declines have done (Ghent and Kudlyak 2011). In other cases, default implies the destruction of intangible collateral, as described above. But in all cases, loan guarantees fundamentally concern unsecured lending.

modeling default, in all the preceding work, default is involuntary.[9] Our article is also related to recent work of Jia (2013), who in turn builds on Li (1998) to study loan guarantees for firms in a setting where the government's relative (and absolute) financing advantage in recessions can be put to use by providing loan guarantees for small businesses. Our focus, by contrast, is on households, and a main goal is to provide a quantitative analysis that is rich and aims to evaluate consumption smoothing, not investment in small business as these other articles do.

Third, our work relates to an earlier, relatively stylized class of articles that focus on the role of interventions, including loan guarantees, on outcomes for a general problem of risky-investment in static or near-static settings under asymmetric information. Key landmarks in this category are Chaney and Thakor (1985), Smith and Stutzer (1989), Gale (1990), Innes (1990), and Williamson (1994).[10] Most of this work abstracts from the financing of guarantee programs as well. By contrast, these costs will feature prominently in our analysis.[11]

1. ILLUSTRATIVE MODEL

The friction faced by households stems, ultimately, from their inability to explicitly insure income shocks and their inability to credibly promise

[9] In future work, we aim to analyze the role of guarantees for mortgage lending. However, the central role of aggregate risk in driving home-loan default makes a full quantitative analysis that satisfactorily incorporates the forces we do allow for here—partially endogenously asymmetric information and limited commitment—currently infeasible. But that model would have the same fundamental structure as the one developed here.

[10] In related work, Lacker (1994) investigates whether adverse selection problems necessarily justify government intervention in credit markets. When cross-subsidization between private contracts is not feasible, intervention is generally welfare-improving.

[11] While substantially different than our model, it is important to note the early work of Smith and Stutzer (1989), who provide a simple argument for the use of loan guarantees in unsecured commercial credit markets—compared to direct government loans or equity purchases, loan guarantees are the only option that does not worsen the private information problem. The interest rate reductions apply to all risk types, so high-risk types do not find any particular advantage, beyond what they already have, for pretending to be low risk. Other programs, such as direct loans to those unable to obtain credit (who are low risk in their model), will lead to additional incentives by high-risk borrowers to claim the contracts intended for low-risk ones, a situation that is harmful to efficiency. Two important distinctions between our work and theirs are worth keeping in mind—the nature of the commitment problem and the issue of government revenue balance. In Smith and Stutzer (1989), limited commitment is a trivial consideration: Default occurs when the borrower receives zero income and is costless (in terms of direct costs). In contrast, U.S. bankruptcy procedures are voluntary and clearly not costless: There is a filing fee in addition to substantial time costs and some form of stigma/nonpecuniary costs appear relevant as well (see Fay, Hurst, and White [1998] or Gross and Souleles [2002]). Smith and Stutzer (1989) do not consider the financing of such payments; any welfare gains from the guarantee could easily be wiped out by the cost of taxation. In contrast, a central aspect of our analysis is the requirement that transfers required to implement loan guarantees be paid for via taxes.

to always repay loans. Given default risk, competitive lenders will be forced to price loans in a way that allows them to break even. As a result, in general, households with differing levels of default risk will face different prices for credit. However, the fact that loans in our model will be priced to reflect default risk also means that some borrowers will find themselves facing expensive credit terms precisely when they most need to borrow. It is these groups who will find guarantees most helpful.

Before turning to the quantitative setting in the next section, it is useful to describe a simple two-period variant of our model to more clearly identify the types of individuals who are affected by risk-based pricing that, by definition, makes borrowing expensive when future income levels might remain low, and who may therefore gain from loan guarantees. Let c_i denote consumption in period $i = 1, 2$, e_i denote the endowment of the consumption good received by the agent in period $i = 1, 2$, and $d_2 \in \{0, 1\}$ denote the default decision in period 2. Defaulting implies that the consumer incurs a nonpecuniary cost $1/\lambda$. To remain consistent with the quantitative model on which the final results are based, and for mnemonic ease, *a high value of* λ *implies a high risk of default*, all else equal, because it implies a low value for the term $1/\lambda$, which is what gets subtracted from utility in the second period in the event of default. This cost is a stand-in for the variety (and entirety) of costs associated with defaulting and is meant to tractably encompass not only the explicit costs (e.g., bankruptcy filing costs, legal costs, etc.) but also the nonpecuniary costs (e.g., difficulty renting durable goods, obtaining employment, emotional distress, etc.). The dependence of default risk on loan size leads loan prices to depend on loan size.

Households are modeled as borrowing through the issuance of debt with a face value $b < 0$. "Face value" refers to the amount that the household is obligated to repay and is the value that it would deliver if it did not default. However, the household may, in period 2, elect to exercise its default option. As a result, the face value of debt, by virtue of being risky, will be discounted by lenders. The term $q(b) \in [0, 1]$ is the discount factor applied to a debt issuance of face value b and is determined by competitive markets. To see how this discount is determined, consider a lender wishing to price a loan with face value b. Let the default probability for this loan be given by $\pi(b)$. Thus, the expected value of the loan is $(1 - \pi(b))b$. In facing this, the lender must decide what discount $q(b)$ to apply. Competition among lenders implies that the discount allow the bank to, at best, break even on average. This implies that $q(b)b$, the real expected value of resources transferred to the borrower, must have the same cost for the lender to obtain as the expected value of the loan. Let the cost of funds for the lender be given by $(1+r)$. Thus, the cost of making a loan with face value b with discount $q(b)$ is $(1+r)q(b)b$. Equating this with the expected value of the loan, $(1-\pi(b))b$, and simplifying gives

$$q(b) = \frac{1-\pi(b)}{1+r}.$$

This is intuitive. As $\pi(b)$ rises, $q(b)$ falls and hits zero when $\pi(b)$ reaches unity (certain default). For loans where $\pi(b) = 0$, the discount is simply $\frac{1}{1+r}$, the competitive price of a risk-free loan.

Given this credit market and endowment structure, households choose consumption, borrowing, and saving to maximize standard expected utility preferences:

$$\max_{c_1,\{c_2(e_2),d_2(e_2)\}} \{u(c_1) + \beta E_{e_2}[u(c_2(e_2)) - d_2(e_2)/\lambda]\}.$$

Default risk arises in the model from the fact that endowments are probabilistic and structured as follows: All households receive e_1 in the first period and this value is a known *constant*. Households face uncertainty with respect to income *only in the second period*—e_2 is drawn from a two-point distribution: $e_2 \in \{e_L, e_H\}$ with probabilities p_L and $1 - p_L$.

Households may save or borrow in the first period, denoted b, with $b_1 > 0$ corresponding to saving and $b_1 < 0$ being borrowing. In period 2, they first draw income e_2, and then elect to default $(d_2(e_2) = 1)$ or not $(d_2(e_2) = 0)$. Note clearly that if they default, they repay nothing (default is total). This is a useful simplification, but can be relaxed to allow partial default. If they do not default, households must repay the face value they issued in the first period, b. As a result, the households' choices are restricted by the following pair of budget constraints. In period 1,

$$c_1 + q(b)\,b \leq e_1.$$

In period 2, they face

$$c_2(e_2) \leq b(1 - d_2(e_2)) + e_2.$$

Consider a case with two types of households. Let one type be those whose second-period endowments e_2 have a high mean relative to their period 1 value and (relatively) *small variance*. In other words, income in the future is expected to be higher than today and relatively safe as well; this group roughly corresponds in the data to relatively highly educated borrowers. This group values access to credit because it helps them bring their high, safe, future income into the present. Thus, loan prices for such households will be at the risk-free rate for a relatively wide range of borrowing levels, as the household will elect

Figure 1 Equilibrium in Two-Period Model

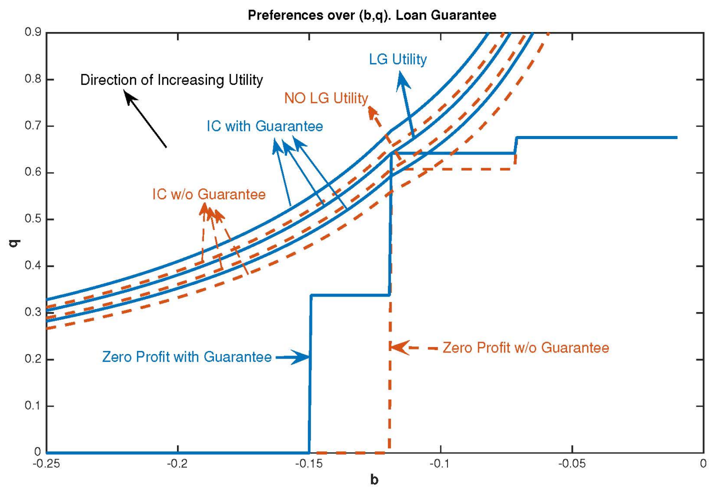

to repay irrespective of the realization of e_2, and then fall abruptly when the loan size reaches a threshold where households would first begin to consider default. For households facing only small uncertainty about e_2, this means that once default in one state becomes attractive, it generally becomes attractive in the other state (since they are, by construction, similar). This is easiest to see in the limit where e_2 is known in period 1 with certainty: For a given loan size b, default in period 2 either occurs none of the time or all of the time.

The second type of household we are interested in has small mean and large variance of e_2; one can think roughly of this group as being relatively less educated and facing greater risk of unemployment in period 2. For these households, borrowing is not particularly useful, but if undertaken, can yield more variation in terms because the default decisions will differ more substantially across realizations of second period income e_2.

Figure 1 shows a typical situation faced by either type of household.[12] The indifference curves are monotone (over the range of interest

[12] The figures represent outcomes under the following parameterization for the endowments of each group. For the first group of agents, three conditions hold: (i) the amount that can be feasibly repaid in the bad state is large (that is, e_L is relatively big); (ii) the household will default in both states under risk-free pricing (in the case where λ is small relative to e_L and e_H, and the latter are close together); and (iii) the

at least), and reflect the fact that a household can, in principle, receive additional consumption today in two ways: hold b fixed, as long as it faces a higher q, or increase borrowing b, but accept that the discount q will fall as default risk rises (though not so rapidly that $q(b)b$, which is what the household receives, falls).[13] At the optimum the household is constrained, in the sense that additional borrowing is desired but not feasible due to the increase in the probability of default; this situation will be typical in the quantitative model as well. Thus, local to that optimal b there are welfare improvements available to households if q can be held fixed while b is increased. This is an obvious point, perhaps, but it is useful to keep in mind as it is the source of the ability of loan guarantees to assist households in smoothing consumption.

Loan Guarantees

We now introduce a publicly funded (via taxes that, in this section, will remain unmodeled) loan guarantee into this model economy. Loan guarantees will be defined by two parameters: (i) a "replacement rate" θ that determines the fraction of defaulted obligations b that the lender receives as a transfer from the government, and (ii) a "coverage limit" ϑ that determines the largest (riskiest) loan that the government will insure. Only loans smaller than, or equal to, ϑ in size qualify for any compensation; lenders making loans larger than the ceiling receive nothing in the event of default.[14] Households may borrow more than

household would borrow if asset markets were complete ($\beta(1+r) < 1$ and $E(e_2)$ significantly larger than e_1). This group is (weakly) harmed by the intertemporal disruptions that default options create; because the two states tomorrow are very similar, the household would either default in both states and thus be unable to borrow at all ($q = 0$), or it would not default in either state and thus care not at all about default options. As a result, the outcome may be worse than if bankruptcy were banned since, in the absence of a default option, feasibility would permit borrowing against the (relatively high) value of e_L.

For the second type's endowments, three conditions are assumed: (i) the amount of debt that can feasibly be repaid in all states is small (that is, e_H is low); (ii) the household will default only in the low state (λ intermediate and e_L and e_H far apart); and (iii) the household would borrow if asset markets were complete ($\beta(1+r) < 1$ and $E(e_2) = e_1$). A member of this group can gain from the default option because she actually can borrow more with a bankruptcy option, as she does not intend to repay in the low state; thus, feasibility is limited only by the amount that can be repaid in the high state and additional consumption smoothing is feasible. This is a manifestation of what might be referred to as a "supernatural" debt limit, as opposed to the "natural" debt limit (e.g., Aiyagari 1994): Feasibility involves what can be repaid in the *best* state instead of the *worst*.

[13] Although not shown in the figure, the typical indifference curve turns upward at very low levels of b, but these lie well outside the budget set.

[14] This program assumes that the household cannot obtain a qualifying loan of size greater than ϑ by visiting multiple lenders; that is, we attach the qualification criterion to the borrower, not the lender.

ϑ, but if they do, the discount on these loans will jump discretely downward as expected rewards to lenders fall discretely due to the "non-conforming" nature of loans exceeding the program limit.

Given that the loan guarantee covers θ percent of the repayments lost to default for the portion of any loan less than ϑ, competitive pricing of loans for a conforming loan—one with face value $b \leq \vartheta$—must obey

$$q\left(b\right) = \frac{1 - \pi\left(b\right)}{1 + r} + \frac{\pi\left(b\right)}{1 + r}\theta.$$

Non-conforming loans are priced exactly as if there was no loan guarantee program in place. In the absence of taxes needed to compensate lenders for default under any guarantee program, it is clear that *both* types of households would gain from the introduction of loan guarantees. Assuming default probabilities don't change, the guarantee increases the bond price for the first group from 0 to $\frac{\theta}{1+r}$ and for the second group from $\frac{1-p_L}{1+r}$ to $\frac{1-p_L+p_L\theta}{1+r}$. This increase expands the set of feasible consumption paths and raises welfare; default probabilities will not increase if θ is small enough due to the discreteness of the income process. To illustrate how a loan guarantee works, Figure 1 shows how the pricing function shifts weakly upward, which clearly raises utility because the household is currently constrained and the deadweight loss from default is the same.

Asymmetric Information

Our analysis has so far focused on limited commitment alone as an impediment to credit access. We now allow for asymmetric information to further hinder lending. To adapt the model to deal with asymmetric information, suppose now that default risk varies according to some characteristic that is not observable to the lender; for concreteness, let there be two such groups, and keep in mind that within each, there are two groups of borrowers with respect to their current cost of default λ. Private information forces (barring a rich menu of screening contracts) the lender to offer a uniform pricing function to both types of households based on the invariant measure of each type (let $\delta \in (0, 1)$ be the measure of the first type); the function is contingent on the costly signal b sent by the household.[15] The pricing function without

[15] We assume that no costless and credible signals are available and that some additional hidden characteristic, such as initial wealth, thwarts the lender's attempts to infer from b the exact value of π.

the guarantee would be

$$q(b) = \delta \frac{\widehat{\pi}(b|1)}{1+r} + (1-\delta)\frac{\widehat{\pi}(b|2)}{1+r},$$

where the hatted variable $\widehat{\pi}$ reflects the fact that under asymmetric information default, probabilities are no longer necessarily known with certainty since the agent who asked for the loan is not necessarily known, in equilibrium, with certainty. Instead, default risk is an imperfect estimate that reflects the uncertainty over which agent-type attempted to take a given loan. For a risk-neutral lender, what is then relevant is the conditional probability of a given loan request having come from either type of borrower (and reflecting the fact that each type will not default with the same probability at any given level of debt). In a more elaborate model (such as the one used later for quantitative analysis), the debt level b would lead lenders to update the estimate of default risk since not all types would find it optimal to issue b; in such a case one can think of lenders computing conditional probabilities of the borrower being of a particular type given their requested loan size b (that is, updating δ) and using those probabilities to compute default risk. Equilibrium then requires that updated beliefs are consistent with the population of borrowers of a given type issuing b.

The "bad" type of borrower—that is, the borrower with the high value of $\widehat{\pi}$—will want to reduce b in order to look more like the good borrower, all things being equal. As discussed more completely in Athreya, Tam, and Young (2012b), pooling is potentially an equilibrium if the pricing function is relatively flat just to the right of the equilibrium choice; in that case, the indifference curves of both types lie above the break-even curve for the lenders so deviations to lower debt levels do not occur. Separating equilibria occur when pricing functions are steep (relative to indifference curves), because then the good type would be better off reducing b while the bad type would not. Loan guarantees reduce the desire of bad types to pool with good types because they break the link between pricing and type; this disincentive is welfare-improving because it improves the allocation of consumption, and so under asymmetric information loan guarantees will have even better welfare properties. But as before, we must consider whether the costs outweigh the gains, and under asymmetric information the costs will increase more than under symmetric information because default is initially lower. Whether the costs or benefits are larger is the main focus of our quantitative model, which we describe in the next section.

The Irrelevance of Actuarially Fair Loan Guarantees

In this article, we study fully subsidized guarantees. However, in practice, many loan guarantee schemes ask the borrower to pay the guarantee fee (such as SBA loans and FHA-guaranteed home loans).[16] An important point we now develop is that any private loan guarantee scheme that is also actuarially fair, and therefore will survive competition, will necessarily be irrelevant.

To see this result, consider a competitive economy in which, notionally, the borrower is obligated to pay the loan guarantee fee, as observed in practice. Let $\tau(b)$ be the insurance premium on a loan with face value (i.e., what is paid outside of default) b. Let q_f be the reciprocal of the risk-free interest rate, i.e., $q_f = 1/(1+r)$, where r is the risk-free rate of interest on savings. As before, let $\pi(b)$ be the probability of default on a loan of size b.

A borrower who issues b units of face value then gets, after the insurance payment of $\tau(b)$, $q_f b - \tau(b)$ units of resources in period t, and is free to default or not in period $t+1$. So what does the guarantee fee have to be? If it is set to break even across all borrowers of the given type of borrower who issued b units of debt, then the premium must be $\tau(b) = \pi(b) b q_f$ (the last term appears since the lender will only get paid next period and so must discount), which equals the expected loss on the loan. Therefore, the net resources an agent gets for issuing b units of debt, after paying the loan guarantee premium, is $q_f b - \pi(b) b q_f$, or $q_f b\,(1 - \pi(b))$. But this is exactly the pricing function that would arise in a competitive setting *without* guarantees.

This result follows naturally from competition between lenders: If the borrower pays the insurance premium and leaves the lender insured, the loan is then risk free to the lender. As a result, a lender can, under competitive conditions, only charge the risk-free rate for the loan. Thus, if loan guarantee schemes are to matter for allocations, they must carry a subsidy with them, such as the one that comes with public provision of the guarantee. This implication is why we study fully subsidized guarantees (no premium). Our approach ensures that guarantees don't merely lead to a reinterpretation of existing contracts, but rather are capable of changing household budget sets.[17]

---

[16] For example, the FHA loan guarantee fee structure is given here: www.sba.gov/community/blogs/community-blogs/small-business-cents/understanding-sba-7a-loan-fees.

[17] Our neutrality result holds in the asymmetric information signaling model we study here. Whether it holds in a screening environment (such as Sánchez [2009]) is unclear, since it may be possible to offer (q, τ) pairs that separate types.

In our quantitative model we study programs that insure only a fraction θ of loan losses, using tax revenue to fund payments to lenders. The argument here is unchanged—the private sector cannot offer meaningful guarantees for the fraction $1-\theta$ of the loan that the government does not cover.

2. QUANTITATIVE ANALYSIS

Given that the implications of guarantees depend on the opposing forces we have isolated above, we now turn to the quantitative analysis of guarantees for consumer credit in order to determine the ultimate effects they may have. The general framework we employ is a standard life-cycle model of consumer debt with default, and aside from the budget-constraint-related complications arising from loan guarantees, the model we use is essentially identical to that of Athreya, Tam, and Young (2012b).[18]

Because the contribution of our article lies in its application of a standard type of model to understand loan guarantees, we relegate the technical description of the model and its parameterization, *except for the pricing implications of guarantees*, to the Appendix. We will note here only the essential model features, as follows. First, there is a large number (continuum) of households who each live for a finite number of periods. Second, households differ, ex ante, and permanently, in their earnings prospects. This is to reflect differences in the population with respect to educational attainment, and we will therefore allow for three classes of households: those who have not completed high school ("NHS"), those who have completed high school ("HS"), and those who have completed college ("Coll"). Third, households face shocks to their labor productivity throughout life, with shocks having both transitory and persistent components. Fourth, households face the risk of needing to spend suddenly in a manner that is involuntary. These "expenditure shocks" capture the idea that households sometimes face sudden health care expenses or legal obligations that make spending more or less involuntary. Such expenses have been viewed by researchers as relevant in at least a portion of observed default so, for completeness, we include them. Fifth, households have access to credit markets in which they may borrow (and where they will receive guarantees), and may save in a single risk-free asset that earns a constant interest rate.

The features just described lead households in our model to solve a consumption-savings problem over a well-defined life cycle in which

[18] Other related work includes Chatterjee et al. (2007) and Livshits, MacGee, and Tertilt (2007), though the former uses an infinite horizon.

Table 1 Model Versus Data

	Model	Target/Data
Discharge/Income Ratio	0.2662	0.5600
Fraction of Borrowers	0.1720	0.1250
Debt/Income Ratio | NHS	0.1432	0.08
Debt/Income Ratio | HS	0.1229	0.11
Debt/Income Ratio | COLL	0.0966	0.15
Default Rate | NHS	1.237%	1.228%
Default Rate | HS	1.301%	1.314%
Default Rate | COLL	0.769%	0.819%

their productivity has both deterministic and stochastic components. The risks faced by households include, most importantly, those that alter labor earnings, but also those that govern the marginal value of default. Because of default risk, lenders discount household promises according to their estimate of repayment likelihood.

We assume that the economy is small and open, so that the risk-free rate is exogenous. There is a representative firm that takes prices and wages as given, and demands labor as a function of its relative price (the wage). In equilibrium, the wage rate is part of the fixed point with the property that the representative firm's first-order condition governing the level of desired labor input at that wage is equal to household labor supply at that wage.

We will restrict attention throughout to stationary equilibria of the model—i.e., steady states. Stationary equilibria are, as usual, those outcomes where prices and the distribution of households over the state space remain constant under optimal household and firm decisionmaking. More intuitively, aggregate outcomes from our model can be viewed as averages (across households of all ages) for a single large cohort whose members each begin life with zero wealth, draw initial shocks from the unconditional distribution of shocks, and then draw shocks according to the stochastic processes we will specify further below.

To generate predictions from the model, we parameterize the environment to match a set of salient features, as displayed in Table 1. These targets are the ones most relevant for our analysis and collectively cover debt use and default-related features. The main message of Table 1 is that our baseline quantitative accurately captures key features of the data, and hence can be seen as reliable in its implications for the counterfactual exercises we will examine for alternative loan guarantee programs. Table 2 displays the specific parameters used in the quantitative model. We turn next to a description of how, in

Table 2 Calibration

Parameter	Value	Parameter	Value
x_{low}	0.0000	Prob(x_{low})	0.9244
x_{median}	0.0888	Prob(x_{median})	0.0710
x_{high}	0.2740	Prob(x_{high})	0.0046
$\lambda_{\text{low}}^{NHS}$	0.7675	$\lambda_{\text{high}}^{NHS}$	0.9087
$\lambda_{\text{low}}^{HS}$	0.7309	$\lambda_{\text{high}}^{HS}$	0.9320
$\lambda_{\text{low}}^{Coll}$	0.7830	$\lambda_{\text{high}}^{Coll}$	0.9017
$\pi_{HH}^{\lambda} = \pi_{LL}^{\lambda}$	0.9636	J	65
j^*	45	Λ	0.0300
σ	2.0000	ϕ	0.0300
α	0.3000	r	0.0100

the quantitative model, loan guarantees will affect the terms on which credit is available to households.

