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

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

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 over-saving. 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

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

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*

$$\begin{aligned} T_0^* &= G_0 & (23) \\ T_1^* &= G_1 - \tau^* \left(r_1^* K_1^* + \pi_1^* (y_{11}^* - c_{11}^* + c_{10}^*) \right). & (24) \end{aligned}$$

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

$$\begin{aligned} f(y_1) &= \chi_{\{0\}}(y_1), \\ \bar{e} &= r_1^* K_1^* + \pi_1 (w_1^* l_{11}^* - c_{11}^* + c_{10}^*), \end{aligned}$$

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

$$\begin{aligned} r_t &= r_t^*, \; w_t = w_t^* \text{ for } t = 0, 1, & (25) \\ q &= 1/r_1^*, & (26) \\ R &= r_1^*/\pi_1, & (27) \end{aligned}$$

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:

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

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.