Loan Pricing and Loan Guarantees

Loan guarantee regimes are defined by two parameters: the "replacement rate" θ and the "coverage limit" ϑ. Only loans smaller than ϑ qualify for any compensation; lenders making loans larger than the ceiling receive nothing in the event of default.[19] Conditional on default occurring, the lender, having made a loan of qualifying size, will receive partial compensation whereby the fraction θ will be paid to the lender for each unit of face value.[20]

We focus throughout on competitive lending whereby intermediaries utilize all available information to offer one-period debt contracts with individualized credit pricing that is subject to meeting a zero profit condition. Denote by I the information set of lenders. The information set is, under symmetric information, the entire state vector as understood by the household. In other words, I is the set of items that *fully summarizes* default risk to a lender for whatever level of borrowing the household requests.

Under asymmetric information, only a subset of these features is known. Specifically, the household's cost of default will not be known to the lender, who must then conjecture it, given the signal embedded

[19] This restriction seems to be standard practice in markets where some form of loan guarantee program exists. For example, FNMA (Fannie Mae) will not issue guarantees on loans that do not conform to their pre-set standards, which include a restriction on the loan-to-value ratio.

[20] As we noted earlier, qualification actually applies to the total debt of the borrower, not the total loan emanating from any one lender. An implicit assumption is therefore that this debt burden is observable.

in the households requested borrowing amount b. Denote this estimate by $\widehat{\pi}(b, I)$. It is the lenders' best estimate of a household's default risk on a debt issuance with face value b. Of course, $\widehat{\pi}(b, I)$ is identically zero for positive levels of net worth, and is also equal to 1 for some sufficiently large debt level. Denote by r the exogenous risk-free saving rate. In order to capture the costs associated with lending, we will also assume henceforth that lenders face a constant (i.e., proportional) transaction cost when lending. This implies that $r + \phi$ is the risk-free borrowing rate.

Given the preceding and the loan guarantee program parameters (θ, ϑ), the break-even pricing function on loans $(b < 0)$ will depend on the size of the loan relative to the guarantee limit ϑ as follows. Letting, as before, $q(\cdot)$ denote the price, or discount, applied to a bond issuance by a household, we note first that since only loans smaller than the guarantee ceiling entitle lenders to compensation, qualifying loans (those with $b \in (-\vartheta, 0)$) are priced as follows:

$$q(b, I) = \psi_{j+1|j} \left[\frac{(1 - \widehat{\pi}(b, I))}{1 + r + \phi} + \frac{\widehat{\pi}(b, I)\,\theta}{1 + r + \phi} \right] \text{ if } 0 > b \geq -\vartheta. \quad (1)$$

The first term, $\psi_{j+1|j}$, is new and represents the conditional probability of surviving to age $j + 1$ given survival to age j. Its presence in the pricing of loans reflects the fact that repayment occurs, if at all, only if the borrower survives. Conditional on survival, the payoff to a loan of face value b will be complete in the event of no default, which occurs with probability $1 - \widehat{\pi}(b, I)$, and partial, according to the guarantee, if default occurs. These payoffs are then discounted according to the cost of funds, inclusive of transactions costs, $1+r+\phi$. For any loans exceeding the guarantee qualification threshold, lenders will receive nothing in the event of default. As a result, the preceding zero-profit loan price collapses (the second term goes to zero), yielding the simpler expression

$$q(b, I) = \psi_{j+1|j} \left[\frac{(1 - \widehat{\pi}(b, I))}{1 + r + \phi} \right] \text{ if } 0 > -\vartheta > b. \quad (2)$$

Lastly, savings are trivial to price, as they carry no transactions costs or default risk. Therefore, for $b \geq 0$, we have

$$q(b, I) = \frac{1}{1 + r} \text{ if } b \geq 0.$$

As for the fiscal implications of loan guarantees, the budget constraint for the government is straightforward: The flat tax rate on all earnings must be sufficient to make payments to all lenders whose borrowers defaulted on their debts.

Allocations

We will study four types of allocations. First, we examine our benchmark setting, where information is symmetric and there is no loan guarantee program. Second, we introduce various loan guarantee programs and examine how credit market aggregates, default rates, and welfare are altered. Third, we relax the assumption of symmetric information and study allocations without loan guarantees; in this setting we permit lenders to use all *observable* characteristics to infer as much as they can about borrowers. Finally, we examine the introduction of loan guarantees into this asymmetric information environment. We will refer to these four allocations as full information without loan guarantees (FI), full information with loan guarantees (FI-LG), asymmetric or "partial" information without loan guarantees (PI), and partial information with guarantees (PI-LG).

To preview the results, we find that introducing a small loan guarantee program into a symmetric information economy (comparing FI with FI-LG) can benefit all households, independent of type, but that increasing generosity quickly eliminates the gains for skilled types. To be clear, our measures of welfare, throughout, will be "ex ante": They are the gains of losses that a household entering the economy at the beginning of its life, i.e., as "newborns", as it were, would obtain. In the environments with asymmetric information (comparing PI to PI-LG), welfare gains are larger for any given generosity, but the same pattern emerges. Thus, a general lesson from these experiments is that loan guarantees are welfare-improving, and in fact can be welfare-improving for all newborns, provided they are not too generous.

Symmetric Information

As we noted at the outset, unsecured credit markets are most vital for the consumption smoothing needs of the least wealthy members of any society. This is obvious for any household with liquid wealth, but even those whose wealth is illiquid will, in general, be able to pledge at least a portion of that wealth to obtain credit. Moreover, as we noted, existing work suggests that information asymmetries may not be central, relative to the limited-commitment problem, in explaining current U.S. unsecured credit market activity. We therefore first isolate the role that loan guarantees and limited commitment play in dealing with the effects of such a friction by studying the model under symmetric information. Moreover, since loan guarantee programs require two parameters for their specification, we simplify the results by focusing throughout the analysis—and unless otherwise stated—on the case

Table 3 Aggregate Effects of Loan Guarantee Program

	$\theta = 0.50$				
ϑ	0.00	0.10	0.30	0.60	0.70
τ_2	0.0000	0.0005	0.0174	0.0531	0.0664
Discharge/Income Ratio	0.2662	0.2691	0.5430	0.9907	1.1172
Fraction of Borrowers	0.1720	0.2039	0.2400	0.4023	0.4466
Debt/Income Ratio | NHS	0.1432	0.1648	0.4765	0.6562	0.7118
Debt/Income Ratio | HS	0.1229	0.1372	0.3707	0.5369	0.5934
Debt/Income Ratio | COLL	0.0966	0.1140	0.2532	0.3858	0.4124
Default Rate | NHS	1.237%	1.768%	11.651%	19.691%	20.877%
Default Rate | HS	1.301%	1.751%	11.658%	16.609%	17.836%
Default Rate | COLL	0.769%	0.987%	5.668%	11.569%	13.100%

where the replacement rate is set to cover 50 percent of lender losses, i.e., $\theta = 0.5$.

Allocations and Pricing

Our first main result is that loan guarantees are powerful tools in altering the use of unsecured credit. In Table 3, we see that as we move away from the case with no loan guarantees ($\vartheta = 0$), equilibrium borrowing rises for all households and the increase in debt is nonlinear. In particular, for small qualifying loan sizes (e.g., $\vartheta = 0.1$, or $4,000), allocations are fairly similar to a setting with no guarantees. In large part, this similarity reflects the presence of bankruptcy costs that serve as a form of implicit collateral. In particular, the fixed cost component of bankruptcy (Λ) will ensure the existence of a region of risk-free debt. Therefore, under a small qualifying loan size, few individuals will see their access to credit substantially altered; in fact, setting $\vartheta < -\Lambda$ would have no effect on credit, since those loans are always risk free. Once the qualifying loan size grows large enough to make large loans "cheap" relative to default risk, matters are different. The compensation to lenders for default disproportionately subsidizes large loans and thereby generates the significant additional default seen in Table 3.

The differential distortion to loan pricing is displayed in Figure 2, and our model suggests that this feature helps account for the striking distributional consequences seen in Table 3. In particular, borrowing behavior changes in different ways across the education groups. Relative to income, debt rises by far the most for the least skilled (NHS) households. The differential increase in debt relative to income for the lowest skilled is also reflected in the disproportionate rise in bankruptcy rates within this group. While remaining modest under small qualifying loan ceilings, more generous ceilings create greatly increased

Figure 2 Pricing Functions with and without Loan Guarantees

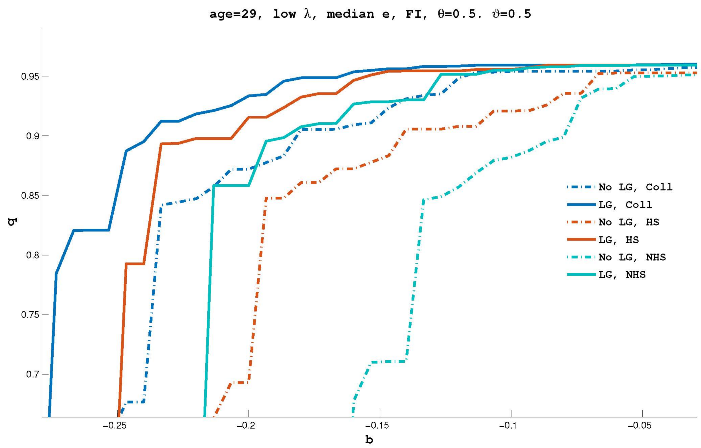

default rates. The preceding suggests in part that the pricing of debt is a meaningful barrier to nearly all households, but especially NHS households. An additional force at work is that high-skilled households have less reason to use unsecured credit beyond early life. As a result, any distortion in the pricing of debt will affect them less than their NHS counterparts. In particular, all NHS households who have income below their age-specific mean will find "artificially" cheap credit useful, while the well-educated, many of whom wish to save less for precautionary reasons (i.e., to hedge against possible bad outcome for income in the future) and more for life-cycle (keeping consumption stable as household age) reasons, will be less sensitive to credit conditions. The latter insensitivity arises from the fact that an individual or household with a pronounced hump in their average earnings shown in Figure 3 will wish to save less or borrow when young, and save in the peak earnings years in order to have a comfortable retirement period (which the model captures by making households incapable of working beyond a certain age). Lastly, under high ceilings for qualifying loan guarantees, the high tax rate will also meaningfully compress the intertemporal profile of earnings, and therefore attenuate the incentives of the skilled to borrow for pure life-cycle smoothing. This will make loan guarantees even less valuable.

Figure 3 Efficiency Units of Labor

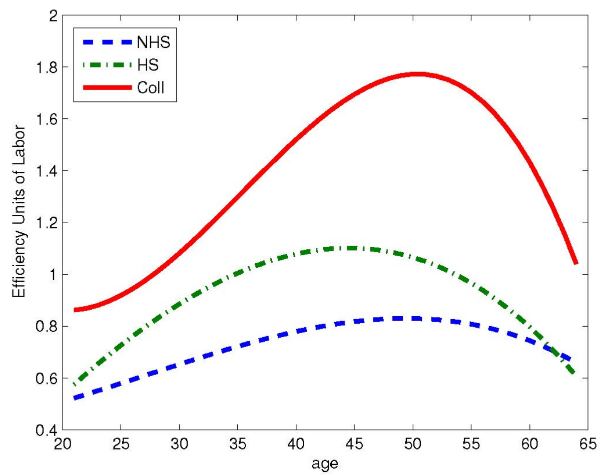

Welfare

Having shown results suggesting that loan guarantees will likely have sizeable and nonlinear effects on credit use and default, we now turn to the issue that motivated us at the outset: Can loan guarantees, by breaking the link between credit risk and loan pricing, improve welfare? And if so, for whom? Our metric for measuring welfare is standard: it is the change to consumption at all dates and states needed to make the household indifferent, in terms of *ex ante* expected utility, between the benchmark economy and the one with loan guarantees.

A fact that will be important for welfare is that households in our economy who borrow are always *constrained*. Figure 4 plots indifference curves in (b, q) space along with the zero-profit pricing function; the optimal amount of borrowing and the resulting price lies where the highest indifference curve intersects this zero-profit curve. At this point, the slope of the indifference curve is strictly smaller than the slope of the pricing function (which is infinite). This implies that borrowing more is desirable at the current interest rate, but the increase in the default rate that a marginal increase in b would generate means that lenders must charge a higher rate. As a result, by reducing the slope of the pricing function at the optimal point, loan guarantees can improve utility at the margin. What we are contemplating, however, are not

Figure 4 Optimal Choice of Borrowing

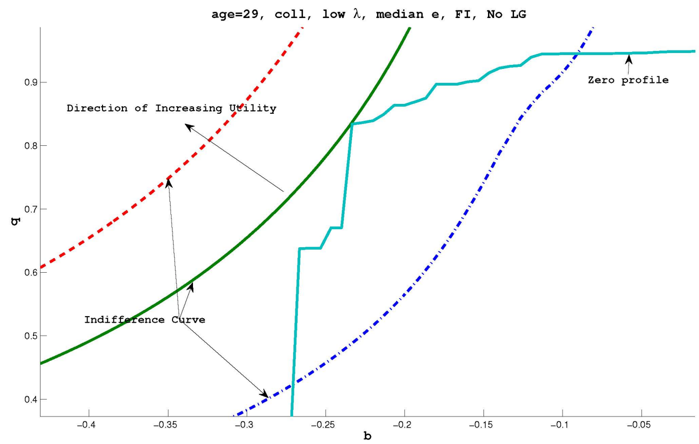

marginal changes; thus, whether a discrete change is welfare-improving is a quantitative question.

We see first, from Table 3, that more generous loan guarantees come with higher taxes, and that the taxes also naturally reflect the nonlinearity in household borrowing and default behavior. However, not all households pay the same amount in taxes, and, as we noted, proportional taxes—which are used here—will by themselves provide some risk-sharing benefits. Moreover, the loan guarantee may allow for an effective form of insurance for some households, especially the low-skilled. The transfers from loan guarantees come "at the right time" for households, but require households to pay a cost, which, intuitively, is akin to a deductible on an insurance policy. Therefore, while households pay more in taxes under a generous loan guarantee scheme, they also receive transfers in a manner that is effective in providing insurance.

Turning to welfare in Table 4, we see that this is precisely what is at work. In this table a positive value indicates a gain to welfare from moving to loan guarantees, and vice versa. In particular, we see that generous loan guarantee schemes mainly represent transfers to the very unskilled. These are, in turn, the groups with the most to gain from improved credit access. As a result, the most skilled households lose in welfare terms from any qualifying loan sizes in excess of approximately \$4,000 ($\vartheta = 0.1$). Conversely, HS households continue to gain, and gain substantially in welfare terms, from loan guarantees of up to

Table 4 Optimal Generosity of Loan Guarantee Program

	$\theta = 0.50$		
	COLL	HS	NHS
$\vartheta = 0.00 \rightarrow \vartheta = 0.10$	0.02%	0.08%	0.13%
$\vartheta = 0.00 \rightarrow \vartheta = 0.20$	−0.24%	0.20%	0.22%
$\vartheta = 0.00 \rightarrow \vartheta = 0.30$	−1.41%	0.27%	0.39%
$\vartheta = 0.00 \rightarrow \vartheta = 0.40$	−1.60%	0.19%	0.78%
$\vartheta = 0.00 \rightarrow \vartheta = 0.50$	−2.24%	−0.11%	1.06%
$\vartheta = 0.00 \rightarrow \vartheta = 0.60$	−2.84%	−0.35%	1.26%
$\vartheta = 0.00 \rightarrow \vartheta = 0.70$	−3.60%	−0.44%	1.02%

$16,000 ($\vartheta = 0.4$). Most strikingly, NHS households gain for very large loan guarantee levels, even to levels exceeding their mean income level. In summary, our results suggest that modest loan guarantee programs can improve welfare for all households, even those households who likely will pay the bulk of the taxes needed to finance them. However, our model also suggests that qualifying loan size is likely to be quite important in determining whether a particular guarantee program serves all households or instead functions as a very significant redistributive mechanism. In the absence of definitive means for detecting the sensitivity of aggregate credit use and default to the size of qualifying loans, instituting a program that is too generous will lead to significant welfare losses for some groups.

Where do the welfare gains come from? Table 5 shows mean consumption and decomposes the variance of consumption into two moments: the variance of mean consumption by age, a measure of *intertemporal* consumption smoothing, and the mean of consumption variance by age, a measure of *intratemporal* consumption smoothing.[21] What we mean here is the following. "Intertemporal" smoothing refers to how much variation of consumption or living standards individuals experience. We measure it by calculating how much average consumption varies over the life cycle, and we average consumption as a natural measure of what the individual can expect at any given age. This is an intuitive measure of consumption smoothing through time: If the variance of average consumption over the life cycle were high, this would mean that young and old households were, on average, consuming quite different amounts. As for "intratemporal" smoothing, our measure answers the question of how much variability there is among households of any given age, when averaged across individuals of all ages. In other

[21] Specifically, we use the decomposition: $\text{var}(\log(c)) = \text{var}(\text{E}[\log(c)\,|j]) + \text{E}[\text{var}(\log(c)\,|j)]$.

Table 5 Distribution of Consumption

	$\mathrm{E}(c)$	$\mathrm{var}(\log(c))$	$\mathrm{E}(\mathrm{var}(\log(c)\,\vert age))$	$\mathrm{var}(\mathrm{E}(\log(c)\,\vert age))$
		Aggregate		
NO LG	0.8455	0.1894	0.1671	0.0223
LG $\vartheta = 0.5, \theta = 0.5$	0.8016	0.1977	0.1755	0.0222
		College		
NO LG	1.0918	0.1776	0.1293	0.0481
LG $\vartheta = 0.5, \theta = 0.5$	1.0521	0.3874	0.3354	0.0520
		High School		
NO LG	0.7767	0.2279	0.1907	0.0372
LG $\vartheta = 0.5, \theta = 0.5$	0.7575	0.3926	0.3749	0.0180
		Non-High School		
NO LG	0.6579	0.2807	0.2582	0.0225
LG $\vartheta = 0.5, \theta = 0.5$	0.6514	0.3932	0.3849	0.0083

words, what is the average variability of consumption that one would expect to observe if one drew a sample of households of any given age? The decomposition of total variance into these components is a complete one: Together they account for the total variance of consumption in the model (and this is due to a simple statistical fact known as the "law of total variance").

Loan guarantees reduce average consumption due to the combination of higher taxes, more borrowing, and more frequent default. The gain comes through a better distribution of consumption over the life cycle. We see here that this gain is driven entirely by a reduction in the intertemporal dimension as intratemporal consumption volatility actually *increases*.

We note here that our welfare results differ significantly from those in Athreya et al. (2012), where the role of the out-of-pocket costs of default, Λ, in restricting access to bankruptcy is explored. High values of Λ restrict access to bankruptcy to high income types (who typically do not want to default), and in a wide range of models the optimal value (from an *ex ante* perspective) is infinite for all types; that is, from the perspective of a newborn household, permitting any bankruptcy in equilibrium is suboptimal. The largest gains are experienced by the college types, because they have the strongest demand to borrow for purely intertemporal reasons (i.e., reasons unrelated to the effects of uncertainty) and this demand is thwarted by risk-based pricing. There are a number of reasons to view that result as impractical from a policy

Table 6 Welfare Decomposition, Symmetric Information

	$\vartheta = 0.5, \theta = 0.50$		
	COLL	HS	NHS
$(q^{NLG}, \tau_2 = 0.0) \rightarrow (q^{LG}, \tau_2 = 0.0)$	4.86%	8.30%	10.69%
$(q^{LG}, \tau_2 = 0.0) \rightarrow (q^{LG}, \tau_2 = 0.0386)$	−6.74%	−7.76%	−8.69%

perspective. Loan guarantees, in contrast, are clearly policy-feasible and benefit the least-skilled more than the more-skilled.[22]

Decomposing the Effect of Taxes on Welfare

In this subsection we decompose the net effect of the loan guarantee program. We consider two experiments, presented in Table 6, where we ask how welfare changes if we confront an individual with the pricing emerging from the presence of a loan guarantee, with and without the taxes needed to finance the program. Starting in the top row of Table 6 we display the effect of a move from the benchmark setting to one in which a tax-free loan guarantee is provided. Welfare increases quite substantially, again by least for the skilled and by most for the unskilled. Since their income profile is flat, the NHS households experience the largest gain because they use unsecured debt over most of their life cycle. By contrast, the more-skilled types decrease unsecured borrowing as they age (see Figure 5).

Turning next to the bottom row of Table 6, we present the welfare implications of a move from a setting with a tax-free loan guarantee to one where, including taxes, the program must now break even. As seen earlier (Table 5), once taxes are imposed only the unskilled benefit from a program this generous, and they lose proportionally more from taxes than do the college types. Why are the costs of a small tax so large in this model? With taxes, permanent income is reduced, leaving households more exposed to the expenditure shock. As a result, they "involuntarily" default more frequently, leading to more deadweight loss and a much larger welfare loss than one would expect from a tax of less than 4 percent. Due to the accumulation pattern of net worth,

[22] Dávila et al. (2012) shows that utilitarian constrained efficient allocations in a model with uninsurable idiosyncratic shocks are skewed toward improving the welfare of "consumption-poor" households (since they have higher marginal utility). While we do not attempt to characterize constrained efficient allocations here, it seems clear that this intuition would apply—thus, policies that raise the utility of the least-skilled would seem to be preferable from a social welfare perspective.

Figure 5 Net Worth Over the Life Cycle

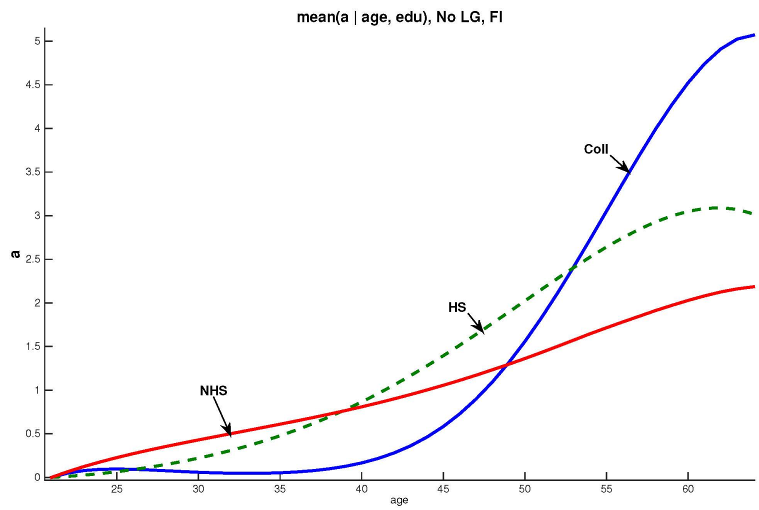

on average NHS households are more exposed to this risk (again, see Figure 5).

Table 7 decomposes the costs of the program by type. The loan guarantee program transfers resources along two dimensions. First, loan guarantees transfer resources from skilled households to less-skilled; college types pay into the program, via taxes, significantly more than they collect in terms of lower interest rates. Second, loan guarantees transfer resources from individuals who pose little default risk (those with *low* λ) to those with a high value for λ, as the latter pose more default risk, all else equal. This transfer occurs because the high-risk types would pay substantially higher interest rates without intervention and therefore gain a lot from the program.

Asymmetric Information

Returning to the problem noted at the outset of the previous subsection, recall that the cost of limited access to unsecured credit is likely largest for the least wealthy. This is particularly likely to be true in a society that lacks the information storage, sharing, and data analysis available in developed nations to effectively identify credit risk at the time of loan origination (and then update it regularly). As a first step

Table 7 Distribution of Net Costs Paid by Type

	$\vartheta = 0.50, \theta = 0.50$, **FI**			
	High λ		Low λ	
	Taxes	Transfer	Taxes	Transfer
Coll	0.1366	0.1050	0.1366	0.0384
HS	0.2995	0.5082	0.2995	0.1512
NHS	0.0639	0.1333	0.0639	0.0639
	$\vartheta = 0.50, \theta = 0.50$, **PI**			
	High λ		Low λ	
	Taxes	Transfer	Taxes	Transfer
Coll	0.1366	0.1155	0.1366	0.0341
HS	0.2995	0.4971	0.2995	0.1239
NHS	0.0639	0.1711	0.0639	0.0583

in getting a sense of the quantitative potential of loan guarantees to alter outcomes in such settings, we now study stationary equilibria of our model under asymmetric information.

To remind the reader, in our economy, asymmetric information will mean that the borrower will have characteristics that are not observable to the lender; specifically, we assume neither current stigma, λ, nor current net worth, a, can be directly observed. However, any information about these variables that can be *inferred* from the observable components of the state vector, as well as from the desired borrowing level, b, is available to the lender.[23] We focus on two representative examples: one that represents a relatively modest loan guarantee program and results in welfare gains for all types under symmetric information ($\theta = 0.1$ and $\vartheta = 0.1$), and one that is more generous and reduces the welfare of college-educated types ($\theta = 0.5$ and $\vartheta = 0.4$). Our key finding is that the presence of asymmetric information will increase the gains available from loan guarantees, no matter how generous.

Allocations and Pricing

We first compare outcomes in the FI and PI economies. Table 8 shows that a move from symmetric to asymmetric information has the following effects. First, default falls for all types, and default skews more strongly toward the high λ type; these individuals are treated relatively better under asymmetric information, since they get terms that reflect the average default risk instead of their own, and therefore end up

[23] We assume that credit markets are anonymous, so that past borrowing is also not observable to the current lender. In Athreya, Tam, and Young (2012b) we introduce a flag that tracks whether a household is likely to have recently defaulted. Due to computational considerations we do not examine this case here.

Table 8 Aggregate Effects of Loan Guarantees—Asymmetric Information

	$\vartheta = 0.40$			
	FI		PI	
$\theta =$	0.0000	0.5000	0.0000	0.5000
τ_2	0.0000	0.0245	0.0000	0.0196
Discharge/Income Ratio	0.2662	0.6965	0.2021	0.6497
Fraction of Borrowers	0.1720	0.3109	0.1614	0.3036
Debt/Income Ratio | NHS	0.1432	0.4880	0.1209	0.4762
Debt/Income Ratio | HS	0.1229	0.3897	0.0909	0.3755
Debt/Income Ratio | COLL	0.0966	0.2691	0.0801	0.2389
Default Rate | NHS	1.237%	13.170%	0.956%	12.704%
Default Rate | HS	1.301%	12.310%	0.957%	11.407%
Default Rate | COLL	0.769%	6.304%	0.658%	5.412%

borrowing amounts that induce relatively high default rates. Second, overall the credit market shrinks, in the sense that we observe fewer borrowers (of each type) and lower discharged debt aggregates.

Figure 6 shows that pricing is significantly worse for the high λ (*low* bankruptcy cost) borrower and *better* for the low λ borrower. Under asymmetric information, the two types will be pooled together, so that the default premium at a given debt level reflects the average default risk. The result is that good borrowers face significantly tighter credit limits and higher interest rates, while bad borrowers face the same credit limit but lower interest rates. The shift in pricing accounts for the smaller credit market size.

Third, as noted at the outset, our model features expenditure shocks. These shocks take on a larger role in defaults under asymmetric information (see Table 9). With tighter credit limits, big expenditure shocks that hit when the household is young are hard to smooth, since income is relatively low. The result is that essentially all defaults are done by households who have received an expenditure shock, despite this group being only 7.56 percent of the population. Information has less of an impact on these defaults, since they are defaults on debt that has been acquired involuntarily.

We now turn to the effects of loan guarantees under asymmetric information. Table 8 shows that the change induced by the introduction of the particular program is larger for all credit market aggregates under asymmetric information, with the exception of the debt-to-income ratio for college-educated households (in which case it is of only slightly smaller magnitude). Figure 7 shows the increased access to credit that guarantees provide in these two cases. Note that the increase in the default rate is smaller under asymmetric information for every

Figure 6 Pricing with Symmetric and Asymmetric Information

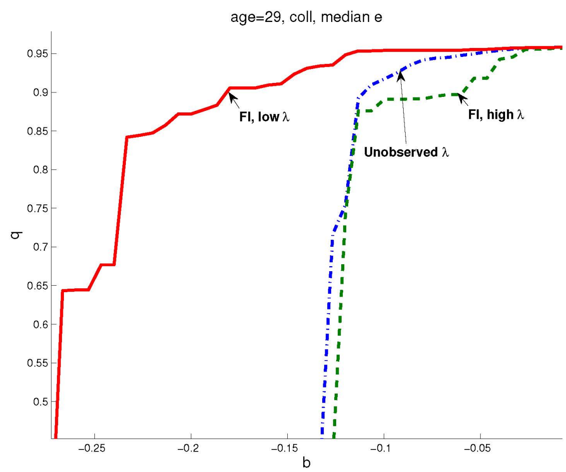

education group. As a result, the taxes required to finance the program are lower than under symmetric information.

Welfare

Table 10 displays the welfare effects of two different loan guarantee programs. Relative to the symmetric information case, loan guarantees are uniformly better when information is asymmetric; this result holds for every case we have computed. The larger gain is partly due to the lower tax burden required in the asymmetric information cases and partly due to the severe pricing distortion caused by asymmetric information evident in Figure 6.

To more directly describe the transfers between agents induced by loan guarantees, Table 7 collects the proportion of costs paid by each group. Now the loan guarantee program subsidizes the high λ (low stigma cost) types much more than under symmetric information. This result is exactly what we would expect, given that this type is receiving better credit terms under asymmetric information.

Targeted Loan Guarantees

Our results suggest that loan guarantees have the potential to become primarily a means of transferring resources from the rich to the poor.

Table 9 Distribution of Default by State

	FI		PI	
	High λ	Low λ	High λ	Low λ
Low x	0.2315	0.0092	0.0089	0.0000
Median x	0.5811	0.0670	0.8033	0.0399
High x	0.0666	0.0446	0.0890	0.0588

Moreover, our findings suggest that they may also lower welfare, often of all types of agents, unless their generosity is modest. In our results, default is disproportionately driven by those who have received an expenditure shock. A natural question therefore is whether the benefits of loan guarantees discussed at the outset can be preserved by limiting compensation to lenders only when a borrower has suffered such a shock. Expenditure shocks represent large increases in debts that are rare and involuntarily acquired. As a result, a policy of guaranteeing loans only under these conditions is unlikely to alter loan pricing substantially (since these states are rare) but may substantially aid households who find themselves in those rare states. Moreover, targeted guarantees are unlikely to induce significant additional deadweight loss because the default decision is more frequently heavily influenced by expenditure shocks, which again, are rare.

To investigate this question, we study a case where $\vartheta = 0.50$ and $\theta = 0.50$, but where lenders only receive compensation in the event that a bankruptcy coincides with a positive expenditure shock ($x > 0$). Table 11 shows that all groups gain from the introduction of a loan guarantee program restricted in this manner. As before, the NHS households gain most and the highly skilled gain the least. Nonetheless, the ability of the conditionality of the program to overturn what was initially a very large welfare loss to the skilled into a gain is striking.[24]

To see the effect on aggregates more generally, we turn to Table 12. It is immediately clear that the tax rate needed to sustain the restricted loan guarantee program is very small relative to the unrestricted case, even though the debt discharged in bankruptcy is similar to the unrestricted guarantee case. Nonetheless, the overall level of debt responds to the restricted guarantee far more modestly than the unrestricted case. For example, under restricted guarantees, the mean

[24] We are implicitly assuming that expenditure shocks are likely to be easy to observe; we doubt that agents could easily hide one from the government, given the size and nature of these shocks. Our calibration, as noted above, equates x to a combination of medical and legal bills plus unplanned family costs; these expenses should be relatively easy to monitor in practice.

Figure 7 Pricing with Loan Guarantee, Symmetric versus Asymmetric Information

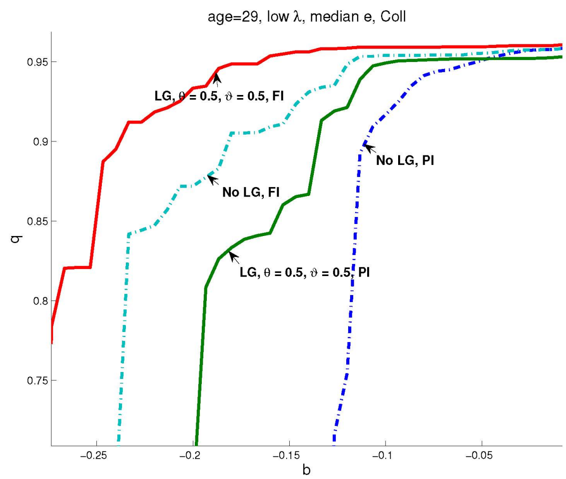

debt-to-income ratio among high-school educated borrowers is less than half that under unrestricted guarantees (0.2256 versus 0.4707). The central reason for the low tax rate is that the default rate responds by far less than with an unrestricted program, even though borrowing does increase nontrivially, relative to the benchmark case. Under restricted guarantees, the bankruptcy rate roughly doubles, while the unrestricted program implies a nearly ten-fold increase.

3. DISCUSSION

We have made a few assumptions in our model that require some additional discussion. First, we have assumed that factor prices are fixed. General equilibrium calculations would imply higher r and lower W would prevail under loan guarantee systems, since they produce more borrowing and less aggregate wealth (as well as increasing the amount of transactions costs that works like a reduction in aggregate supply of goods). Factor price movements of this sort are likely to make the welfare costs larger (gains smaller), since the higher risk-free interest rate would make borrowing more costly and the lower wages would reduce mean consumption. Despite these effects, we choose to abstract from equilibrium pricing because it is well known that income processes representative of the vast majority of households will, in environments

Table 10 Welfare Effects of Loan Guarantees

	COLL	HS	NHS
NO LG$\rightarrow \theta = 0.50, \vartheta = 0.40$,FI	−1.60%	0.19%	0.78%
NO LG$\rightarrow \theta = 0.50, \vartheta = 0.40$,PI	−1.02%	0.98%	1.59%
NO LG$\rightarrow \theta = 0.10, \vartheta = 0.10$,FI	0.01%	0.02%	0.03%
NO LG$\rightarrow \theta = 0.10, \vartheta = 0.10$,PI	0.04%	0.08%	0.11%

such as ours, produce less wealth concentration than observed (see Castaneda, Díaz-Giménez, and Ríos-Rull 2003), meaning that the mean wealth position will be too similar to the median, implying larger factor price changes than would occur if the distribution of wealth were matched. Given the immense computational burden that matching the U.S. Gini coefficient of wealth would impose on our OLG setup, and given that the factor price adjustments should be small, we feel justified in ignoring them.[25]

Second, we have financed the program using proportional labor income taxes. An obvious alternative would be to finance the program using progressive income taxes, where high income (college) types would pay higher marginal tax rates. This approach would increase the gains to the NHS types, who already gain substantially, and reduce (or even eliminate) any gains to college types. We expect a similar result from capital income taxation as well, since it will tend to tax the wealthier college types more heavily. In contrast, a regressive income tax would imply the types who benefit the most, the NHS, would pay a higher marginal tax rate. Regressive tax systems seem unlikely to be implemented on equity grounds, even if they are welfare-improving within a specific model. We could also introduce separate programs for each education group, so that the cross-subsidization that makes the program so attractive to NHS types would be eliminated; we conjecture that this case would result in larger gains for college types and smaller for NHS types.

Third, there is a conceptual issue of the right benchmark allocation. The U.S. corporate income tax rate is 35 percent and banks are permitted to deduct losses due to nonperforming loans from their taxable income. As a result, it may be that the appropriate benchmark is a case where the loan guarantee program is not zero, but rather has a large value of ϑ and $\theta = 0.35$. We can of course easily express the welfare

[25] In Chatterjee et al. (2007), the model is calibrated to the U.S. distribution of wealth; the resulting effects of an endogenous risk-free rate are quantitatively unimportant.

Table 11 Welfare Effects of Restricted Loan Guarantees

	$\vartheta = 0.5, \theta = 0.50$		
	COLL	HS	NHS
NO LG→Restricted LG	0.40%	0.77%	0.99%
Restricted LG→Unrestricted LG	−2.66%	-0.88%	−0.07%

gains relative to this benchmark instead; a more detailed investigation of this issue is part of ongoing work.

There are some natural extensions of our model that seem useful to pursue. Given our results regarding the effect of loan guarantees to redistribute toward the unskilled from the skilled, it would be productive to know if the least skilled, for example, would benefit from a loan guarantee program that was required for self-financing via taxes on only the unskilled. Such an extension would be along the lines explored in Gale (1991), who studies *targeted* loan guarantees designed to facilitate credit access for certain identifiable subpopulations (such as minority borrowers). Targeted programs would be related to the regulations we mentioned earlier that require certain characteristics not be reflected in credit terms; exactly how the dual goals of encouraging access to these groups without allowing their characteristics to alter credit terms would affect welfare is unknown and worth studying. It would also be straightforward to investigate loans targeted to individual borrowers who are deemed constrained by competitive lenders.[26] In our model, since borrowers are at a "cliff" in the pricing function, they would benefit from government loans *at their existing interest rate*, provided the tax costs are not "too high."

Also, our work is a step in the direction that, in the future, will allow us to analyze the role of guarantees for mortgage lending. However, the central role of aggregate risk in driving home-loan default makes a full quantitative analysis that satisfactorily incorporates the forces we *do* allow for here—asymmetric information and limited commitment—currently infeasible. But we note that such a model would have the same fundamental structure as that developed here.

4. CONCLUDING REMARKS

A significant share of the U.S. population appears credit constrained. These households usually lack collateral and must therefore rely on the unsecured credit market to help them smooth consumption in the

[26] A stylized approach to this is taken in Smith and Stutzer (1989).

Table 12 Aggregate Effects of Restricted Loan Guarantees

	$\vartheta = 0.50$		
$\theta =$	0.00	0.50	0.50
	No LG	Restricted LG	Unrestricted LG
τ_{LG}	0.0000	0.0004	0.0386
Discharge/Income Ratio	0.2662	0.7208	0.8657
Fraction of Borrowers	0.1720	0.2408	0.3527
Debt/Income Ratio | NHS	0.1432	0.2649	0.5738
Debt/Income Ratio | HS	0.1229	0.2256	0.4707
Debt/Income Ratio | COLL	0.0966	0.1681	0.3285
Default Rate | NHS	1.237%	2.755%	16.797%
Default Rate | HS	1.301%	2.586%	14.619%
Default Rate | COLL	0.769%	1.643%	9.072%

face of life-cycle and shock-related movements in income. However, the unsecured credit market in the United States appears significantly impeded by forces that keep the costs of unsecured debt default low, and thereby make lending risky and, hence, expensive. Perhaps the most widely used route to increase credit flows to target groups is via the use of loan guarantees whereby public funds defray private lenders' losses from default. Aside from their direct effects on credit access and pricing, guarantees are likely to be particularly useful in unsecured credit markets given limitations on the ability of policies to directly influence borrowers' default incentives. In this article, we assess the consequences of extending loan guarantees to unsecured consumer lending to improve allocations.

Our article attempts to quantify the impact of loan guarantees in a model that incorporates both meaningful private information and a limited commitment problem into a rich life-cycle model of consumption and savings. Our quantitative analysis focuses on evaluating the impact of introducing loan guarantees into unsecured consumer credit markets. These markets have large consequences for household welfare because they influence the limits on smoothing faced by some of the least-equipped subgroups in society, particularly the young and the unlucky.

Our calculations suggest first that, under symmetric information, loan guarantees can actually improve the *ex ante* welfare of *all* households if they are not too generous (meaning only small loans qualify). This welfare gain is disproportionately experienced by low-skilled households who face flat average income paths and relatively large shocks. Indeed, such households gain from very generous programs, but higher-skilled types rapidly begin to experience welfare losses as loan guarantees are made more generous. These results arise because loan guarantees induce a transfer from skilled to unskilled, and this transfer can be substantial, while the gains to the skilled from seeing loan pricing terms improve as a result of guarantees is relatively small. Second, we find that allocations are quite sensitive to the size of qualifying loans: Even modest limits on qualifying loan size invite very large borrowing—as perhaps intended by proponents—but also spur very large increases in default rates. As a result, loan guarantee programs transfer resources in significant amounts from all households to the lifetime poor. Under asymmetric information, the welfare gains are larger for all households, as the taxes required to finance the programs are smaller. Our work provides an answer for why, despite the potential welfare gains from expanding guarantees to consumer credit that thereby alleviate credit constraints for a marginalized population otherwise lacking collateral, public guarantees on unsecured consumer credit have not yet been implemented. The value of the program depends on how elastically credit demand and supply respond to default risk, which may be hard to estimate, and the programs are quite costly if too generous.[27] As a practical matter, the forces at work in our model may well be part of explaining why student loan default rates hit 25 percent in the early 1990s, at which point the government increased monitoring and enforcement (recall also the similar findings of Lelarge, Sraer, and Thesmar [2010] in the French entrepreneurship context).

The preceding intuition will likely carry over to markets beyond the one for unsecured consumer credit, in particular for two areas that have seen some form of loan guarantee: federal student loans and home loans. It suggests that loans of the size guaranteed by a federal student loan program would have been likely to default at high rates, even under a relatively "partial" nature of the guarantee. Similarly, the FHA/VA and others have historically provided loan guarantees for mortgage loans. The calibrated costs of default measured in our model suggest strongly that larger loans, especially if covered more fully by a loan guarantee program, would lead to even greater debt and default than that predicted for the consumer credit market. Therefore, unless such loans are vetted carefully, one should expect a high take-up rate, a high subsequent failure rate, and nontrivial transfers from better-off households. Nonetheless, despite the risks involved, a main result of the article is that a limited program, specifically one where loan guarantees are made contingent on certain rare but disastrous events, *can*

[27] An important caveat here is that in our model, the costs of default are assumed invariant to the level of default in the economy. Thus, a major loan guarantee program may meaningfully affect default costs. This is surely subject to at least some Lucas Critique-related problems. Nonetheless, endogenizing these costs is beyond the scope of the article.

deliver net gains for all households. Such a policy seems worth exploring further. Of course, a caveat to the conclusion that targeting guarantees to those who have suffered a bad expense shock is that it may require additional resources to battle any moral hazard that might be present, especially when default is allowed upon getting any shock that is not a genuine catastrophe to households. Taken as a whole, our results suggest that loan guarantees can help, but care must be taken if policymakers intervene in credit markets through the use of loan guarantees.

Lastly, because the results reported in this article suggest that loan guarantees for household credit may be a powerful tool for altering steady-state consumption, our work should be of help for future examinations of the extent to which consumer lending and more importantly, consumer *willingness* to borrow, can be amplified to spur current consumption in business cycle contexts. The model of Gordon (2015) could possibly be adapted to this question.

APPENDIX: QUANTITATIVE MODEL

We now provide a detailed description of the quantitative model used here. As noted at the outset, it is essentially that of Athreya, Tam, and Young (2012b), modified to accommodate changes in loan pricing and taxes necessitated by loan guarantees.

Preferences

Households in the model economy live for a maximum of $J < \infty$ periods and face stochastic labor productivity and mortality risk. Households supply labor inelastically.[28] Households differ along several dimensions over their life cycles according to an index of type, denoted y and defined in what follows. Each household of age j and type y has a conditional probability $\psi_{j,y}$ of surviving to age $j+1$. Households retire exogenously at age $j^* < J$. Let n_j denote the number of "effective" members in a household. Households value consumption per effective

[28] We abstract from elastic labor supply because it is known (e.g., Pijoan-Mas 2006) that under incomplete markets, households borrowing significant amounts tend to supply labor relatively inelastically, and for our study, this margin is unlikely to be crucial. It naturally implies that our welfare cost measurements may be biased, but it is unclear which direction that bias would go.

household member $\frac{c_j}{n_j}$. They have identical additively separable isoelastic felicity functions with parameter σ, and possess a common discount factor β. To smooth consumption, all households have access to risk-free savings, and also debt that they may fully default on, subject to some costs. These costs reflect the variety of consequences that bankruptcy imposes on households, and need not be interpreted solely as "stigma," but include any such costs. A portion of these costs are represented by a nonpecuniary cost of filing for bankruptcy, denoted by $\lambda_{j,y}$, which we also permit to depend on household type y. Household preferences are therefore given by

$$U\left(\left\{\frac{c_j}{n_j}\right\}_{j=1}^{J}\right) = \frac{1}{1-\sigma} E\left[\sum_{j=1}^{J} \beta^{j-1} \psi_{j,y} \left(\left(\left[\lambda_{j,y} d_j + (1-d_j)\right] \frac{c_j}{n_j}\right)^{1-\sigma}\right)\right], \tag{3}$$

where d_j is the indicator function that equals unity when the household chooses to default in the current period (in which case $d_j = 1$).

The existence of nonpecuniary costs of bankruptcy is strongly suggested by the calculations and evidence in Fay, Hurst, and White (1998) and Gross and Souleles (2002). The first article shows that a large measure of households would have "financially benefited" from filing for bankruptcy but did not, while both articles document significant unexplained variability in the probability of default across households after controlling for a large number of observables.

In this specification, a household with a relatively low value of $\lambda_{j,y}$ will obtain low value from any given *expenditure* on consumption (c_j) in a period in which they file for bankruptcy. This is meant to reflect the increased transactions cost associated with obtaining utility via consumption expenditures in the period of a bankruptcy. Examples include increased "shopping time" arising from difficulty in obtaining short-term credit and payments services, locating rental housing and car services, as well as any stigma/psychological consequences. For convenience, we will sometimes refer to $\lambda_{j,y}$ as stigma in what follows; we intend it to be more encompassing.[29] Because of the breadth of costs that λ represents, we will allow it to vary stochastically over time and across individuals as a function of their type y, according to a transition function p_λ.

At the time of obtaining a loan, a household who expects to have a relatively low value of λ next period will know that filing for

[29] Another possibility is that these households gain the benefits from bankruptcy without filing, as suggested by Dawsey and Ausubel (2004). Athreya et al. (2012) extends the benchmark model to include a delinquency state in which households do not formally file for bankruptcy but also do not service their debt.

bankruptcy will result in a relatively high cost of obtaining any given level of marginal utility in the next period. Given the current marginal utility of consumption, consumption smoothing (i.e., keeping marginal utility in accordance with the standard Euler equation) under bankruptcy will therefore be costlier, all else equal, than for a household with a high value of λ. This is further amplified by the fact that households are not allowed to borrow in the same period as when they file for bankruptcy. For convenience, we will therefore refer to those whose value of $\lambda_{j,y}$ is relatively low as "low-risk" borrowers, and vice versa.

In addition to this nonpecuniary cost, there is an out-of-pocket pecuniary resource cost Λ that represents all formal legal costs and other procedural costs of bankruptcy. Lastly, households are not allowed to borrow or save in the same period as a bankruptcy filing, to capture provisions guarding against fraud that are routinely applied in court. There are no other costs of bankruptcy in the model.

Endowments

Our focus on consumer credit makes it critical to allow for both uninsurable idiosyncratic risk. Consumer default, and hence the value of loan guarantees, is by all accounts strongly tied to *individual*-level uninsurable risk (see, e.g., Sullivan, Warren, and Westbrook [1999, 2000] and Chatterjee et al. [2007]).[30] There are two sources of such risk in our model. First, households face shocks to their labor productivity, and because they are modeled as supplying labor inelastically, face shocks to their labor earnings. Second, households are susceptible to shocks to their net worth. The former represent shocks arising in the labor market more generally, and the latter represent sudden required expenditures arising from unplanned events such as sickness, divorce, and legal expenses.

In addition to the use of credit to deal with stochastic fluctuations in income and expenditures, consumer credit also likely serves, as noted earlier, as a tool for longer-term, more purely intertemporal smoothing in response to predictable, low-frequency changes in labor income, such as those coming with increased age and labor market experience. This leads us to specify, in addition to transitory and persistent shocks to income, a deterministic evolution in average labor productivity over the life cycle. This component of earnings will reflect most obviously

[30] In mortgage lending, loan guarantees protect lenders against house price fluctuations, which in turn are strongly tied to *aggregate risk* (or at least city-level risk). The full incorporation of the aggregate risk, private information, and limited commitment needed to analyze this specific class of guarantees remains an important topic for future work.

one's final level of educational attainment, which is represented in the model as part of an agent's type, y.

Specifically, log labor income will be determined as the sum of four terms: the aggregate wage index W, a permanent shock y realized prior to entry into the labor market, a deterministic age term $\omega_{j,y}$, and a persistent shock e that evolves as an AR(1) process. The log of income at age-j for type$-y$ is therefore given by

$$\log W + \log \omega_{j,y} + \log y + \log e + \log \nu,$$

where

$$\log\left(e'\right) = \varsigma \log\left(e\right) + \epsilon', \quad (4)$$

and a purely transitory shock $\log(\nu)$. Both ϵ and $\log(\nu)$ are independent mean zero normal random variables with variances that are y-dependent and have distributions p_e and p_v, respectively.

As for the risk of stochastic expenditures, we follow the literature (e.g., Chatterjee et al. 2007 and Livshits, MacGee, and Tertilt 2007), and specify a process x_j to denote the expense shock to net worth that takes on three possible values $\{0, x_1, x_2\}$ from a probability distribution $p_x(\cdot)$ with i.i.d. probabilities $\{1 - p_{x1} - p_{x2}, p_{x1}, p_{x2}\}$.

We will take agents' permanent type y to reflect differences between households with permanent differences in human capital. Specifically, we will consider agents with three types of human capital: those who did not graduate high school, those who graduated high school, and those who graduated college.[31] This partition of households follows Hubbard, Skinner, and Zeldes (1994). The central reason for allowing this heterogeneity is that the observed differences in mean life-cycle productivity for each of these types of agents gives them different incentives to borrow over the life cycle. In particular, college workers will have higher survival rates and a steeper hump in earnings; the second is critically important as it generates a strong desire to borrow early in the life cycle. They also face smaller shocks than the other two education groups. The life-cycle aspect of our model is key; in the data, while bankruptcies occur late into the life cycle for some (see, e.g., Livshits, MacGee, and Tertilt 2007), defaults are still skewed toward young households.[32]

[31] Mortality rates also differ by education, although this heterogeneity is of no consequence for our questions.

[32] See Sullivan, Warren, and Westbrook (2000).

Market Arrangement

As stated earlier, to smooth consumption and save for retirement, households have access to both risk-free savings as well as one-period defaultable debt. The issuance and pricing of debt is modeled as a two-stage game in which households at any age j first announce their desired asset position b_j, after which a continuum of lenders simultaneously announces a loan price q. As a result, a household issuing b_j units of face value receives qb_j units of the consumption good today. A household who issues debt with face value b_j at age-j is agreeing to pay b_j in the event that they fully repay the loan, and pay zero otherwise (i.e., when they file for bankruptcy). The fact that nonrepayment can occur with positive probability in equilibrium means that lenders will not be willing to pay the full face value, even after adjusting for one-period discounting. Therefore, given any gross cost of funds $\widehat{R}$, we must have $q \leq 1/\widehat{R}$.

As we will allow for both symmetric and asymmetric information, we introduce the following notation. Let I denote the information set for a lender and $\widehat{\pi} : b \times I \to [0, 1]$ denote the function that assigns a probability of default to a loan of size b_j given information I. Clearly, since default risk assessed by lenders will depend in general on both their information and the size of the loan taken by a household, so will loan prices. Therefore, let loan pricing be given by the function $q(b_j, I)$. Under asymmetric information, we allow lenders to use the information revealed by the size of the loan request and lenders' knowledge of the distribution of household net worth in the economy to update their assessment of all current unobservables. Thus, lenders use their knowledge of both (i) optimal household decision making (i.e., their decision rules as a function of their state), and (ii) the endogenous distribution of households over the state vector. We will describe the determination of this function in detail below.

The household budget constraint during working life, as viewed immediately after the decision to repay or default on debt has been made, is given by

$$c_j + q\left(b_j, I\right) b_j + \Lambda d_j \leq a_j + \left(1 - \tau_1 - \tau_2\right) W \omega_{j,y} y e \nu. \tag{5}$$

a_j is net worth *after* the current-period default decision d_j. Therefore, $a_j = b_{j-1} - x_j$ if $d_j = 0$ and 0 if $d_j = 1$. Households' default decisions also determine their available resources beyond removing debt, because default consumes real resources Λ, arising from court costs and legal fees. The last term, $\left(1 - \tau_1 - \tau_2\right) W \omega_{j,y} y e \nu$, is the after-tax level of current labor income, where τ_1 is the flat-tax rate used to fund pensions and τ_2 is the rate used to finance the loan guarantee program. Keep in mind also that implicit in the specification of the loan pricing function

$q(\cdot)$ is the fact that if the household borrows an amount in excess of the guarantee limit, the price is that of an entirely nonguaranteed loan.

The budget constraint during retirement is

$$c_j + q\left(b_j, I\right) b_j + \Lambda d_j \leq a_j + \upsilon W \omega_{j^*-1,y} y e_{j^*-1} \nu_{j^*-1} + \Upsilon W, \qquad (6)$$

where for simplicity we assume that pension benefits are composed of a fraction $\upsilon \in (0,1)$ of income in the last period of working life plus a fraction $\Upsilon \in (0,1)$ of average income W (we normalize average individual labor earnings to 1).

Consumer's Problem

The timing is as follows. In each period, all uncertainty is first realized. Thus, income shocks e and υ, the default cost λ, and the current expense shock x are all known before any decisions within the period are made. Following this, households must decide, if they have debt that is due in the current period, to repay or default. This decision, along with the realized shocks, then determines the resources the household has available in the current period. Given this, the household chooses current consumption and debt or asset holding with which to enter the next period, and the period ends.

Prior to making the current-period bankruptcy decision, a household can be fully described by b_{j-1}, the debt, if any, that is due in the current period, their type y, the pair of currently realized income shocks e and υ, their cost of default λ, the current realization of the shock to expenses, x_j, and their age j.[33]

Letting $V(\cdot)$ denote the household's value function prior to the decision to default or repay, with primed variables denoting objects one period ahead, we have the following recursive description. If the household chooses to repay its debt b_{j-1}, and therefore sets $d_j = 0$, then the value they derive from state $(b_{j-1}, y, e, \nu, \lambda, x, j)$ is

$$v^{d=0}\left(b_{j-1}, y, e, \nu, \lambda, x, j\right) =$$
$$\max_{c, b_j} \left\{ \begin{array}{c} \left(\frac{c_j}{n_j}\right)^{1-\sigma} + \\ \beta \psi_{j,y} \left(\sum_{e',\nu',\lambda',x'} p_x(x') p_e\left(e'|e\right) p_\nu\left(\nu'\right) p_\lambda\left(\lambda'|\lambda\right) V\left(b, y, e', \nu', \lambda', x', j+1\right)\right) \end{array} \right\} \qquad (7)$$

subject to the budget constraint

[33] To avoid repetition, we display only the value functions during working life; retirement is entirely analogous.

$$c_j + q\left(b_j, I\right) b_j + x_j \leq b_{j-1} + \left(1 - \tau_1 - \tau_2\right) W \omega_{j,y} y e \nu. \qquad (8)$$

If the household *has* chosen bankruptcy for the current period ($d_j = 1$), since we disallow credit market activity in the period of bankruptcy, which implies $b_j = 0$, we obtain

$$v^{d=1}\left(b_{j-1}, y, e, \nu, \lambda, x, j\right) = \left\{ \begin{array}{c} \left(\lambda_{j,y} \frac{c_j}{n_j}\right)^{1-\sigma} + \\ \beta \psi_{j,y} \left(\sum_{e',\nu',\lambda',x'} p_x(x') p_e\left(e'|e\right) p_\nu\left(\nu'\right) p_\lambda\left(\lambda'|\lambda\right) V\left(0, y, e', \nu', \lambda', x', j+1\right)\right) \end{array} \right\},$$

subject to the budget constraint:

$$c_j + \Lambda \leq \left(1 - \tau_1 - \tau_2\right) W \omega_{j,y} y e \nu. \qquad (9)$$

Notice that both debt due in the current period, b_{j-1}, and the current expenditure shock realization, x_j, get removed by bankruptcy, and hence disappear, when comparing the budget constraint under bankruptcy to one under nonbankruptcy. By contrast, the resource- and nonpecuniary costs, Λ, and $\lambda_{j,y}$, respectively, both appear.

Given this, prior to the bankruptcy decision, the current-period value function is

$$V\left(b_{j-1}, y, e, \nu, \lambda, x, j\right) = \max\{v^{d=1}(b_{j-1}, y, e, \nu, \lambda, x, j), v^{d=0}(b_{j-1}, y, e, \nu, \lambda, x, j)\}.$$

For the full information setting we assume I contains the entire state vector for the household; let $I = (y, e, \nu, x, \lambda, j)$. Abusing notation slightly, let $d(\cdot)$ now denote the decision *rule* governing default. As described earlier, this function drives the decision to repay a given debt or not, and hence depends on the full household state vector. Letting non-primed objects represent current period decisions, and using primed variables for objects dated one period ahead, we have the following zero profit condition for the intermediary. Simply put, it requires that the probability of default used to price debt must be consistent with that observed in the stationary equilibrium, implying that

$$\widehat{\pi}^{fi}\left(b, y, e, \lambda, j\right) = \sum_{e',\nu',\lambda',x'} d(b, e', \nu', x', \lambda', j+1) p_e\left(e'|e\right) p_\nu\left(\nu'\right) p_\lambda\left(\lambda'|\lambda\right) \pi_x\left(x'\right). \qquad (10)$$

Since $d\left(b, e', \nu', x', \lambda', j+1\right)$ specifies whether or not the agent will default in state $\left(e', \nu', x', \lambda'\right)$ tomorrow at debt level b, integrating over all such events *one period hence* produces the relevant estimated default risk $\widehat{\pi}^{fi}$. This expression also makes clear that knowledge of the persistent components (e, λ) is relevant for predicting default probabilities, and the more persistent these characteristics are, the more useful they become in assessing default risk.

Asymmetric Information

As we noted at the outset, earlier work, starting with Narajabad (2012), and including the work of Sánchez (2009) and Athreya, Tam, and Young (2012a), found that in past decades, unsecured credit market outcomes may well have been affected by informational frictions. In the latter article, asymmetric information governing *individual-level costs of bankruptcy* were shown to be consistent with a variety of features of the data from the 1980s and earlier. Thus, to evaluate the implications of loan guarantees under asymmetric information, we assume that nonpecuniary default costs, $\lambda_{j,y}$, is unobservable. With the exception of current household net worth following the bankruptcy decision in a period (which we denoted by a) all other household attributes, including educational attainment, age, and the current realization of the persistent component of income are assumed observable. To be clear, using household decisions rules and the distribution of households over the state space to infer a borrower's current net worth, a, is not useful because the net worth a is relevant to forecasting income, default risk, or anything else; it is not. Rather, it is because lenders want to draw a more precise inference on the current values of the persistent aspects of a household's state. In this case the inference is about the current realization of a household's λ, something that is clearly relevant to assessing default risk.

Let $p^*(\lambda|b, y, e, \nu, x, j)$ denote the equilibrium conditional probability of a household having a realized value of λ, given that they have observable characteristics y, e, ν, x, j, *and* that they have issued bonds of b units of face value. To construct the equilibrium assessment of default risk, $\pi^*(\cdot)$, lenders use their knowledge of household decision making and the joint (conditional) distribution of households over the state space to arrive at a probability distribution for the current value of a household's nonpecuniary default cost.[34] The best estimate of default risk is then given by

$$\widehat{\pi}^{pi}(b, y, e, \nu, x, j) = \sum\nolimits_{\lambda} p^*(\lambda|b, y, e, \nu, x, j)\, \widehat{\pi}^{fi}(b, y, e, \nu, x, \lambda, j).$$

[34] See Athreya, Tam, and Young (2012b) for details.

Equilibrium in the Credit Market

Here, we follow Athreya, Tam, and Young (2012b), and employ the Perfect Bayesian Equilibrium (PBE) concept to define equilibrium in the game between borrowers and lenders. Denote the state space for households by $\Omega = \mathcal{B}\times\mathcal{Y}\times\mathcal{E}\times\mathcal{V}\times\mathcal{L}\times\mathcal{J}\times\{0,1\} \subset \mathcal{R}^6\times\mathcal{Z}_{++}\times\{0,1\}$ and space of information as $\mathcal{I} \subset \mathcal{Y}\times\mathcal{E}\times\mathcal{V}\times\mathcal{L}\times\mathcal{J}\times\{0,1\}$. Let the stationary joint distribution of households over the state be given by $\Gamma(\Omega)$. Let the stationary equilibrium joint distribution of households over the state space Ω *and* loan requests b' be derived from the decision rules $\{b'^*(\cdot)$, $d^*(\cdot)\}$ and $\Gamma(\Omega)$, and be denoted by $\Psi^*(\Omega, b')$. Given $\Psi^*(\Omega, b')$, let $\mu^*(b')$ be the fraction of households (i.e., the marginal distribution of b') requesting a loan of size b'. Lastly, let the common beliefs of lenders on the household's state, Ω, given b', be denoted by $\Upsilon^*(\Omega|b')$.[35]

Definition 1 *A PBE for the credit market game of incomplete information consists of (i) household strategies for borrowing $b'^* : \Omega \to \mathcal{R}$ and default $d^* : \Omega\times\lambda\times\mathcal{E}\times\mathcal{V} \to \{0,1\}$, (ii) lenders' strategies for loan pricing $q^* : \mathcal{R}\times\mathcal{I} \to \left[0, \frac{1}{1+r}\right]$ such that q^* is weakly decreasing in b', and (iii) lenders' common beliefs about the borrower's state Ω given a loan request of size b', $\Upsilon^*(\Omega|b')$, that satisfy the following:*

1. ***Households optimize**: Given lenders' strategies, as summarized in the locus of prices $q^*(b', I)$, decision rules $\{b'^*(\cdot), d^*(\cdot)\}$ solve the household problem.*

2. ***Lenders optimize given their beliefs**: Given common beliefs $\Upsilon^*(\Omega|b')$, q'^* is the pure-strategy Nash equilibrium under one-shot simultaneous-offer loan-price competition.*

3. ***Beliefs are consistent with Bayes' rule wherever possible**: $\Upsilon^*(\Omega|b')$ is derived from $\Psi^*(\Omega, b')$ and household decision rules using Bayes rule whenever b is such that $\mu^*(b') > 0$.*

Equilibria are located through an iterative procedure. The interested reader is directed to the online appendix in Athreya, Tam, and Young (2012b), where we discuss the computational procedure used to solve for equilibria. As a quick summary, we define an iterative procedure that maps a set of pricing functions back into themselves, whose fixed points are PBE of the game between lenders and borrowers. This

[35] Recall that the stationary distribution of households over the state space alone is given by $\Gamma(\cdot)$.

procedure is monotonic, so starting from the upper limit yields convergence to the largest fixed point.[36]

Government

The government's budget constraint is motivated by two expenditures it must finance. Most importantly, it must finance payments to a lender to honor the loan guarantee program. Letting $\Gamma(a, y, e, \nu, x, \lambda, j)$ denote the invariant cumulative distribution function of households over the states, this is given by tax τ_2, which must satisfy

$$\tau_2 W \int y\omega_{j,y} e\nu d\Gamma(a, y, e, \nu, x, \lambda, j < j^*) =$$
$$\int \psi_{j+1|j} \frac{\widehat{\pi}(b(a, y, e, \nu, x, \lambda, j), I)}{1 + r + \phi} \max(0, b(a, y, e, \nu, x, \lambda, j) + \vartheta) \times$$
$$\frac{\theta b(a, y, e, \nu, x, \lambda, j)}{b(a, y, e, \nu, x, \lambda, j) + \vartheta} d\Gamma(a, y, e, \nu, x, \lambda, j). \tag{11}$$

In addition to financing loan guarantees, the government funds pension payments to retirees and to finance the loan guarantee system. The government budget constraint for pensions is

$$\tau_1 W \int (y\omega_{j,y} e\nu) d\Gamma(a, y, e, \nu, x, \lambda, j < j^*) =$$
$$W \int (\upsilon\omega_{j^*-1,y} y e_{j^*-1} \nu_{j^*-1} + \Upsilon) d\Gamma(a, y, e, \nu, x, \lambda, j \geq j^*). \tag{12}$$

Wage Determination

For both simplicity and substantive reasons, we assume constant and exogenous factor prices in our welfare calculations. In particular, we assume that the risk-free rate r is exogenous and determined by the world market for credit. Our approach follows several articles in the literature in abstracting from feedback effects onto risk-free rates of saving coming from changes in borrowing in the unsecured credit market, including Livshits, MacGee, and Tertilt (2007). This is a convenient abstraction and will be reasonable as long as guarantee programs are not inordinately generous.

[36] Uniqueness cannot be ensured, since $q = 0$ is a fixed point of our mapping. However, simple sufficient conditions exist to rule out $q = 0$ as the maximal fixed point; $\Lambda > 0$ is enough to guarantee the existence of an interval $[-\Lambda, 0]$ of risk-free debt. Sufficient conditions that ensure the existence of nontrivial default risk in equilibrium are not known.

Specifically, given r, profit maximization by domestic production firms implies that

$$W = (1-\alpha)\left(\frac{r}{\alpha}\right)^{\frac{\alpha}{\alpha-1}}, \tag{13}$$

where α is capital's share of income in a Cobb-Douglas aggregate production technology.

Stationary Equilibrium

We have already given the definition of equilibrium for the game between borrowers and lenders. The outcomes of that interaction were, of course, part of a larger fixed-point problem that included, among other things, the joint distribution of households over the state space, $\Gamma(\cdot)$, and the tax rates τ_1 and τ_2 needed to fund transfers and loan guarantees, respectively. But this joint distribution depended on household borrowing behavior, which in turn influenced the construction of $\Gamma(\cdot)$. Given this feedback, we will focus throughout on stationary equilibria in which all aggregate objects including, critically, the joint distribution $\Gamma(\cdot)$, remain constant over time under the decision rules that arise from household and creditor optimization.

Computing stationary equilibria requires two layers of iteration. We first specify the wage rate, interest rate, tax rates, and public sector transfer and loan guarantee policies. This allows us to solve the household's decision problem and locate the associated stationary distribution of households over the state space—all for a given *guess* of the equilibrium loan-pricing locus $q(\cdot)$. Our use of a risk-free rate-taking open economy allows us to iterate on the function $q(\cdot)$ without having to deal with any additional feedback from loan pricing to risk-free interest rates and wages. Once we have located a price function that is a fixed point under the stationary distribution induced by optimal household decision making (which we can denote by $q^*(\cdot)$), we need to check if the government budget constraint holds. Here, we must iterate again, this time on transfers and taxes. We use Brent's method to solve for the tax rate that satisfies the government budget constraint (re-solving for the fixed-point loan pricing function $q^*(\cdot)$ each time); whenever Laffer curve considerations arise, we choose the lower tax rate.

Parametrization

To assign values to model parameters, we proceed first by imposing standard values from the literature for measures of income risk, out-of-pocket expenses, risk aversion, and demographics. We then calibrate the remaining model parameters, which are those governing bankruptcy costs and the discount factor. The goal is to match, as well as possible, key facts about bankruptcy and unsecured credit markets in the United States, given income risk, risk aversion, and demographics. As discussed earlier, we follow the literature by calibrating to recent data and assuming *symmetric* information between borrowers and lenders.

The parametrization is relatively parsimonious and largely standard. First, as mentioned above, we directly assign values to household level income risk and risk aversion at values standard in the literature. The model period is taken to be one year. The income process is taken from Hubbard, Skinner, and Zeldes (1994), who estimate separate processes for non-high school (NHS), high school (HS), and college-educated (Coll) workers for the period 1982–1986.[37] Figure 3 displays the path $\omega_{j,y}$ for each type; the large hump present in the profile for college-educated workers implies that they will want to borrow early in life to a greater degree than the other types (despite their effective discount factor being somewhat higher because of higher survival probabilities). The process is discretized with 15 points for e and 3 points for ν. The resulting processes are

$$\begin{aligned} \log\left(e'\right) &= 0.95\log\left(e\right) + \epsilon' \\ \epsilon &\sim N\left(0, 0.033\right) \\ \log\left(\nu\right) &\sim N\left(0, 0.04\right) \end{aligned}$$

for non-high school agents,

$$\begin{aligned} \log\left(e'\right) &= 0.95\log\left(e\right) + \epsilon' \\ \epsilon &\sim N\left(0, 0.025\right) \\ \log\left(\nu\right) &\sim N\left(0, 0.021\right) \end{aligned}$$

for high school agents, and

$$\begin{aligned} \log\left(e'\right) &= 0.95\log\left(e\right) + \epsilon' \\ \epsilon &\sim N\left(0, 0.016\right) \\ \log\left(\nu\right) &\sim N\left(0, 0.014\right) \end{aligned}$$

for college agents. We normalize average income to 1 in model units, and in the data one unit roughly corresponds to $40,000 in income. When we construct the invariant distribution of the model, we assume

[37] In Athreya, Tam, and Young (2009) we study the effect of the rise in the volatility of labor income in the United States and find the effect on the unsecured credit market to be quantitatively small; the key parameter for default is the persistence of the shocks. We would find similar numbers if we adjusted the variance of the shocks upward to conform to more recent data.

households are born with zero assets and draw their first shocks from the stationary distributions.

To assign values for the idiosyncratic risk of out-of-pocket expenses, we choose the parameters for the expenditure shock x_j to be the annualized equivalent of those used in Livshits, MacGee, and Tertilt (2007). For pensions, we set $\upsilon = 0.35$ and $\Upsilon = 0.2$, yielding an average replacement rate of 55 percent, and assume an exogenous retirement age of $j^* = 45$. Relative risk aversion is set to $\sigma = 2$, as is standard, and a value that also avoids overstating the insurance problem faced by households. Lastly, with respect to demographics, we set the measures of the college (Coll), high school (HS), and non-high school (NHS) agents to 20, 58, and 22 percent, respectively, and the maximum lifespan to $J = 65$, corresponding to a calendar age of 85 years.

Table 1 in the main text displays the targeted moments and the implied ones from the model.[38] Table 2 in the main text displays the parameters associated with this calibration, along with the other parameters of the model (such as the cost of default Λ, which is set to match the observed \$1,200 filing cost). First, the default rates, measured as filings for Chapter 7 bankruptcy, are very close to the data. Second, the model does fairly well at matching the debt/income ratios in the data, measured as credit card debt divided by income (from the Survey of Consumer Finances 2004), although it reverses the order by understating debt for college types and overstating it for non-high school types. Lastly, the model generates a somewhat higher proportion of the observed fraction of borrowers while yielding smaller value of discharged debt to income ratio than currently measured.[39]

To parameterize the nonpecuniary costs of bankruptcy while limiting free parameters, we represent λ by a two-state Markov chain with realizations $\{\lambda_{L,y}, \lambda_{H,y}\}$ that are independent across households, but serially dependent with a symmetric transition matrix P_λ:

$$P_\lambda = \begin{bmatrix} p_\lambda & 1 - p_\lambda \\ 1 - p_\lambda & p_\lambda \end{bmatrix}.$$

The calibrated process suggests that nonpecuniary costs of bankruptcy are largely in the nature of a "type" for any given household. This interpretation arises because the benchmark calibration reveals λ to be very persistent, and therefore very unlikely to change during the part

[38] The calibrated parameters are obtained by minimizing the (equally weighted) sum of squared deviations between the data and moments from the invariant distribution of the model. Since the model is not linear, we cannot guarantee that there exists a set of parameters that makes this criterion zero; indeed, we find that such a vector does not seem to exist.

[39] If we had data on discharge by education type, we could permit the persistence of λ to vary by type and possibly match the aggregates more closely.

of life where unsecured credit is useful. This persistence is also what makes the model consistent with the observed ability of households to borrow substantial amounts but still default at a nontrivial rate. Despite this "implicit collateral," debts discharged in bankruptcy are still higher in the data; however, the discharge ratio from the data (obtained as the median debts discharged in bankruptcy divided by the median income of filers taken from the survey data of Sullivan, Warren, and Westbrook [2000]) is likely an overestimate, as it includes small business defaults that are generally large and not present in the model. The size of the values for λ are relatively large, implying that even the low cost types view default as equivalent to a loss of nearly 10 percent of consumption; thus, the primary source of implicit collateral in this model is stigma rather than pecuniary costs.

Table 3 in the main text presents a decomposition of defaults according to the various combinations of expense shock and stigma. The median shock for x and the high value of λ constitute only 3.55 percent of the population but are responsible for 58.11 percent of the defaults under symmetric information, while the high shock for x and high value for λ are 0.23 percent of the population and 6.66 percent of the defaults. Thus, defaults are clearly skewed toward households that experience an expenditure shock, consistent with the model of Livshits, MacGee, and Tertilt (2007).

Lastly, while omitted from the tables for brevity, the other relevant probability is that of the likelihood of default *given the receipt of an expenditure shock*. This distribution yields two pieces of information about the model. First, getting an expenditure shock, particularly the largest one, greatly increases the likelihood of default, all else equal. Second, the vast majority of households who receive such a shock still do not default. The reason for this is that the power of such shocks to drive default, while nontrivial, is still naturally limited by the wealth positions households take on as they move through the life cycle. Default is most likely to happen when one has substantial debts at the same time that one receives such a shock. This rules out relatively older households from being very susceptible; as seen in Figure 6, they have, in the main, already begun saving for retirement.[40]

[40] For agents with the relatively high value for λ in the model:
High expense shock: 26%
Median expense shock: 15%
Low expense shock (a value of zero)=1%
For agents with the relatively low value for λ in the model:
High expense shock: 17%
Median expense shock: 2%
Low expense shock (a value of zero)=0%
The numbers are very similar under asymmetric information.

REFERENCES

Aiyagari, S. Rao. 1994. "Uninsured Idiosyncratic Risk and Aggregate Saving." *Quarterly Journal of Economics* 109 (August): 659–84.

Andolfatto, David. 2002. "A Theory of Inalienable Property Rights." *Journal of Political Economy* 110 (April): 382–93.

Athreya, Kartik B. 2008. "Default, Insurance, and Debt over the Life-Cycle." *Journal of Monetary Economics* 55 (May): 752–74.

Athreya, Kartik B., Juan M. Sánchez, Xuan S. Tam, and Eric R. Young. 2012. "Bankruptcy and Delinquency in a Model of Unsecured Credit." Federal Reserve Bank of St. Louis Working Paper 2012-042A.

Athreya, Kartik B., Xuan S. Tam, and Eric R. Young. 2009. "Unsecured Credit Markets Are Not Insurance Markets." *Journal of Monetary Economics* 55 (January): 83–103.

Athreya, Kartik B., Xuan S. Tam, and Eric R. Young. 2012a. "Debt Default and the Insurance of Labor Income Risk." Federal Reserve Bank of Richmond *Economic Quarterly* 98 (Fourth Quarter): 255–307.

Athreya, Kartik B., Xuan S. Tam, and Eric R. Young. 2012b. "A Quantitative Theory of Information and Unsecured Credit." *American Economic Journal: Macroeconomics* 4 (July): 153–83.

Becker, Gary S. 1967. *Human Capital and the Personal Distribution of Income: An Analytical Approach (Woytinsky lecture)*. Ann Arbor, Mich.: Institute of Public Administration.

Calem, Paul S., Michael B. Gordy, and Loretta J. Mester. 2006. "Switching Costs and Adverse Selection in the Market for Credit Cards: New Evidence." *Journal of Banking and Finance* 30 (June): 1,653–85.

Castaneda, Ana, Javier Díaz-Giménez, and José-Víctor Ríos-Rull. 2003. "Accounting for Earnings and Wealth Inequality." *Journal of Political Economy* 111 (August): 818–57.

Chaney, Paul, and Anjan Thakor. 1985. "Incentive Effects of Benevolent Intervention: The Case of Government Loan Guarantees." *Journal of Public Economics* 26 (March): 169–89.

Chatterjee, Satyajit, P. Dean Corbae, Makoto Nakajima, and José-Víctor Ríos-Rull. 2007. "A Quantitative Theory of Unsecured Consumer Credit with Risk of Default." *Econometrica* 75 (November): 1,525–90.

Dávila, Julio, Jay H. Hong, Per Krusell, and José-Víctor Ríos-Rull. 2012. "Constrained Efficiency in the Neoclassical Growth Model with Uninsurable Idiosyncratic Shocks." *Econometrica* 80 (November): 2,431–67.

Dawsey, Amanda E., and Lawrence M. Ausubel. 2004. "Informal Bankruptcy." Working Paper (April).

Dubey, Pradeep, John Geanakoplos, and Martin Shubik. 2005. "Default and Punishment in General Equilibrium." *Econometrica* 73 (January): 1–37.

Fay, Scott A., Erik Hurst, and Michelle J. White. 1998. "The Bankruptcy Decision: Does Stigma Matter?" University of Michigan Working Paper 98-01 (January).

Gale, William G. 1990. "Federal Lending and the Market for Credit." *Journal of Public Economics* 42 (July): 177–93.

Gale, William G. 1991. "Economic Effects of Federal Credit Programs." *American Economic Review* 81 (March): 133–52.

Ghent, Andra C., and Marianna Kudlyak. 2011. "Recourse and Residential Mortgage Default: Evidence from U.S. States." *Review of Financial Studies* 24 (June): 3,139–86.

Gordon, Grey. 2015. "Evaluating Default Policy: The Business Cycle Matters." *Quantitative Economics*. Available at: www.qeconomics.org/ojs/forth/372/372-2.pdf.

Gross, David B., and Nicholas S. Souleles. 2002. "An Empirical Analysis of Personal Bankruptcy and Delinquency." *Review of Financial Studies* 15 (March): 319–47.

Hubbard, R. Glenn, Jonathan Skinner, and Stephen P. Zeldes. 1994. "The Importance of Precautionary Motives for Explaining Individual and Aggregate Saving." *Carnegie-Rochester Conference Series on Public Policy* 40 (June): 59–125.

Innes, Robert. 1990. "Investment and Government Intervention in Credit Markets when there is Asymmetric Information." *Journal of Public Economics* 46 (December): 347–81.

Jappelli, Tullio. 1990. "Who is Credit Constrained in the U.S. Economy?" *Quarterly Journal of Economics* 105 (February): 219–34.

Jeske, Karsten, Dirk Krueger, and Kurt Mitman. 2010. "Housing and the Macroeconomy: The Role of Implicit Guarantees for Government Sponsored Enterprises." Working Paper.

Jia, Ye. 2013. "Small Business Loan Guarantees as Insurance Against Aggregate Risks." *The B.E. Journal of Macroeconomics* 13 (January): 455–79.

Lacker, Jeffrey M. 1994. "Does Adverse Selection Justify Government Intervention in Loan Markets?" Federal Reserve Bank of Richmond *Economic Quarterly* 80 (Winter): 61–95.

Lelarge, Claire, David Sraer, and David Thesmar. 2010. "Entrepreneurship and Credit Constraints: Evidence from a French Loan Guarantee Program." In *International Differences in Entrepreneurship*, edited by Josh Lerner and Antoinette Schoar. Chicago: University of Chicago Press, 243–73.

Li, Wenli. 1998. "Government Loan Guarantee and Grant Programs: An Evaluation." Federal Reserve Bank of Richmond *Economic Quarterly* 84 (Fall): 25–51.

Li, Wenli. 2002. "Entrepreneurship and Government Subsidies: A General Equilibrium Analysis." *Journal of Economic Dynamics and Control* 26 (September): 1,815–44.

Livshits, Igor, James MacGee, and Michèle Tertilt. 2007. "Consumer Bankruptcy: A Fresh Start." *American Economic Review* 97 (March): 402–18.

Malysheva, Nadezhda, and John R. Walter. 2013. "How Large Has the Federal Financial Safety Net Become?" Federal Reserve Bank of Richmond Working Paper 10-03R (March).

Narajabad, Borghan N. 2012. "Information Technology and the Rise of Household Bankruptcy." *Review of Economic Dynamics* 15 (October): 526–50.

Pijoan-Mas, Josep. 2006. "Precautionary Savings or Working Longer Hours?" *Review of Economic Dynamics* 9 (April): 326–52.

Sánchez, Juan. 2009. "The Role of Information in the Rise of Consumer Bankruptcy." Federal Reserve Bank of Richmond Working Paper 2009-4 (April).

Smith, Bruce D., and Michael J. Stutzer. 1989. "Credit Rationing and Government Loan Programs: A Welfare Analysis." *Real Estate Economics* 17 (2): 177–93.

Sullivan, Teresa A., Elizabeth Warren, and Jay Lawrence Westbrook. 1999. *As We Forgive Our Debtors: Bankruptcy and Consumer Credit in America*. New York: Oxford University Press.

Sullivan, Teresa A., Elizabeth Warren, and Jay Lawrence Westbrook. 2000. *The Fragile Middle Class: Americans in Debt*. New Haven, Conn.:Yale University Press.

Tam, Xuan S. 2009. "Long Term Contracts in Unsecured Credit Markets." Manuscript.

Telyukova, Irina A. 2013. "Household Need for Liquidity and the Credit Card Debt Puzzle." *Review of Economic Studies* 80 (January): 1,148–77.

Walter, John R., and John A. Weinberg. 2002. "How Large is the Federal Financial Safety Net?" *CATO Journal* 21 (Winter): 369–94.

Williamson, Stephen. 1994. "Do Informational Frictions Justify Federal Credit Market Programs?" *Journal of Money, Credit and Banking* 26 (August): 523–44.

Zame, William R. 1993. "Efficiency and the Role of Default when Security Markets are Incomplete." *American Economic Review* 83 (December): 1,142–64.

Zeldes, Stephen P. 1989. "Consumption and Liquidity Constraints: An Empirical Investigation." *Journal of Political Economy* 97 (April): 305–46.

Economic Quarterly—Volume 100, Number 4—Fourth Quarter 2014—Pages 353–385

Optimal Institutions in Economies with Private Information: Exclusive Contracts, Taxes, and Bankruptcy Law

Borys Grochulski and Yuzhe Zhang

In economies with private information, it is typically optimal to prohibit or otherwise discourage a subset of trades that individual agents want to enter. Economists often refer to such optimal distortions as wedges. In this article, we use a simple private-information Mirrleesian economy to, first, show examples of these wedges and, second, discuss institutions that may be used to implement them in practice. Implementation of wedges has received a lot of attention in the literature recently because it can lead to large improvements in economic outcomes.[1]

In a Mirrleesian economy, agents are privately informed about their own skills or productivity.[2] By exerting less effort, a highly productive agent can supply the same amount of effective labor services as a less productive one. To an outside observer, these two agents will appear

Zhang is an associate professor with Texas A&M University. Grochulski is a senior economist with the Richmond Fed. The authors would like to thank Arantxa Jarque, Joe Johnson, Ned Prescott, and Zhu Wang for their helpful comments. The views expressed in this article are those of the authors and not necessarily those of the Federal Reserve Bank of Richmond or the Federal Reserve System. E-mail: zhangeager@tamu.edu; borys.grochulski@rich.frb.org.

[1] See Huggett and Parra (2010) and Farhi and Werning (2013) for calibrations showing significant welfare gains.

[2] Mirrlees (1971) started a large literature studying an optimal redistribution problem with private skills.

identical because effort is not observable and the quantity of effective labor services produced by the two agents is the same.

Private information about agents' productivity and effort becomes a problem when agents seek insurance against shocks to their individual productivity levels.[3] Under full information, it is efficient to equate all agents' marginal utility of consumption, i.e., to have more productive agents supply more labor but consume the same as the less productive ones. With private information, full consumption insurance is impossible due to moral hazard. Productive agents can shirk, i.e., put in little effort, produce only as much effective labor as the unproductive ones, and claim a bad realization of the productivity shock. In order to elicit effort, higher labor supply must be compensated with higher consumption. From the ex ante perspective, this is costly as risk-averse agents prefer stable consumption profiles.

Moreover, moral hazard becomes more severe as society gets richer. With diminishing marginal utility of consumption, the amount of consumption compensation needed to elicit a given amount of effort increases as the average level of consumption increases. Richer agents must be exposed to more consumption risk. For this reason, in the intertemporal setting, it is efficient to suppress wealth accumulation. A benevolent social planner would front-load the agents' consumption and suppress the accumulation of wealth in order to moderate the cost of providing effort incentives in the future.[4] At the optimal allocation, therefore, for incentive reasons, agents are exposed to risky and low future consumption profiles. This front-loading of consumption leads to the so-called intertemporal wedge in the optimal allocation: The agents' shadow interest rate is smaller than the real interest rate (the rate of return on physical capital).

In a decentralized setting, an individual agent does not internalize future costs of incentives. Instead, due to the intertemporal wedge, the real interest rate is high enough for the agent to want to smooth her consumption profile by accumulating more wealth. In other words, the agent wants to self-insure by increasing her savings. If the optimal allocation with its intertemporal wedge is to be supported as an equilibrium of a decentralized economy—one in which agents make their own consumption and savings decisions—some distortions must be introduced that prohibit or otherwise prevent agents from trading away from the optimum.

[3] Diamond and Mirrlees (1978) is perhaps the first study of the dynamic insurance problem with private information about skills. The more recent literature surveyed in Kocherlakota (2010) extends the analysis of this problem.

[4] Rogerson (1985) makes this point in a related repeated moral hazard model, where agents learn the productivity of their effort only after they exert it.

In this article, we discuss three sets of institutions that can provide a distortion needed to make agents' individual optimization consistent with the private-information constrained optimal allocation. We begin by formally defining a simple, two-period Mirrleesian economy in Section 1. We provide a concise characterization of optimal wedges in Section 2.

In Section 3, we discuss competitive equilibrium with exclusive contracts originally studied in Prescott and Townsend (1984) and Atkeson and Lucas (1992). In this model of market interaction, firms sign agents to comprehensive lifetime labor and consumption contracts. These contracts are exclusive: Once under contract with a firm, an agent is not to trade with anyone else. The firm takes over production and savings/capital accumulation decisions and gives the agent a comprehensive schedule for future consumption and labor supply required of the agent. That schedule is incentive compatible, i.e., robust to agents' moral hazard. Lifetime utility delivered by this contract is the price at which agents sell and over which firms compete. In order to preserve incentives, the comprehensive lifetime contract must be exclusive, i.e., it must prevent the agent from retrading or simply postponing his consumption by saving. To enforce this broad exclusivity, the firm must monitor the agent's consumption, and in particular it must make sure the agent consumes and does not save.

There are two limitations of this approach to decentralization. First, if decentralization is defined as a model in which agents make their own consumption and savings decisions, the Prescott-Townsend-Atkeson-Lucas exclusive contracts model is not a decentralization. Rather, as Atkeson and Lucas (1992) put it, it is a model of competing principals, or social planners, each of whom, while maximizing own profit from the relationship with the agent, internalizes the incentive costs in designing the comprehensive contracts, exactly as does the benevolent social planner. Second, exclusive contracts prohibiting agents' private savings and committing them to not quitting the firm are not observed in practice, which makes the model very unrealistic (i.e., not relevant empirically).

In Section 4, we discuss the taxation model studied in the recent literature known as New Dynamic Public Finance (see Kocherlakota 2010). In this model, agents make their own consumption, savings, and labor supply decisions subject to taxes. The government designs a system of taxes and transfers to provide optimal incentives while insuring the agents against their productivity shocks. In this model, trade is decentralized and the assignment of the monitoring duty is realistic: In practice, governments typically monitor people's savings and wealth (capital they hold) in order to assess capital income taxes. The optimal tax prescriptions generated by this analysis, however, are not very realistic. The model implies that marginal capital tax rates should vary in complicated ways with labor income in order to deter joint deviations consisting of simultaneously shirking and saving. The lack of realism, of course, is not a basis for rejection of a piece of normative analysis. It does, however, invite the question of what other mechanisms could also implement the optimal allocation in this intertemporal Mirrleesian environment.

In Section 5, we discuss an alternative institutional setup capable of providing the needed distortions in agents' private consumption, savings, and labor supply decisions. We follow Grochulski (2010) in considering a model with private extension of unsecured credit and default regulated by government-enforced bankruptcy law. In this model, agents obtain insurance by taking out unsecured, defaultable loans while simultaneously investing in riskless bonds (which can also be interpreted as safe deposits). Insurance is provided to unproductive agents by granting them discharge of the loan they owe while letting them keep the payoff from the bonds they hold, up to a limit that is determined by the optimal allocation. Critically, this limit must apply to all wealth the agent holds in order to, again, prevent the agent from over-saving. In particular, because unsecured loans are priced under the presumption of no-shirking, saving behind the back of the agent's unsecured lenders must be prevented because saving is complementary with shirking. In this decentralization, taxes are only used to fund government spending, and not to provide incentives or insurance to the taxpayers. An attractive feature of this decentralization is that the assignment of the monitoring duty (to the bankruptcy court), the provision of insurance through unsecured credit, and the restrictions on debt discharge necessitated by moral hazard are all realistic.

Section 6 concludes. Our analysis underscores the multiplicity of possible implementation and, therefore, the difficulty in using private information as a basis for normative analysis of any one such institution. Yet, the differences in the degree of decentralization and empirical relevance of the three sets of institutions we discuss make the implementation exercise considered here useful in thinking about the implications of private information for economic outcomes and observed institutions.

1. ENVIRONMENT

We use a simple Mirrleesian environment very similar to that studied in Section 3 of Kocherlakota (2005). There are two dates, $t = 0, 1$, and a unit measure of ex ante identical agents. At each date, a single consumption good is produced from capital and labor. Technology of production is described by the production function $F : \mathbb{R}_+ \times \mathbb{R}_+ \to \mathbb{R}_+$, where $F(k, y)$ denotes the amount of the consumption good produced from k units of capital and y units of labor. We assume that F is strictly concave, twice differentiable, exhibits constant returns to scale, and satisfies the usual Inada conditions. For simplicity, we also assume that capital used in production depreciates fully.

The initial aggregate endowment of physical capital, K_0, is distributed uniformly, i.e., each agent holds $k_0 = K_0$ units of capital at $t = 0$. In addition to physical capital, agents supply effective labor services, which they generate from skill and labor effort. At $t = 0$, all agents have identical skills: One unit of labor effort produces one unit of the effective labor input. Thus, if all agents provide l_0 units of labor effort, the aggregate supply of effective labor is $Y_0 = y_0 = l_0$ at this date. The aggregate labor supply Y_0 and the initial capital K_0 produce a total of $F(K_0, Y_0)$ units of the consumption good at $t = 0$. This amount can be consumed or saved as capital available at $t = 1$, K_1. At $t = 1$, agents' skills are subject to a stochastic, individual skill shock θ. Thus, agents become heterogeneous at $t = 1$. For simplicity, we assume that θ takes the value of either 0 or 1. Agents whose individual skill realization is $\theta = 1$ can convert one unit of labor effort into one unit of effective labor input, just like at $t = 0$. Agents whose individual realization of skill is $\theta = 0$, however, can convert a unit of effort into zero units of effective labor, i.e., are completely unproductive.[5] We will denote the probability of the shock realization θ by $\pi_\theta > 0$, for both $\theta \in \{0, 1\}$. Each agent's individual realization of θ is his private information.

In this environment, a (type-identical) allocation A is a list of nonnegative numbers

$$\{c_0, (c_{1\theta})_{\theta \in \{0,1\}}, l_0, (l_{1\theta})_{\theta \in \{0,1\}}, Y_0, Y_1, K_1\},$$

where c_0 is per capita consumption at $t = 0$, and $c_{1\theta}$ is per capita consumption of agents with skill θ at $t = 1$. The expected ex ante utility a representative agent obtains under allocation A is given by

$$u(c_0) - v(l_0) + \beta \sum_{\theta \in \{0,1\}} \pi_\theta (u(c_{1\theta}) - v(l_{1\theta})), \tag{1}$$

where u and v are strictly increasing, twice differentiable functions with $u'' < 0$, $v'' > 0$, and $v(0) = 0$.

[5] Our analysis can be generalized to the case of positive productivity in both states without changing our main results.

An allocation is resource feasible if it satisfies the following resource constraints

$$c_0 + K_1 + G_0 \leq F(K_0, Y_0), \tag{2}$$

$$Y_0 = l_0, \tag{3}$$

$$\sum_{\theta \in \{0,1\}} \pi_\theta c_{1\theta} + G_1 \leq F(K_1, Y_1), \tag{4}$$

$$Y_1 = \sum_{\theta \in \{0,1\}} \pi_\theta \theta l_{1\theta}, \tag{5}$$

where G_t is a fixed level of government spending in period $t = 0, 1$. It is without loss of generality to only consider allocations with $l_{10} = 0$. This is because the agents whose skill at $t = 1$ is $\theta = 0$ cannot provide any effective labor into production, and so it would be a waste to have them exert a positive amount of labor effort.

Because skills at $t = 1$ are private information, we restrict attention to allocations that satisfy the following incentive compatibility (IC) constraint

$$u(c_{11}) - v(l_{11}) \geq u(c_{10}). \tag{6}$$

This constraint requires that the skilled agents at $t = 1$ do not prefer to mimic the unskilled ones by providing zero effective labor and consuming the amount that allocation A assigns to unskilled agents, c_{10}. By the Revelation Principle, restricting attention to IC allocations is without loss of generality.

2. OPTIMAL ALLOCATIONS AND WEDGES

Allocation A is incentive-optimal, or optimal for short, if it is resource feasible, incentive compatible, and if among all resource feasible and incentive compatible allocations it maximizes the ex ante welfare of the representative agent. Thus, A is optimal if and only if it solves the following social planning problem (SPP): maximize (1) subject to the resource constraints (2)–(5) and the IC constraint (6).

Denote an optimal allocation by

$$A^* = \{c_0^*, (c_{1\theta}^*)_{\theta \in \{0,1\}}, l_0^*, (l_{1\theta}^*)_{\theta \in \{0,1\}}, Y_0^*, Y_1^*, K_1^*\}$$

and by U^* the level of ex ante expected utility, i.e., the value of the objective (1), attained at the optimal allocation. As noted above, since effort of an agent whose $\theta = 0$ is unproductive, $l_{10}^* = 0$. It is straightforward to use the first-order conditions of the SPP to demonstrate the following properties of A^* (see Kocherlakota 2005, Section 3):

$$l_0^* > 0,\ l_{11}^* > 0,\ c_{11}^* > c_{10}^* > 0,$$

$$\begin{aligned} u(c_{11}^*) - v(l_{11}^*) &= u(c_{10}^*), & (7)\\ u'(c_0^*) &= \frac{v'(l_0^*)}{w_0^*}, & (8)\\ u'(c_{11}^*) &= \frac{v'(l_{11}^*)}{w_1^*}, & (9)\end{aligned}$$

and

$$\frac{1}{u'(c_0^*)} = \frac{1}{r_1^*\beta}\mathbb{E}\left[\frac{1}{u'(c_1^*)}\right], \tag{10}$$

where

$$r_t^* = F_1(K_t^*, Y_t^*),\ w_t^* = F_2(K_t^*, Y_t^*) \text{ for } t = 0, 1.$$

Equation (7) says that the IC constraint is binding at A^*. Equations (8) and (9) tell us that, if labor services are paid their marginal product, the productive agents' disutility of making one extra dollar of labor income at A^* is equal to the utility of consuming it. This means that there are no intratemporal distortions (wedges) at A^*: If agents are given the allocation A^* and can earn in period $t = 0, 1$ wages $w_t^* = F_2(K_t^*, Y_t^*)$, they would not want to deviate from the optimal allocation by working a different amount than what the optimal allocation prescribes.

Equation (10) is the so-called Inverse Euler Equation.[6] Bringing the expectation inside the inverse function, using Jensen's inequality and the fact that $c_{11}^* \neq c_{10}^*$, we get that the optimal allocation A^* satisfies

$$u'(c_0^*) < r_1^*\beta\mathbb{E}[u'(c_1^*)]. \tag{11}$$

This inequality tells us that there is a distortion in the intertemporal margin at A^*. This distortion is often referred to as the *intertemporal wedge*: If capital services are paid their marginal product, $r_t^* = F_1(K_t^*, Y_t^*)$, the disutility of reducing consumption c_0 by a small amount and investing it in capital K_1 is smaller than the resulting expected benefit of having more capital at $t = 1$. Thus, if agents are given allocation A^* and can save (accumulate capital) without any distortions, they would like to trade away from the optimum A^* by saving more than what is socially optimal. In this sense, agents are savings-constrained at A^*. Inequality (11) is important because it makes clear that if A^*

[6] This condition is first obtained in Diamond and Mirrlees (1978). Rogerson (1985) derives this equation in a moral hazard model. Golosov, Kocherlakota, and Tsyvinski (2003) derive this equation in a general dynamic Mirrlees economy with privately evolving skills.

is to be consistent with agents' individual utility maximization—a necessary condition for equilibrium—the intertemporal margin cannot be left undistorterd.

In addition, we will use the following two properties of the optimum A^*:

$$r_1^* \beta u'(c_{11}^*) < u'(c_0^*) < r_1^* \beta u'(c_{10}^*), \tag{12}$$

and

$$w_1^* l_{11}^* - c_{11}^* > w_1^* l_{10}^* - c_{10}^*. \tag{13}$$

The inequalities in (12) tell us that agents are insurance-constrained at A^*. If agents could insure (or hedge) their individual shocks θ at a fair-odds premium, they would like to trade away from A^* by purchasing additional insurance. Inequality (13) shows that the optimal allocation A^* delivers a state-contingent transfer from the productive agents to the unproductive ones at $t = 1$.[7]

3. IMPLEMENTATION WITH EXCLUSIVE PRIVATE CONTRACTS

In this section, we discuss decentralization of the optimal allocation as an equilibrium in a competitive market economy with comprehensive, exclusive contracts. This decentralization follows Prescott and Townsend (1984), Atkeson and Lucas (1992), and Golosov and Tsyvinski (2007).

Firms sign agents to comprehensive lifetime-utility contracts. In such a contract, the firm promises lifetime utility $\bar{U}$ in return for the agent's capital k_0. Utility $\bar{U}$ is delivered by an assignment of consumption, which the agent gets from the firm, and of effective labor, which the agent is to deliver to the firm. The firm combines these inputs to produce output according to the production function F. Importantly, the contract is exclusive, i.e., the agent signs off her right to trade with anybody else. The utility value $\bar{U}$ is determined in equilibrium; each firm takes it as given. Without loss of generality, we assume that government expenditure G_t is funded by non-distortionary, lump-sum taxes $T_t = G_t$, $t = 0, 1$.

Firms maximize profits. The problem they solve is to design a lifetime-utility contract that delivers to the agent the market level of

[7] Proof of (12) follows simply from the first-order conditions of the social planning problem. To prove (13), note that $c_{10}^* - w_1^* l_{10}^* - c_{11}^* + w_1^* l_{11}^* = \frac{u_{11}'^*)(c_{10}^* - c_{11}^*) - v_{11}'^*)(l_{10}^* - l_{11}^*)}{u_{11}'^*)} > \frac{(u(c_{10}^*) - v(l_{10}^*)) - (u(c_{11}^*) - v(l_{11}^*))}{u_{11}'^*)} = 0$, where the first equality uses (9) and the inequality follows from the property that if function $f(\cdot)$ is strictly concave, then $f'(y)(x-y) > f(x) - f(y)$ for any $x \neq y$.

utility $\bar{U}$ at minimum cost. To do so, the firm chooses a consumption-labor plan for the agent,

$$\psi = \{c_0, (c_{1\theta})_{\theta\in\{0,1\}}, l_0, (l_{1\theta})_{\theta\in\{0,1\}}\},$$

and the amount of capital it saves for date 1, K_1^{f}. Given that each firm provides the same $\bar{U}$, agents are indifferent among firms. Hence, each active firm is able to attract a non-zero mass of agents. The equilibrium number of firms is indeterminate. In sum, the firm's problem is as follows:

$$\begin{aligned}
\Pi_0(\bar{U}) \quad = \quad & \max_{\psi, K_1^{\mathrm{f}}} \; F(K_0, Y_0) - (c_0 + G_0) - K_1^{\mathrm{f}} \\
& s.t. \\
& u(c_0) - v(l_0) + \beta \sum_{\theta\in\{0,1\}} \pi_\theta (u(c_{1\theta}) - v(l_{1\theta})) = \bar{U}, \quad (14) \\
& \sum_{\theta\in\{0,1\}} \pi_\theta c_{1\theta} + G_1 \leq F(K_1^{\mathrm{f}}, \pi_1 l_{11}), \quad (15) \\
& u(c_{11}) - v(l_{11}) \geq u(c_{10}). \quad (16)
\end{aligned}$$

The objective is the profit the firm makes, measured here in units of date-0 capital. Condition (14) is the promise-keeping constraint requiring that the contract indeed deliver $\bar{U}$ to each agent who signs with the firm. Condition (15) ensures that the firm has enough capital and effective labor at $t = 1$ to cover its obligations toward the agents. In particular, the fraction π_1 of the agents signed by the firm are productive at that date. The IC constraint (16) ensures that the productive agents prefer to supply labor into the production process.

Definition 1 ***Competitive equilibrium with exclusive contracts*** *consists of a comprehensive contract*

$$\hat{\psi} = \{\hat{c}_0, (\hat{c}_{1\theta})_{\theta\in\{0,1\}}, \hat{l}_0, (\hat{l}_{1\theta})_{\theta\in\{0,1\}}\},$$

the firm's capital plan $\hat{K}_1^{\mathrm{f}}$*, and a price* $\bar{U}$ *such that (a) the pair* $(\hat{\psi}, \hat{K}_1^{\mathrm{f}})$ *attains* $\Pi_0(\bar{U})$ *(profit maximization), and (b)* $\Pi_0(\bar{U}) = 0$ *(free entry).*

Theorem 1 *Competitive equilibrium with exclusive contracts is efficient, i.e.,* $\bar{U} = U^*$ *and the pair* $(\hat{\psi}, \hat{K}_1^{\mathrm{f}})$ *replicates* A^**.*

Proof. In the Appendix. ■

The proof of this version of the First Welfare Theorem follows immediately from the duality between the SPP and the firm's profit-maximization problem. With exclusive contracts, the constraints in the two problems are the same except that the planning problem takes the aggregate resource constraint as given and maximizes the utility of the agent, while the firm's problem takes the utility to be delivered to the agent as given and maximizes profit. In equilibrium with free entry, profits must be zero, which guarantees resource feasibility of the solution.

The efficiency of the outcome of competition with exclusive comprehensive contracts is appealing. However, this decentralization concept has shortcomings. First, as pointed out by Atkeson and Lucas (1992, 444), "The difficulty with using such an equilibrium as a model of observed market arrangements stems from the capability to monitor individual wealth positions granted to this intermediary, relative to the capabilities of actual financial institutions." In particular, the intertemporal wedge (11) implies that under the optimal contract ψ agents would benefit from saving a portion of c_0 for consumption at $t = 1$. Under the comprehensive contract, failing to consume the whole c_0 constitutes a breach of contract.

Second, the exclusivity requirement used in this competition concept seems like a strong one. Trade only happens ex ante. This feature, on the one hand, delivers efficiency of the outcome really easily because with an exclusive right to trade with an individual the firm can internalize the incentives, which is evident in the IC constraint entering the firm's problem in exactly the same form as it enters the SPP. But, on the other hand, it leads the model to predict that agents make no economic decisions in their lives other than their initial signing with a firm and subsequently reporting their productivity shocks θ to that firm.

For these reasons, it is useful to examine other, more decentralized implementations of the incentive-optimal allocation. In particular, it is natural to consider the possibility that monitoring of the agents' trades is delegated to the government, as in practice the government does monitor people's savings for the purpose of collecting taxes. Savings, or wealth more broadly, are also monitored in personal bankruptcy. In the next two sections, we discuss implementations with, respectively, distortionary capital income taxes and personal bankruptcy.

4. IMPLEMENTATION WITH INCOME-CONTINGENT CAPITAL TAXES

In this section, we discuss a decentralization in which the optimal intertemporal wedge is implemented by distortionary capital income taxes. This approach to implementing wedges has been used extensively in the literature known as *New Dynamic Public Finance*.[8] We will therefore refer to this approach as the NDPF decentralization. In particular, we follow Kocherlakota (2005) closely.

[8] Kocherlakota (2010) provides a comprehensive survey.

The NDPF decentralization does not use exclusive contracts. Instead, agents rent capital out and supply labor to firms in spot markets every period. Agents also make their own consumption and saving decisions, subject to taxes. In particular, in order to prevent over-saving (which is complementary with shirking), the capital income agents earn at $t = 1$ is taxed at a constant marginal rate τ. We start out with a discussion of the natural tax rate, τ^*, that closes the intertemporal wedge present at the optimal allocation (recall inequality 11).

The Natural Intertemporal Tax Rate

NDPF starts out with a natural connection between the intertemporal wedge and a tax on savings and points out a problem with it. The natural connection is as follows. The agents' incentive to over-save, relative to A^*, can be removed if a proportional tax is imposed on capital income in period 1 with the tax rate

$$\tau^* = 1 - \frac{u'(c_0^*)}{r_1^* \beta \mathbb{E}[u'(c_1^*)]}. \tag{17}$$

The fact that τ^* is strictly positive could provide an efficiency-based role for positive capital taxes. Such an efficiency-based argument is something that the optimal taxation literature started by Ramsey (1927) has been lacking.

The problem with this connection that NDPF points out is as follows. There is one more wedge at the optimum A^*:

$$u'(c_0^*) < r_1^* \beta u'(c_{10}^*). \tag{18}$$

We can call this wedge a *shirker's intertemporal wedge*. If an agent shirks, i.e., decides at $t = 0$ that he will exert zero effort at $t = 1$ in both states θ (that means, he will supply zero units of effective labor even when his skill shock realization is $\theta = 1$), then the intertemporal tradeoff relevant to him is not one between the marginal utility $u'(c_0^*)$ and the expected marginal utility $\mathbb{E}[u'(c_1^*)]$, but rather that between $u'(c_0^*)$ and $u'(c_{10}^*)$, because a shirker knows already at $t = 0$ the consumption he will be assigned at $t = 1$. Since θ is not publicly observable, the allocation A^* assigns consumption in period 1 on the basis of the agent's report or, equivalently, the agent's observed effective labor input $y_1 = \theta l_1$. A shirking agent will be assigned at $t = 1$ consumption c_{10}^* with probability one because he always produces $y_1 = 0$.

Since $u'(c_{10}^*) > \mathbb{E}[u'(c_1^*)]$, the "natural" tax rate τ^* in (17), although high enough to deter over-saving by an agent who does not shirk, is not high enough to deter a shirker from over-saving. Thus, a shirker prefers to over-save and shirk over simply shirking. Because the IC constraint (6) is binding at A^*, shirking without over-saving gives an agent as much utility as non-shirking. Thus, the "joint deviation" plan of both shirking and over-saving gives the agent more utility than non-shirking. Therefore, under the simple proportional tax τ^* agents would choose to over-save and shirk. Thus, τ^* is not sufficient to implement A^*.

Labor-Income-Contingent Capital Tax Rates

To deal with this problem, Kocherlakota (2005) uses a tax system in which the marginal tax rates applied to capital income are contingent on labor income of the agent. In particular, let τ_{10} be the capital income tax rate applied to all agents whose labor income is zero at $t = 1$, and τ_{11} be the rate applied to those with positive labor income at that date. The joint deviation of shirking and over-saving can now be deterred with the tax rate τ_{10}, while the tax rate τ_{11} can be set so as to balance out the saving incentives for a non-shirker.

In particular, the after-tax Euler equation of a shirking agent is

$$u'(c_0) = r_1\beta(1 - \tau_{10})u'(c_1). \tag{19}$$

Thus, if the capital tax rate conditional on zero labor income at date 1 is set at

$$\tau_{10} = 1 - \frac{u'(c_0^*)}{r_1^*\beta u'(c_{10}^*)}, \tag{20}$$

the shirker's wedge present at the optimal allocation, (18), is closed, i.e., the Euler equation (19) holds and a shirking agent no longer desires to over-save relative to A^*. The shirker thus can no longer obtain more discounted utility than $u(c_0^*) + \beta u(c_{10}^*) = U^*$, so agents' incentive to shirk is removed. The other tax rate, τ_{11}, can now be set so as to deter the non-shirker from over- or under-saving. In particular, with taxes $\tau_1 = (\tau_{10}, \tau_{11})$, the non-shirker's Euler equation must hold:

$$\begin{aligned} u'(c_0^*) &= r_1^*\beta\mathbb{E}[(1 - \tau_1)u'(c_1^*)] \\ &= r_1^*\beta\pi_0(1 - \tau_{10})u'(c_{10}^*) + r_1^*\beta\pi_1(1 - \tau_{11})u'(c_{11}^*). \end{aligned} \tag{21}$$

Using (20) and solving for τ_{11}, we obtain

$$\tau_{11} = 1 - \frac{u'(c_0^*)}{r_1^*\beta u'(c_{11}^*)}. \tag{22}$$

By making use of the information contained in labor income earned by each agent at $t = 1$, the two-rate tax system can deter over-saving for both a shirker and a non-shirker, which allows for implementation of the optimum A^* as a competitive equilibrium without a need for the fully exclusive contracts discussed in the previous section.

Formally, agents rent their capital to firms, choose their own labor supply, savings, and consumption to maximize their expected utility

(1) subject to the budget constraints

$$\begin{aligned} c_0 + k_1 &\leq w_0 l_0 + r_0 k_0 - T_0, \\ c_{1\theta} &\leq w_1 \theta l_{1\theta} + (1 - \tau_1(w_1 \theta l_{1\theta})) r_1 k_1 - T_1(w_1 \theta l_{1\theta}), \quad \theta = 0, 1, \end{aligned}$$

where

$$\tau_1(w_1 \theta l_{1\theta}) = \begin{cases} \tau_{10} & \text{if} \quad w_1 \theta l_{1\theta} = 0, \\ \tau_{11} & \text{if} \quad w_1 \theta l_{1\theta} > 0, \end{cases}$$

T_0 is a lump-sum tax at $t = 0$, and $T_1(w_1 \theta l_{1\theta})$ is a quasi-lump-sum tax at $t = 1$:

$$T_1(w_1 \theta l_{1\theta}) = \begin{cases} T_{10} & \text{if} \quad w_1 \theta l_{1\theta} = 0, \\ T_{11} & \text{if} \quad w_1 \theta l_{1\theta} > 0. \end{cases}$$

Definition 2 *Given a set of taxes* $(\tau_{10}, \tau_{11}, T_{10}, T_{11})$, ***competitive equilibrium with taxes*** *consists of an allocation*

$$\hat{A} = \{\hat{c}_0, (\hat{c}_{1\theta})_{\theta \in \{0,1\}}, \hat{l}_0, (\hat{l}_{1\theta})_{\theta \in \{0,1\}}, \hat{Y}_0, \hat{Y}_1, \hat{K}_1\},$$

the agent's individual saving choice $\hat{k}_1$, *and prices* $\{r_t, w_t\}_{t=0,1}$ *such that: (a) the values* $\hat{c}_0$, $(\hat{c}_{1\theta})_{\theta \in \{0,1\}}$, $\hat{l}_0$, $(\hat{l}_{1\theta})_{\theta \in \{0,1\}}$, *and* $\hat{k}_1$ *solve the agent's utility maximization problem, (b) capital and labor are paid their respective marginal products*

$$r_t = F_1(\hat{K}_t, \hat{Y}_t), \; w_t = F_2(\hat{K}_t, \hat{Y}_t) \text{ for } t = 0, 1,$$

and (c) consumption, labor, and capital markets clear

$$\begin{aligned} \hat{c}_0 + \hat{K}_1 + G_0 &= F(\hat{K}_0, \hat{Y}_0), \\ \hat{K}_1 &= \hat{k}_1, \\ \hat{Y}_0 &= \hat{l}_0, \\ \sum_{\theta \in \{0,1\}} \pi_\theta \hat{c}_{1\theta} + G_1 &= F(\hat{K}_1, \hat{Y}_1), \\ \hat{Y}_1 &= \sum_{\theta \in \{0,1\}} \pi_\theta \theta \hat{l}_{1\theta}. \end{aligned}$$

Note that this definition implies that in equilibrium the government balances its budget every period. Taxes $(\tau_{10}, \tau_{11}, T_{10}, T_{11})$ implement an optimum A^* if there exists a competitive equilibrium with taxes such that the equilibrium allocation $\hat{A}$ coincides with the optimal allocation A^*.

Theorem 2 *Let capital income tax rates* τ_{10}, τ_{11} *be as in, respectively, (20) and (22), let the lump-sum tax at* $t = 0$ *be*

$$T_0 = w_0^* l_0^* + r_0^* K_0^* - K_1^* - c_0^*,$$

and the quasi-lump-sum taxes at $t = 1$ *be*

$$T_{1\theta} = \frac{u'(c_0^*)}{\beta u'(c_{1\theta}^*)} K_1^* + w_1^* \theta l_{1\theta} - c_{1\theta}^* \text{ for } t = 0, 1.$$

Then the optimal allocation A^**, savings* $k_1 = K_1^*$*, and prices* $r_t = r_t^*, w_t = w_t^*, t = 0, 1$ *are a competitive equilibrium with taxes.*

Proof. In the Appendix. ■

The capital tax rate and the quasi-lump-sum tax applied to the agent at $t = 1$ depend on whether or not he earns positive labor income at that date. If the agent's idiosyncratic productivity shock is $\theta = 0$, the agent is unable to earn positive labor income. Ex ante, hence, there are just two lifetime labor supply plans the agent may follow. Plan 1: produce zero income if unproductive and positive income if productive. Plan 2 (shirking): produce zero income if unproductive or productive. The strategy of the proof of this implementation theorem is to show that if the agent follows Plan 1, his utility is maximized by the exact consumption, labor supply, and savings choices that are prescribed by A^*, i.e., $c_0 = c_0^*$, $l_0 = l_0^*$, $k_1 = k_1^*$, and $c_{1\theta} = c_{1\theta}^*$, $l_{1\theta} = l_{1\theta}^*$ for both θ, hence giving him expected utility U^*; if he follows Plan 2, his optimal choices are $c_0 = c_0^*$, $l_0 = l_0^*$, $k_1 = k_1^*$, and $c_{11} = c_{10} = c_{10}^*$, $l_{11} = l_{10} = 0$, which, because the IC constraint (6) binds at A^*, also gives him expected utility U^*. Thus, the agent has no profitable deviation from the optimal allocation, inducing the joint deviation of over-saving at $t = 0$ and shirking at $t = 1$.

Further Properties and Comparison to Exclusive Contracts

The Inverse Euler Equation (10) implies that the expected marginal tax rate is zero:

$$\begin{aligned}
\mathbb{E}[(1 - \tau_1)] &= \pi_0 \frac{u'(c_0^*)}{r_1^* \beta u'(c_{10}^*)} + \pi_1 \frac{u'(c_0^*)}{r_1^* \beta u'(c_{11}^*)} \\
&= \frac{u'(c_0^*)}{r_1^* \beta} \mathbb{E}\left[\frac{1}{u'(c_1^*)}\right] \\
&= \frac{u'(c_0^*)}{r_1^* \beta} \frac{r_1^* \beta}{u'(c_0^*)} \\
&= 1.
\end{aligned}$$

The government therefore collects zero net revenue from capital taxes and all revenue needed to fund government expenditures G_t is collected via lump-sum and quasi-lump-sum taxes. The role for capital income taxes with state-contingent rates τ_{10}, τ_{11} is here purely to deter oversaving. It is immediate from (20) and (22) that $\tau_{10} > 0 > \tau_{11}$. Taxes τ_{10}, τ_{11} discourage savings not by decreasing the average return on savings (as the average tax rate is zero) but rather by introducing a negative correlation between the after-tax return on savings and the agent's marginal utility. In state $\theta = 0$, the agent's marginal utility of consumption, $u'(c_{10}^*)$, is high (as $c_{11}^* > c_{10}^*$ and u is strictly concave). Precisely in this state, however, the capital income tax τ_{10} is high and hence the after-tax return on savings is low. Conversely, in state $\theta = 1$ the capital income tax τ_{11} is low, so the after-tax return on savings is high, but the agent's marginal utility of consumption $u'(c_{11}^*)$ is low in this state because his consumption is high. Capital taxes are therefore designed to make savings a poor self-insurance tool, paying off more when the return is worth less to the agent, which discourages savings.

From the resource constraint at $t = 0$ we get that $T_0 = G_0$. Also we have

$$\begin{aligned} T_{11} - T_{10} &= u'(c_0^*)\frac{K_1^*}{\beta}\left(\frac{1}{u'(c_{11}^*)} - \frac{1}{u'(c_{10}^*)}\right) + w_1^* l_{11}^* - c_{11}^* + c_{10}^* \\ &> 0, \end{aligned}$$

where $w_1^* l_{11}^* - c_{11}^* + c_{10}^* > 0$ follows from (13). In this implementation, thus, agents who are poor at $t = 1$ (i.e., those with $\theta = 0$ and zero labor income) pay a high (positive) capital income tax and a low quasi-lump-sum tax at that date. Those with high (i.e., positive) labor income at $t = 1$ receive a subsidy to their capital income and pay a high quasi-lump-sum tax.

The structure of trade in this implementation with taxes is more decentralized and realistic than the one with lifetime exclusive contracts we discussed in the previous section. Here, agents trade without being monitored by anyone except the government for the purpose of collecting taxes. The strictly positive intertemporal wedge (11) implies that no mechanism decentralizing the private-information constrained optimum A^* can leave agents' savings decisions unmonitored. Using the government to monitor savings in order to collect capital income taxes is much closer to actual monitoring arrangements than what exclusive contracts discussed in the previous section require.

The structure of quasi-lump-sum taxes is also pretty realistic. In fact, instead of applying T_{11} to all agents with positive incomes and T_{10} to the agents with zero income at $t = 1$, we could equivalently apply a uniform lump sum tax $T_1 = T_{11}$ to all agents (regardless of income) and give a disability (on unemployment) benefit $B = T_{11} - T_{10}$ to those who earn no labor income at $t = 1$.[9]

Capital income taxes obtained here, however, are not very intuitive. In this decentralization, a complicated structure mapping labor income into tax rates is needed, whereas in practice capital tax rates are often flat. In particular, the subsidy to capital income (a negative tax rate τ_1) given to rich agents at $t = 1$ is hard to reconcile with actual capital income structures.

In the next section, we discuss another implementation mechanism in which the monitoring of savings is done by a court only in the event of the agent filing for bankruptcy. The set of bankruptcy rules needed to implement the optimum is very realistic, and the capital income tax rate is flat.

5. IMPLEMENTATION WITH UNSECURED CREDIT, BANKRUPTCY, AND SIMPLE TAXES

In this section, we study an implementation mechanism in which agents use unsecured credit and bankruptcy to obtain insurance against their productivity shocks, while taxes are used to fund government spending.

In this implementation, unsecured credit markets work as in Grochulski (2010). Competitive intermediaries trade with the agents using two financial instruments: unsecured, defaultable loans h, and riskless bonds b. Agents borrow using loans h and intermediaries borrow using bonds b.[10] Each agent faces a limit $\bar{h}$ on the amount of unsecured loans that he can take out with the intermediaries (it is his total credit limit with the whole industry).[11] The intermediaries hold loans h as assets and issue bonds b as their liabilities. Bonds b sell at $t = 0$ at the discount price q. The gross interest rate charged on the defaultable loans is R. Intermediaries diversify away the individual-specific risks by holding large (i.e., positive-measure) portfolios of defaultable loans. Intermediaries face a competitive market for unsecured loans. They decide whether to enter or not. If they do, they put in a credit offer on competitive terms. On these terms, the intermediaries expect a loan demand volume h^e with a fraction D^e of the loans going into default

[9] Golosov and Tsyvinski (2006) consider such a disability benefit.

[10] These bonds can be thought of as interest-bearing deposits.

[11] Grochulski (2010) discusses how the industry-wide credit limit $\bar{h}$ can be obtained as an outcome of strategic competition between financial intermediaries. As such, $\bar{h}$ is an object endogenous to the model. Critically, the competing intermediaries must be able to fully observe unsecured credit extended to the agent by other intermediaries. That is, in addition to limited debt discharge in bankruptcy, the model requires that a full credit report be available for each agent.

at $t = 1$ and with an expected principal recovery rate given default γ^e. In equilibrium, these expectations will be fulfilled.

The bankruptcy code works as follows. There is an eligibility criterion f and an asset exception level $\bar{e}$. The eligibility condition simply says that only agents with zero labor income can be granted discharge of their debts, h, in bankruptcy. The bankruptcy law also says that discharge can be granted only to agents who surrender their assets, $r_1 k_1 + b$. Assets up to the value $\bar{e}$ are exempt, i.e., are returned to the agent. Assets in excess of $\bar{e}$ are non-exempt, i.e., are distributed to the lenders whose unsecured loans h are being discharged. Those distributions are the basis for the lenders' expected principal recovery rate γ^e.

Intermediaries take as given prices q, R, and the credit limit $\bar{h}$. They form correct expectations of h^e, D^e, γ^e. Since they have zero external equity at $t = 0$, in order to balance assets and liabilities at $t = 0$, the intermediaries must satisfy at $t = 0$ the budget constraint $h^e = qb$. The expected profits are

$$\begin{aligned} \Pi_1 &= (1 - D^e)Rh^e + D^e\gamma^e h^e - b \\ &= \left((1 - D^e)R + D^e\gamma^e - q^{-1}\right) h^e, \end{aligned}$$

where the second line uses the budget constraint. Free entry into intermediation gives us immediately that

$$(1 - D^e)R + D^e\gamma^e = \frac{1}{q},$$

whenever $h^e > 0$. I.e., the expected rate of return on unsecured loans h must be equal to the intermediaries' cost of funding, $1/q$. The number of intermediaries operating in this competitive environment is indeterminate; it can be normalized to one.

Taxes are as follows. There is a proportional, flat-rate wealth tax τ at $t = 1$ and lump-sum taxes T_t at $t = 0, 1$.

With the asset markets, bankruptcy, and taxes as described above, the representative agent's problem is to choose non-negative consumption $c_0, c_{1\theta}$, labor $l_0, l_{1\theta}$, asset positions h, b, k_1, and a discrete bankruptcy filing plan $(d_0, d_1) \in \{0,1\} \times \{0,1\}$ so as to maximize (1) subject to

$$\begin{aligned} h &\leq \bar{h}, \\ c_0 + qb + k_1 &\leq w_0 l_0 + r_0 k_0 + h - T_0, \\ d_\theta &\leq f(w_1 \theta l_{1\theta}), \\ c_{1\theta} &\leq w_1 \theta l_{1\theta} + r_1 k_1 + b - (1 - d_\theta)Rh \\ &\quad - d_\theta \max\{r_1 k_1 + b - \bar{e}, 0\} - \tau(r_1 k_1 + b) - T_1, \end{aligned}$$

where the bankruptcy eligibility condition f is given by the indicator function of the number zero (agent is eligible only if $w_1\theta l_{1\theta} = 0$), and the asset exemption level $\bar{e}$ is a positive number.

Definition 3 *Given a set of taxes (τ, T_0, T_1) and bankruptcy laws $(f, \bar{e})$,* ***competitive equilibrium with taxes and bankruptcy*** *consists of an allocation*

$$\hat{A} = \{\hat{c}_0, (\hat{c}_{1\theta})_{\theta\in\{0,1\}}, \hat{l}_0, (\hat{l}_{1\theta})_{\theta\in\{0,1\}}, \hat{Y}_0, \hat{Y}_1, \hat{K}_1\},$$

the agents' loan and asset positions $\hat{h}$, $\hat{b}$, $\hat{k}_1$ and bankruptcy filing choices $(\hat{d}_\theta)_{\theta\in\{0,1\}}$, prices $q, R, \{r_t, w_t\}_{t=0,1}$, expectations D^e, γ^e, h^e, and a credit limit $\bar{h}$ such that: (a) the values $\hat{c}_0$, $(\hat{c}_{1\theta})_{\theta\in\{0,1\}}$, $\hat{l}_0$, $(\hat{l}_{1\theta})_{\theta\in\{0,1\}}$, $\hat{h}$, $\hat{b}$, $\hat{k}_1$, and $(\hat{d}_\theta)_{\theta\in\{0,1\}}$ solve the agent's utility maximization problem, (b) intermediaries break even: $\Pi_1 = 0$, (c) capital and labor are paid their respective marginal products

$$r_t = F_1(\hat{K}_t, \hat{Y}_t),\ w_t = F_2(\hat{K}_t, \hat{Y}_t) \text{ for } t = 0, 1,$$

(d) consumption, labor, and capital markets clear

$$\begin{aligned}
\hat{c}_0 + \hat{K}_1 + G_0 &= F(\hat{K}_0, \hat{Y}_0),\\
\hat{K}_1 &= \hat{k}_1,\\
\hat{Y}_0 &= \hat{l}_0,\\
\sum_{\theta\in\{0,1\}} \pi_\theta \hat{c}_{1\theta} + G_1 &= F(\hat{K}_1, \hat{Y}_1),\\
\hat{Y}_1 &= \sum_{\theta\in\{0,1\}} \pi_\theta \theta \hat{l}_{1\theta},
\end{aligned}$$

and (e) expectations are correct

$$\begin{aligned}
h^e &= \hat{h},\\
D^e &= \sum_{\theta\in\{0,1\}} \pi_\theta \hat{d}_\theta,\\
\gamma^e &= \frac{\max\{r_1\hat{k}_1 + \hat{b} - \bar{e}, 0\}}{\hat{h}} \text{ if } \hat{h} > 0.
\end{aligned}$$

Note that this definition implies that in equilibrium the government balances its budget every period. We will say that taxes (τ, T_0, T_1) and bankruptcy laws $(f, \bar{e})$ implement an optimum A^* if there exists a competitive equilibrium with taxes and bankruptcy such that the equilibrium allocation $\hat{A}$ coincides with the optimal allocation A^*.

Theorem 3 *Let A^* be optimal. Let taxes be (τ^*, T_0^*, T_1^*), where τ^* is given in (17) and*

$$T_0^* = G_0 \tag{23}$$

$$T_1^* = G_1 - \tau^* \left(r_1^* K_1^* + \pi_1^*(y_{11}^* - c_{11}^* + c_{10}^*)\right). \tag{24}$$

Let the bankruptcy code $(f, \bar{e})$ be given by

$$f(y_1) = \chi_{\{0\}}(y_1),$$

$$\bar{e} = r_1^* K_1^* + \pi_1(w_1^* l_{11}^* - c_{11}^* + c_{10}^*),$$

where χ is the indicator function. These taxes and bankruptcy rules implement A^*.

Proof. In the Appendix. ■

The proof of this theorem is constructive. It specifies a list of objects (prices, credit limits, expectations, agents' loan and asset choices) that along with the allocation $\hat{A} = A^*$ are a conjectured equilibrium. Then it checks that these conjectured choices in fact do satisfy the equilibrium conditions (a)–(e) of Definition 3.

In particular, the conjectured equilibrium prices are

$$r_t = r_t^*, \; w_t = w_t^* \text{ for } t = 0, 1, \tag{25}$$

$$q = 1/r_1^*, \tag{26}$$

$$R = r_1^*/\pi_1, \tag{27}$$

and the unsecured credit limit is

$$\bar{h} = \pi_1 \left(w_1^* l_{11}^* - c_{11}^* + c_{10}^*\right) / r_1^*. \tag{28}$$

The intermediaries' expectations are conjectured to be

$$h^e = \bar{h}, \; D^e = \pi_0, \; \gamma^e = 0,$$

the agent's loan and asset holding choices to be

$$\hat{h} = \bar{h}, \; \hat{b} = \bar{h} r_1^*, \; \hat{k}_1 = K_1^*,$$

and the bankruptcy filing plan to be

$$d_0 = 1, \; d_1 = 0. \tag{29}$$

This bankruptcy rules and the agents' equilibrium filing plan make the pricing of assets clear in this model. The bonds b, as riskless and receiving in bankruptcy the same treatment as physical capital holdings, earn the same return as capital. The return R on defaultable loans h contains a risk premium consistent with the agents' equilibrium default plan and the expected recovery rate. Agents default in and only in the low state $\theta = 0$ at date 1. With the fraction π_0 of agents receiving the low state realization, fraction π_0 of loans will default. The recovery rate γ^e is zero. Thus, in order to provide the required return of r_1, the face return R on the defaultable loan must satisfy $(1-\pi_0)R = r_1$, which implies the gross risk premium of $1/\pi_1$.

Confirming the consistency of the proposed choices and prices with agents' individual optimization is only the first step of the proof. In addition to the proposed bankruptcy filing plan (29), the proof considers three other cases, one for each alternative bankruptcy filing plan available to the agent. In particular, shirking in this model is complementary with maxing out the unsecured credit, producing no income at $t = 1$, and filing for bankruptcy in both states θ. This strategy is considered in Case II of the proof, where it is shown that, similar to the NDPF decentralization, a shirker can do as well as the non-shirker, but the IC constraint implies that he cannot do better.

Interestingly, competition between intermediaries providing unsecured credit sets the agent's credit limit, given in (28). This limit is determined at the maximal level with which repayment of defaultable loans h remains incentive compatible (conditional on not having over-saved). Indeed, any intermediary offering the agent an additional unsecured loan in excess of the equilibrium limit $\bar{h}$ would understand that with that loan the agent's utility would be maximized by producing zero income, filing for bankruptcy, and defaulting in both states θ. In effect, the intermediary would never see this additional loan repaid, hence no credit in excess of $\bar{h}$ is offered in equilibrium. Conversely, the credit limit the agent faces cannot be lower than $\bar{h}$. If it were, the additional loan offered by the intermediary would not trigger default in state $\theta = 1$, and so, with free entry, an intermediary willing to make this credit offer (perhaps pricing it marginally higher than R) would be found.

Differences and Similarities between the Implementations

In thinking about the possible institutional arrangements that can support the optimal allocation A^*, the unsecured credit model we study in this section is an alternative to the NDPF tax model we discussed earlier. This alternative is quite realistic. In the United States, eligibility for discharge of personal debts is income-tested. Agents with labor incomes above a certain threshold are not eligible for discharge in the so-called Chapter 7 liquidation procedure.[12] Our eligibility

[12] Instead, they must go through a repayment plan under the so-called Chapter 13 debt reorganization plan. See U.S. Code, Title 11, Chapter 7, Subchapter I, § 707—Dismissal of a case or conversion to a case under chapter 11 or 13. Available at www.law.cornell.edu/uscode/text/11/707.

condition f corresponds to this feature of the law.[13] Likewise, the Chapter 7 discharge procedure allows for a limited asset exemption in bankruptcy, similar to our allowance $\bar{e}$.

In the implementation with unsecured credit and bankruptcy, the monitoring of agents' savings is done by the bankruptcy court upon the agent's filing for bankruptcy and requesting discharge of his unsecured debts. For the proof of Theorem 3 to work, it is crucial that agents cannot borrow and later obtain discharge while hiding assets from the court. This assumption is quite reasonable, as hiding assets from the court is probably hard in practice, at least in the United States.

By using a different set of restrictions than those used in the NDPF tax model, the bankruptcy mechanism shows a different way to implement the distortions embedded in the optimal allocation A^*, which also helps us to understand those distortions better. The differences between the two mechanisms are as follows.

In the NDPF mechanism, agents trade a single asset, capital. In the bankruptcy mechanism, agents trade capital, unsecured loans, and bonds. The bankruptcy rules $(f, \bar{e})$ and the credit limit $\bar{h}$ support trade in unsecured, defaultable loans in equilibrium. Because agents' bankruptcy filing decisions are state-contingent, the payoff of a portfolio consisting of a loan and bonds is tailored to each agent's individual realization of uncertainty. The asset span generated by these assets is therefore larger than that provided by the single asset used in the NDPF implementation. The agents use this extended asset span to obtain insurance. The government does not use fiscal policy instruments (the quasi-lump-sum taxes and marginal capital income tax rates in the previous section) to provide insurance via ex post redistribution. In fact, in the simple tax structure that funds government expenditures in this section, the present value of lifetime taxes that agents pay does not depend on the realization of uncertainty, which is not the case in the NDPF tax system.[14]

The model with unsecured credit and bankruptcy falls in between the exclusive contracts model and the NDPF tax model in the

[13] Condition f is extreme in restricting eligibility to those with zero income, but this is not essential. In a more general version of the Mirrleesian environment, the low realization of θ could be strictly positive, in which case the low-income agents receiving bankruptcy discharge would make strictly positive income. In this case, the test f would not require zero income for discharge eligibility.

[14] The implementation of optimal ex post transfers with a combination of quasi-lumps-sum taxes and marginal capital tax rates that depend on observed labor income is very similar to the capital- and labor-income tax system studied in Kocherlakota (2005). In Golosov and Tsyvinski (2006), these transfers are implemented via a tax-funded disability benefit. In Grochulski and Kocherlakota (2010), the optimal transfers are implemented via state-contingent marginal capital tax rates and a tax-funded, state-contingent social security benefit.

following sense. In the exclusive contracts model, private firms provide insurance and monitor agents. In the NDPF model, both these functions are taken over by the government (its fiscal agent). In the unsecured credit and bankruptcy model, the government (which means the court system in this case) does the monitoring but the private sector (i.e., the intermediaries extending unsecured credit) provides insurance. Alternatively, we can think of agents as self-insuring using the extended asset span supported by the intermediaries' extension of defaultable credit, which is made possible by the court's monitoring of savings in bankruptcy.

The NDPF tax system studied in the previous section is designed to overcome the problem of agents' joint deviations. The simple tax system we study in this section is not. Similar to Grochulski (2010), agents find joint deviations unprofitable because of the bankruptcy rules, not because of taxes. In particular, the joint deviation in which an agent saves at $t = 0$ more than what is socially optimal and works at $t = 1$ less than the socially optimal amount is not profitable because the bankruptcy rules $(f, \bar{e})$ make it impossible for any agent to keep at $t = 1$ both the transfer that optimally goes from the productive to the unproductive agents and the return on any savings exceeding the optimal amount. To obtain the transfer, the agent must file for bankruptcy, because bankruptcy debt discharge combined with an asset exemption is the means through which this transfer is provided in this model. But in bankruptcy, the agent must give up assets in excess of the exemption $\bar{e}$, and this exemption is set precisely at the socially optimal amount of savings. Thus, the agent cannot benefit from the implicit insurance payment and the return on over-saving simultaneously.

The proportional wealth tax τ^* has a role in discouraging over-saving in that, in contrast to Grochulski (2010), the bankruptcy exemption caps do not bind in the utility maximization problem of an agent who does not shirk (see Case I in the proof of Theorem 3). But this role is not essential. It is straightforward to follow the steps in the proof of Theorem 3 to check that the optimum can be implemented with a simple, proportional capital tax with the marginal rate τ given by any number between zero and the "natural" value τ^* used in Theorem 3.

To see this, note that any tax rate $\tau < \tau^*$ is too low to close the intertemporal wedge of a non-shirker. One might suspect that agents would find it profitable to over-save if the wealth tax rate they face is $\tau < \tau^*$. This, however, is not the case because the exemption cap constraint would become binding. In fact, the value τ^* used in Theorem 3 is already too low to close the intertemporal wedge of a shirker, and the exemption cap constraint binds in the shirker's problem (Case II in the proof of Theorem 3).

It is true, however, that if the marginal wealth tax rate is sufficiently negative, agents will over-save. The threshold value $\tilde{\tau}$ at which this happens satisfies $1-\tilde{\tau} = u'(c_0^*)/\pi_1 r_1^* \beta u'(c_{11}^*) > 1$, i.e., $\tilde{\tau} < 0$. Clearly, if the subsidy to savings is sufficiently large, agents will over-save because they get to keep the after-tax return on savings at least in the state $\theta = 1$ in which they do not file for bankruptcy. But because the threshold number $\tilde{\tau}$ is strictly negative, this will not happen for any non-negative marginal wealth tax rate τ.

If the marginal rate τ exceeds τ^*, however, agents do find it optimal to deviate from the optimal allocation. This is because a tax rate τ higher than τ^* suppresses savings below $\bar{e}$, and with savings smaller than $\bar{e}$ the bankruptcy exemption cap does not bind. The value τ^* given in (17), therefore, is the upper end of the interval containing the marginal wealth tax rates consistent with implementation of the optimum in the tax/bankruptcy mechanism studied in this section. In addition to the natural interpretation of closing the intertemporal wedge (11), τ^* maximizes the amount of revenue raised from wealth taxes, and, thus, minimizes the size of the lump sum tax levied at $t = 1$, for a given level of government spending G_1.

6. CONCLUSION

In this article, we use a simple Mirrleesian model to discuss the implications of private information for optimal allocations and their decentralizations. We focus on the intertemporal wedge and three ways to implement it: exclusive contracts, capital income taxes, and a set of bankruptcy rules.

The multiplicity of possible implementations makes normative analysis challenging. The model pins down optimal wedges but does not determine the institutions that should be used to support them in equilibrium. For example, the bankruptcy model shows that when private credit markets provide insurance, a simple, non-contingent capital income tax τ^* can be optimal in the Mirrleesian economy. Thus, the fact that there is private information in the economy does not imply that capital tax rates should be state-contingent, as the NDPF literature suggests. Likewise, bankruptcy is not essential if firms can sign agents to exclusive, comprehensive lifetime utility contracts.

Despite this difficulty, further study of implementations is valuable and needed for the following three reasons. First, from a purely theoretical perspective, implementation exercises show that some of the many constraints imposed in the social planning problem are inessential. For example, the personal bankruptcy implementation shows that saving must be monitored and prohibited/taxed/confiscated not always but only in the event of the agent claiming a social insurance payout (discharge of unsecured debts in this implementation). Thus, saving is detrimental to incentives only to the extent to which it provides self-insurance.

Second, the Mirrleesian environment admits multiple implementations perhaps because it is not rich enough to determine both the optimum and the institutions needed to implement it. In practice, in addition to private information about individual productivity shocks, other frictions may be affecting the optimal allocation. In a richer model, implementation may be pinned down much more closely.[15] The observation that one implementation coming out of the simple Mirrleesian model fits better with real-life institutions than another can inform us about the direction in which the Mirrleesian model should be enriched to better capture reality.

Finally, the very purpose of studying optimal allocations with their wedges is to provide lessons for the design of better policies and improvements in institutions that affect the actual economic outcomes. That the shadow interest rate of the agent should be strictly lower than the rate of return on capital is a robust implication of private information in the Mirrleesian environment. But absent implementation, this implication does not tell us anything of *practical* value. Therefore, despite the difficulties associated with matching implementation results to real-life institutions, further study of implementations, along with optimal allocations, is needed.

APPENDIX

Proof of Theorem 1

If $\bar{U} > U^*$, then A^* could not have been optimal. If $\bar{U} < U^*$, then there exists a contract $\tilde{\psi}$ and a capital level $\tilde{K}_1^{\mathrm{f}}$ that deliver to the agent utility $\tilde{U} \geq \bar{U}$ while making a strictly positive profit for the firm, $\Pi_0(\tilde{U}) > 0$, which implies that $(\hat{\psi}, \hat{K}_1^{\mathrm{f}})$ cannot be an equilibrium. QED

[15] Cole and Kocherlakota (2001) study a model with two-dimensional private information: hidden income and hidden storage. That environment pins down both the optimal allocation and its implementation, which in that case happens to coincide with pure self-insurance.

Proof of Theorem 2

We need to show that if taxes are as specified in the theorem's statement, then the choices for labor, consumption, and savings prescribed by the optimal allocation A^* solve the agent's utility maximization problem.

If $\theta = 0$, the agent cannot produce positive income at $t = 1$, so the capital income tax rate that applies to him is τ_{10} and the quasi-lump-sum tax he pays is T_{10}. If $\theta = 1$, the agent has a choice. He can produce positive income at $t = 1$, which means that the capital income tax rate that applies to him is τ_{11} and the quasi-lump-sum tax he pays is T_{11}, or he can choose to produce zero income at $t = 1$, so the capital income tax rate that applies to him is τ_{10} and the quasi-lump-sum tax he pays is T_{10}.

Consider first the agent's plan, chosen at of $t = 0$, to produce positive income at $t = 1$ if $\theta = 1$. We will show that the optimal allocation solves the agent's problem conditional on positive income produced in state $\theta = 1$. With positive income in this state, using the expressions for the proposed taxes $\tau_{1\theta}$ and $T_{1\theta}$ and simplifying, the budget constraints the agent faces are

$$\begin{aligned} c_0 - c_0^* + k_1 - K_1^* &\leq w_0^* \left(l_0 - l_0^*\right), \\ c_{10} - c_{10}^* &\leq \frac{u'(c_0^*)}{\beta u'(c_{10}^*)} \left(k_1 - K_1^*\right), \\ c_{11} - c_{11}^* &\leq w_1^* \left(l_{11} - l_{11}^*\right) + \frac{u'(c_0^*)}{\beta u'(c_{11}^*)} \left(k_1 - K_1^*\right). \end{aligned}$$

As we see, the budget constraints are expressed in terms of deviations from the proposed equilibrium choices, which shows that the proposed choices are feasible for the agent. The objective the agent maximizes is his ex ante expected utility given in (1). Using the budget constraints to eliminate consumption from the objective, we boil down the problem to two intratemporal choices, l_0 and l_{11}, and one intertemporal choice, that of k_1. The problem is convex, the first-order (FO) conditions are sufficient. Taking the FO condition with respect to l_0 and l_{11} and evaluating at l_0^* and l_{11}^*, we get (8) and (9), i.e., the FO conditions are satisfied by l_0^* and l_{11}^*. At the same time, taking the FO condition with respect to k_1 and evaluating it at K_1^*, we get the after-tax Euler equation (21). This shows that $k_1 = K_1^*$ satisfies the FO condition (precisely because the tax rates $\tau_{1\theta}$ were chosen so that the Euler equation [(21)] is satisfied). Together, these three conditions are sufficient for optimality. The optimal allocation thus solves the agent's problem, conditional on producing positive income when productive at date 1.

Consider now the other option available to the agent ex ante: the plan to produce zero income at $t = 1$ in state $\theta = 1$. If the agent plans to produce zero income in this state, he produces zero income in both states at $t = 1$. Thus, his budget constraints are

$$\begin{aligned} c_0 - c_0^* + k_1 - K_1^* &\leq w_0^* \left(l_0 - l_0^* \right), \\ c_1 - c_{10}^* &\leq \frac{u'(c_0^*)}{\beta u'(c_{10}^*)} \left(k_1 - K_1^* \right), \end{aligned}$$

where c_1 is consumption at $t = 1$ (the same in both states θ as, again, the agent produces the same income and faces the same taxes in both states at $t = 1$). The objective the agent maximizes is

$$u(c_0) - v(l_0) + \beta u(c_1).$$

We claim that the following choices solve this utility maximization problem: $c_0 = c_0^*, l_0 = l_0^*, k_1 = K_1^*$, and $c_1 = c_{10}^*$. Indeed, substituting consumption out of the objective using the budget constraint, we get a convex maximization problem in l_0 and k_1. Taking the FO condition with respect to l_0 and evaluating at l_0^*, we get again (8). Taking the FO condition with respect to k_1 and evaluating it at K_1^*, we get the shirker's after-tax Euler equation (19). This shows that $l_0 = l_0^*$ and $k_1 = K_1^*$ satisfy the FO conditions that are sufficient for optimality. Conditional on shirking, the agent obtains ex ante utility of

$$u(c_0^*) - v(l_0^*) + \beta u(c_{10}^*).$$

In order to compare the agent's value of non-shirking and shirking, we invoke the binding IC constraint (7). This equation implies that the agent does not gain ex ante utility by shirking. Therefore, the optimal allocation is consistent with individual optimization, under the specified capital income and semi-lump-sum taxes.

Finally, we note that because the optimal allocation satisfies the resource constraints, markets clear and the government raises the requisite amount of revenue, hence we have an equilibrium. QED

Proof of Theorem 3

In this proof, we must check that the optimal allocation A^* along with the proposed equilibrium quantities given in (25) though (29) satisfy the equilibrium conditions of Definition 3.

Because the consistency and market clearing conditions (b)–(e) are expressed simply as algebraic equalities, direct substitution of the proposed equilibrium values into these equalities confirms that conditions (b)–(e) are satisfied. Condition (a), however, is a maximization condition. The remainder of the proof is devoted to showing that 1) the proposed equilibrium choices (25)–(29) belong to the representative agent's budget set, and 2) the agent cannot benefit by deviating from the proposed equilibrium behavior.

First we check that the consumption, unsecured loan, assets, and bankruptcy choices are in the budget set. Substituting the proposed equilibrium values and the tax T_0^* from (23) into the budget constraint at date 0, we get

$$c_0^* + \bar{h}r_1^*/r_1^* + K_1^* = w_0^* l_0^* + r_0^* K_0 + \bar{h} - G_0.$$

Using (28), we see that this equation holds true because $w_0^* l_0^* + r_0^* K_0 = F(K_0, Y_0^*)$ and A^* satisfies (2). Substituting taxes (17) and (24) and the proposed equilibrium choices into the date-1 budget constraints yields

$$\begin{aligned} c_{10}^* &= w_1^* 0 + r_1^* K_1^* + \bar{h}r_1^* - 0\bar{h}r_1^*/\pi_1 \\ &\quad -1 \max\{r_1^* K_1^* + \bar{h}r_1^* - \bar{e}, 0\} - \tau^*(r_1^* K_1^* + \bar{h}r_1^*) - T_1^*, \\ c_{11}^* &= w_1^* l_{11}^* + r_1^* K_1^* + \bar{h}r_1^* - \bar{h}r_1^*/\pi_1 \\ &\quad -0 \max\{r_1^* K_1^* + \bar{h}r_1^* - \bar{e}, 0\} - \tau^*(r_1^* K_1^* + \bar{h}r_1^*) - T_1^*. \end{aligned}$$

Using $T_1^* = G_1 - \tau^* \left(r_1^* K_1^* + r_1^* \bar{h}\right)$ and $\bar{e} = r_1^* K_1^* + r_1^* \bar{h}$, we can simplify these to

$$\begin{aligned} c_{10}^* &= r_1^* K_1^* + \bar{h}r_1^* - G_1, \\ c_{11}^* &= w_1^* l_{11}^* + r_1^* K_1^* + \bar{h}r_1^*(1 - 1/\pi_1) - G_1. \end{aligned}$$

Since $F(K_1^*, Y_1^*) = w_1^*(\pi_1 l_{11}^* + \pi_0 0) + r_1^* K_1^*$, resource feasibility of A^* implies that $G_1 = w_1^* \pi_1 l_{11}^* + r_1^* K_1^* - \pi_0 c_{10}^* - \pi_1 c_{11}^*$. Substituting this into the above two equations and simplifying terms, we get

$$c_{10}^* = \bar{h}r_1^* - w_1^* \pi_1 l_{11}^* + \pi_0 c_{10}^* + \pi_1 c_{11}^*, \tag{30}$$

$$c_{11}^* = \pi_0 w_1^* l_{11}^* + \bar{h}r_1^* \pi_0/\pi_1 + \pi_0 c_{10}^* + \pi_1 c_{11}^*. \tag{31}$$

Using (28) we check that the right-hand side of (30) is indeed c_{10}^* and the right-hand side of (31) is indeed c_{11}^*.

We also need to show that $\hat{b}$ and $\hat{h}$ are non-negative. They both are non-negative if and only if $w_1^* l_{11}^* - c_{11}^* + c_{10}^* \geq 0$. Inequality (13) shows that this inequality holds strictly, so $\hat{b}$ and $\hat{h}$ are in fact strictly positive. This confirms budget-feasibility. Next, we need to show that the agent cannot do better by deviating from the proposed choices.

The agent needs to make a discrete choice of a bankruptcy plan. There are four possible bankruptcy plans the agent can choose among $(d_0, d_1) \in \{(1, 0), (1, 1), (0, 0), (0, 1)\}$. The proposed equilibrium plan is $(d_0, d_1) = (1, 0)$. We will go through all four cases to show that the proposed equilibrium plan is the best for the agent. Also, we will show that conditional on $(d_0, d_1) = (1, 0)$, the rest of the proposed equilibrium behavior maximizes the utility of the representative agent.

Case I. $(d_0, d_1) = (1, 0)$, i.e., the agent uses bankruptcy when $\theta = 0$ and does not when $\theta = 1$.

Conditional on this bankruptcy plan, the agent's budget constraints are as follows:

$$\begin{aligned} h &\leq \bar{h}, \\ c_0 + qb + k_1 &\leq w_0 l_0 + r_0 k_0 + h - T_0^*, \\ c_{10} &\leq w_1 0 + r_1 k_1 + b - \max\{r_1 k_1 + b - \bar{e}, 0\} \\ &\quad -\tau^*(r_1 k_1 + b) - T_1^*, \qquad (32) \\ c_{11} &\leq w_1 l_{11} + r_1 k_1 + b - Rh - \tau^*(r_1 k_1 + b) - T_1^*, \end{aligned}$$

with prices w_t, r_t, q, and R as specified in (25)–(27). We will relax this problem by dropping the non-positive term $-\max\{r_1 k_1 + b - \bar{e}, 0\}$ from the right-hand side of (32). That is, we replace (32) with

$$c_{10} \leq w_1 0 + r_1 k_1 + b - \tau^*(r_1 k_1 + b) - T_1^*.$$

We now show that the proposed equilibrium behavior solves the relaxed problem. Then, we will check that this solution is also feasible in the unrelaxed problem.

The relaxed problem is a concave maximization problem. The FO conditions along with the budget constraints at equality are necessary and sufficient. The FO conditions are as follows:

$$\begin{aligned} v'(l_0) &= u'(c_0) w_0, \\ v'(l_{11}) &= u'(c_{11}) w_1, \\ u'(c_0) &\geq R\beta\pi_1 u'(c_{11}) \text{ with equality if } h < \bar{h}, \\ u'(c_0) &= q^{-1}\beta(1-\tau^*)\mathbb{E}[u'(c_1)], \\ u'(c_0) &= r_1\beta(1-\tau^*)\mathbb{E}[u'(c_1)]. \end{aligned}$$

With prices as in (25)–(27), simple substitution of the proposed equilibrium values for c, l, and h verifies that these values do solve this problem. In particular, the intratemporal conditions for labor follow from (8) and (9). The intertemporal condition with respect to h follows from the left inequality in (12). Using $q^{-1} = r_1 = r_1^*$ and the expression for τ^* in (17), we get that the intertemporal conditions with respect to b and k_1 are satisfied as well. (We note that only the FO condition with respect to h is binding here.)

We now note that the solution to the relaxed problem is also feasible in the unrelaxed problem because $-\max\{r_1^* K_1 + \hat{b} - \bar{e}, 0\} = 0$. This verifies that the proposed equilibrium behavior solves the agent's problem conditional on $(d_0, d_1) = (1, 0)$. The utility the agent obtains in this case is thus

$$u(c_0^*) - v(l_0^*) + \beta\pi_1(u(c_{11}^*) - v(l_{11}^*)) + \beta\pi_0 u(c_{10}^*).$$

Case II. $(d_0, d_1) = (1, 1)$, i.e., the agent goes bankrupt in both individual states θ.

In order to be eligible to go bankrupt in state $\theta = 1$, the agent must choose $l_{11} = 0$. This means that $l_{11} = l_{10}$, i.e., the agent behaves identically in both states $\theta = 0, 1$. The agent thus chooses c_0, l_0, h, b, k_1, and c_1 so as to maximize

$$u(c_0) - v(l_0) + \beta[u(c_1) - v(0)] \tag{33}$$

subject to

$$\begin{aligned} h &\leq \bar{h}, \\ c_0 + qb + k_1 &\leq w_0 l_0 + r_0 k_0 + h - T_0^*, \\ c_1 &\leq r_1 k_1 + b - \max\{r_1 k_1 + b - \bar{e}, 0\} - \tau^*(r_1 k_1 + b) - T_1^*, \end{aligned}$$

with prices as in (25)-(27). We will show that $c_0 = c_0^*$, $l_0 = l_0^*$, $h = \hat{h}$, $b = \hat{b}$, $k_1 = K_1^*$, and $c_1 = c_{10}^*$ solve this problem. Since under the bankruptcy filing plan considered in this case the unsecured loan is repaid in neither state θ, the agent, clearly, chooses $h = \bar{h}$. We can thus rewrite the budget constraints as

$$\begin{aligned} c_0 + qb + k_1 &\leq w_0 l_0 + r_0 k_0 + \bar{h} - T_0^*, \\ c_1 &\leq \min\{r_1 k_1 + b, \bar{e}\} - \tau^*(r_1 k_1 + b) - T_1^*. \end{aligned}$$

Due to the confiscation of nonexempt assets in bankruptcy, it will never be optimal for the agent who files for bankruptcy with probability one to choose $r_1 k_1 + b$ larger than $\bar{e}$. We can therefore rewrite the above budget constraints as

$$\begin{aligned} c_0 + qb + k_1 &\leq w_0 l_0 + r_0 k_0 + \bar{h} - T_0^*, \\ c_1 &\leq (1 - \tau^*)(r_1 k_1 + b) - T_1^*, \\ r_1 k_1 + b &\leq \bar{e}. \end{aligned}$$

The problem of maximization of (33) subject to these budget constraints is a convex problem. The set of FO necessary and sufficient conditions for the maximum consists of the budget equations and

$$\begin{aligned} v'(l_0) &= w_0 u'(c_0), \\ u'(c_0) &\leq (1 - \tau^*) q^{-1} \beta u'(c_1) \text{ with equality if } r_1 k_1 + b < \bar{e}, \\ u'(c_0) &\leq (1 - \tau^*) r_1 \beta u'(c_1) \text{ with equality if } r_1 k_1 + b < \bar{e}. \end{aligned}$$

That values c_0^* and l_0^* satisfy the first of these conditions follows from (8). Using (17), substituting $c_0 = c_0^*$, $c_1 = c_{10}^*$, $q^{-1} = r_1 = r_1^*$, and cancelling out terms, the second and third conditions reduce to a single condition

$$1 \leq \frac{u'(c_{10}^*)}{\mathbb{E}[u'(c_1^*)]} \text{ with equality if } r_1 k_1 + b < \bar{e},$$

which is satisfied because $\mathbb{E}[u'(c_1^*)] < u'(c_{10}^*)$ and $r_1^* K_1^* + \hat{b} = \bar{e}$. Thus, the values we proposed as a solution to this problem in fact solve it. (Note that the constraint $r_1 k_1 + b \leq \bar{e}$ binds in this problem.) In sum, the value that the agent can obtain using the bankruptcy strategy $d_0 = d_1 = 1$ equals

$$u(c_0^*) - v(l_0^*) + \beta u(c_{10}^*).$$

Because the IC constraint holds at the optimum A^* with equality, this amount of utility is exactly equal to what the agent obtains in Case I. Thus, the proposed equilibrium behavior is weakly better for the agent than the strategy of going bankrupt in both states.

Case III. $(d_0, d_1) = (0, 0)$, i.e., the agent never uses the bankruptcy option.

Because the agent never goes bankrupt, any unsecured loan h he takes out at $t = 0$ will be repaid at $t = 1$ with probability one. Thus, as long as the agent's total savings are non-zero, $R > (1 - \tau^*) r_1$ implies that the agent will choose $h = 0$ simply because reducing savings is a cheaper form of borrowing than the unsecured loan h. Also, it is without loss of generality here to take $b = 0$ because $q = 1/r_1$. Thus, the agent's budget constraints reduce to

$$\begin{aligned} c_0 + k_1 &\leq w_0 l_0 + r_0 k_0 - T_0^*, \\ c_{1\theta} &\leq w_1 \theta l_{1\theta} + (1 - \tau^*) r_1 k_1 - T_1^*, \end{aligned} \tag{34}$$

with prices as in (25)–(27). Let $(\tilde{c}_0, \tilde{k}_1, \tilde{l}_0, \tilde{c}_{1\theta}, \tilde{l}_{1\theta})$ be a solution to this problem. We need to show that the value the agent attains in this problem is less than the value delivered by the optimal allocation A^*. Two cases are possible: $r_1^* \tilde{k}_1$ is weakly smaller than $\bar{e}$, or it is larger than $\bar{e}$.

If $r_1^* \tilde{k}_1 \leq \bar{e}$, then the choices $(\tilde{c}_0, \tilde{k}_1, \tilde{l}_0, \tilde{c}_{1\theta}, \tilde{l}_{1\theta})$ along with $b = h = 0$ are feasible in the agent's problem considered in Case I, where the agent uses the bankruptcy plan $(d_0, d_1) = (1, 0)$. In particular, $r_1^* \tilde{k}_1 \leq \bar{e}$ ensures that the agent does not surrender any assets in the "empty" bankruptcy in state $\theta = 1$ (where discharge is zero because $h = 0$). Because these choices are feasible in the problem considered in Case I, they cannot deliver to the agent strictly more utility than a solution to that problem, and we saw in Case I that allocation A^* did solve this problem. Thus, the value attained by choices $(\tilde{c}_0, \tilde{k}_1, \tilde{l}_0, \tilde{c}_{1\theta}, \tilde{l}_{1\theta})$ cannot be larger than the value delivered by A^*.

If $r_1^* \tilde{k}_1 > \bar{e}$, then the agent pays more in taxes here than in Case I. This is because in Case I the agent carries less wealth into period 1, $r_1^* K_1^* + \hat{b} = \bar{e}$, and the tax rate on wealth τ^* is strictly positive. Thus, extending choices $(\tilde{c}_0, \tilde{k}_1, \tilde{l}_0, \tilde{c}_{1\theta}, \tilde{l}_{1\theta})$ to the whole population results with an allocation, to be denoted by $\tilde{A}$, that generates enough taxes to satisfy (2)–(5). Allocation $\tilde{A}$ is also incentive compatible, as the pair $(c_1, l_1) = (\tilde{c}_{10}, 0)$ satisfies the budget constraint (34) with $\theta = 1$ but the pair $(c_1, l_1) = (\tilde{c}_{11}, \tilde{l}_{11})$ is chosen instead when $\theta = 1$. Allocation $\tilde{A}$ is therefore feasible in the social planning problem (SPP) defined in Section 2. As a feasible choice in SPP, $\tilde{A}$ cannot generate a higher value than a solution to SPP, and A^* solves SPP.

Case IV. $(d_0, d_1) = (0, 1)$, i.e., the agent goes bankrupt when $\theta = 1$ and does not when $\theta = 0$.

In order to be eligible for bankruptcy in state $\theta = 1$, the agent must choose $l_{11} = 0$. He obviously also chooses $l_{10} = 0$.

At the solution to this problem, the agent chooses either $h = 0$ or $h > 0$. Suppose first that the agent chooses $h = 0$. In this case, it is weakly better for the agent to not file for bankruptcy at all because the benefit of discharging $h = 0$ in bankruptcy is zero, and the cost of subjecting his savings to the bankruptcy exemption cap $\bar{e}$ is nonnegative. Thus, with $h = 0$, the bankruptcy strategy $(d_0, d_1) = (0, 1)$ is weakly dominated by the strategy $(d_0, d_1) = (0, 0)$, which was shown in Case III to be dominated by the proposed equilibrium strategy $(d_0, d_1) = (1, 0)$.

If $h > 0$, then we will show that the strategy $(d_0, d_1) = (0, 1)$ is dominated by the strategy $(d_0, d_1) = (1, 1)$. Indeed, the budget constraints the agent faces conditional on $(d_0, d_1) = (0, 1)$ are

$$
\begin{aligned}
h &\leq \bar{h}, \\
c_0 + qb + k_1 &\leq w_0 l_0 + r_0 k_0 + h - T_0^*, \\
c_{10} &\leq r_1 k_1 + b - Rh - \tau^*(r_1 k_1 + b) - T_1^*, \\
c_{11} &\leq r_1 k_1 + b - \max\{r_1 k_1 + b - \bar{e}, 0\} - \tau^*(r_1 k_1 + b) - T_1^*,
\end{aligned}
$$

with prices w_t, r_t, q, and R as specified in (25)–(27). We check that at the solution to this problem with $h > 0$, the agent does not save more than $\bar{e}$, i.e., his choices satisfy $r_1 k_1 + b \leq \bar{e}$. Suppose he saves exactly $\bar{e}$, i.e., $r_1 k_1 + b = \bar{e}$, and considers increasing his savings by investing $\varepsilon > 0$ more in capital. (The argument is the same if the agent considers increasing his bond holdings b.) The marginal payoff this extra investment gives at date 1 is $(1 - \tau^*) r_1 \varepsilon < r_1 \varepsilon$ in state $\theta = 0$ and zero in state $\theta = 1$ because savings in excess of $\bar{e}$ are confiscated in bankruptcy (for which the agent files in state $\theta = 1$). The alternative strategy of decreasing $h > 0$ by ε has the same cost at $t = 0$ and pays off $R\varepsilon > r_1 \varepsilon$ in state $\theta = 0$ and zero in state $\theta = 1$. Thus, the agent would prefer to reduce h rather than to increase his savings above $\bar{e}$. It is thus not optimal for the agent to have savings $r_1 k_1 + b$ higher than $\bar{e}$.

We now can show that, keeping all other choices unchanged, the agent can increase his consumption c_{10} by filing for bankruptcy in state

$\theta = 0$. This is simply because

$$\begin{aligned} r_1k_1 + b - Rh - \tau^*(r_1k_1 + b) - T_1^* &< r_1k_1 + b - \tau^*(r_1k_1 + b) - T_1^* \\ &= r_1k_1 + b - \max\{r_1k_1 + b - \bar{e}, 0\} \\ &\quad -\tau^*(r_1k_1 + b) - T_1^*, \end{aligned}$$

where the equality follows from $r_1k_1 + b \leq \bar{e}$ and the strict inequality follows from $h > 0$. (Intuitively, given that the agent chooses at date 0 savings that do not trigger asset confiscation in bankruptcy he files for in state $\theta = 1$, there is no reason to repay the unsecured loan $h > 0$ in state $\theta = 0$, as the agent is eligible for bankruptcy in state $\theta = 0$ because with $\theta = 0$ his labor income $\theta w_1 l_{10}$ is zero anyway.) Thus, the best course of action under the bankruptcy strategy $(d_0, d_1) = (0, 1)$ is improved upon by simply changing the bankruptcy plan to $(d_0, d_1) = (1, 1)$. But it was shown in Case II that no course of action associated with the bankruptcy plan $(d_0, d_1) = (1, 1)$ can be superior to the proposed equilibrium strategy that implements the optimal allocation A^* with with $(d_0, d_1) = (1, 0)$. QED

REFERENCES

Atkeson, Andrew, and Robert E. Lucas, Jr. 1992. "On Efficient Distribution with Private Information." *Review of Economic Studies* 59 (July): 427–53.

Cole, Harold L., and Narayana R. Kocherlakota. 2001. "Optimal Allocations with Hidden Income and Hidden Storage." *Review of Economic Studies* 68 (3): 523–42.

Diamond, Peter A., and James A. Mirrlees. 1978. "A Model of Social Insurance with Variable Retirement." *Journal of Public Economics* 10 (3): 295–336.

Farhi, Emmanuel, and Iván Werning. 2013. "Insurance And Taxation Over The Life Cycle." *Review of Economic Studies* 80 (2): 596–635.

Golosov, Mikhail, Narayana Kocherlakota, and Aleh Tsyvinski. 2003. "Optimal Indirect and Capital Taxation." *Review of Economic Studies* 70 (July): 569–87.

Golosov, Mikhail, and Aleh Tsyvinski. 2006. "Designing Optimal Disability Insurance: A Case for Asset Testing." *Journal of Political Economy* 114 (April): 257–79.

Golosov, Mikhail, and Aleh Tsyvinski. 2007. "Optimal Taxation with Endogenous Insurance Markets." *Quarterly Journal of Economics* 122 (2): 487–34.

Grochulski, Borys. 2010. "Optimal Personal Bankruptcy Design Under Moral Hazard." *Review of Economic Dynamics* 13 (April): 350–78.

Grochulski, Borys, and Narayana Kocherlakota. 2010. "Nonseparable Preferences and Optimal Social Security Systems." *Journal of Economic Theory* 145 (November): 2,055–77.

Huggett, Mark, and Juan Carlos Parra. 2010. "How Well Does the U.S. Social Insurance System Provide Social Insurance?" *Journal of Political Economy* 118 (1): 76–112.

Kocherlakota, Narayana R. 2005. "Zero Expected Wealth Taxes: A Mirrlees Approach to Dynamic Optimal Taxation." *Econometrica* 73 (5): 1,587–621.

Kocherlakota, Narayana R. 2010. *The New Dynamic Public Finance.* Princeton, N.J.: Princeton University Press.

Mirrlees, James A. 1971. "An Exploration in the Theory of Optimum Income Taxation." *Review of Economic Studies* 38 (April): 175–208.

Prescott, Edward C., and Robert M. Townsend. 1984. "Pareto Optima and Competitive Equilibria with Adverse Selection and Moral Hazard." *Econometrica* 52 (January): 21–45.

Ramsey, F. P. 1927. "A Contribution to the Theory of Taxation." *Economic Journal* 37 (March): 47–61.

Rogerson, William P. 1985. "Repeated Moral Hazard." *Econometrica* 53 (1): 69–76.