

Fraud Deterrence in Dynamic Mirrleesian Economies*

Roc Armenter[†]

Thomas M. Mertens[‡]

November 2, 2010

Abstract

Insurance schemes rely on legal action to deter fraudulent claims. We capture this aspect by introducing a random state verification technology in a dynamic economy with private information. With some probability, an agent's skill level becomes known, allowing to punish misreporting agents. We demonstrate how deferring consumption can ease the provision of incentives. As a result, the marginal benefit of investment can be below its marginal cost, which suggests *subsidizing* savings. In an infinite horizon economy, the constrained-efficient allocation converges to a high consumption level, full insurance, and no labor distortions.

Keywords: Private information, state verification, capital taxation.

JEL classification: H21, E61.

*We thank Marco Bassetto, Mike Golosov, Oleg Itshoki, Cyril Monnet, Jaromir Nosal, Thomas Philippon, B. Ravikumar, Aleh Tsyvinski, Ivan Werning, Randy Wright, Yuzhe Zhang, and the audiences at the University of Maryland, University of Iowa, the Philadelphia Fed, the Duke Junior Macro Jamboree, the Cowles Summer Conference, and the Christmas meeting for German economists abroad for their comments and suggestions. The views expressed here do not necessarily reflect the views of the Federal Reserve Bank of Philadelphia or the Federal Reserve System. This paper is available free of charge at www.philadelphiafed.org/research-and-data/publications/working-papers.

[†]Federal Reserve Bank of Philadelphia, Ten Independence Mall, Philadelphia, PA 19206, e-mail: roc.armenter@phil.frb.org

[‡]New York University, Stern School of Business, 44 W 4th St. Suite 9-73, New York, NY 10012, e-mail: mertens@stern.nyu.edu

1 Introduction

Insurance schemes are intended to provide support to people in adverse situations. An individual's condition determines the payment but is typically hard to monitor. Full insurance might thus not be optimal as agents struggle to resist the temptation of pocketing more support than they are entitled to. In a seminal contribution, Mirrlees (1971) formalized the trade-off between insurance and incentives by treating the agent's skill as private information. The celebrated work by Golosov, Kocherlakota and Tsyvinski (2003) extends the analysis of private information to dynamic economies. The optimal scheme in these economies is always designed such that agents pick the payments intended for them, even if misreporting is undetectable.

Yet most of the insurance schemes we observe rely on the threat of legal action to induce truthful reporting. Laws against fraud are as commonplace as they are ancient.¹ Of course, the rule of law requires the means of verifying the validity of claims. For example, insurance companies routinely inspect a subset of claimants after they filed for disability insurance. Similarly, the elderly need to provide proof of eligibility in order to collect on nursing care insurance, and the government routinely selects households for tax audits.

In this paper, we take a first step toward incorporating the threat of punishment as means of incentive provision in an optimal scheme. We do so by modifying an otherwise dynamic standard Mirrleesian economy. Idiosyncratic skill levels evolve over time following a stochastic process; output is given by a production function with constant returns to both labor and capital. The planner observes the amount each agent produces and consumes, but, as in Mirrlees (1971), the effort exerted by the household is not observable. We study the constrained efficient allocations in both finite and infinite horizon economies.

In contrast to the standard Mirrleesian setup, the planner gets to verify the present skill level of a subset of agents selected randomly each period. For the remaining households, the present skill level remains private information: The planner has no means to check past skill levels. Households must put forth their claims before it is known whether their skill level will be observable. This “random state verification” technology effectively assumes that governments and private companies can check the validity of some but not all claims; and whether a particular claim is observable or not is out of the control of both the claimant and the provider.²

We find that the possibility of state verification shapes the decisions and allocations on the equilibrium path. Agents' incentives to save are sensitive to state verification: We find that the marginal benefit may be below the marginal cost of investment; and in infinite horizon

¹For example, see Buckland (1975), p. 594.

²This specification allows us to tackle a full-fledged dynamic private information economy and compare the results with the existing literature. The joint optimization of monitoring strategies and allocations is generally not tractable unless some strong assumptions are imposed on the economy.

economies, the constrained efficient allocation converges to full insurance, high consumption, and no labor distortions. In the standard dynamic Mirrleesian framework with no state verification, in contrast, the marginal benefit is always above the marginal cost of investment, and efficiency requires immiseration as consumption converges to zero for almost all agents.

In equilibrium, no agent ever misreports the type. Not surprisingly then, the planner always prescribes the worst continuation welfare as a punishment to an agent who is found to be misreporting his or her type. Incentive provision improves with higher probabilities of state verification and with sterner punishments. To focus on the implications of state verification, the worst continuation welfare, and thus punishment, is given by an exogenous lower bound on payoffs. Our setup encompasses the dynamic Mirrlees economy described in Golosov, Kocherlakota and Tsyvinski (2003) and the full information economy as special cases by setting the probability of state verification to zero and high enough, respectively. Our interest lies in the intermediate values.

The key insight is that, with random state verification, the provision of incentives can be facilitated by increasing savings and deferring consumption to later periods. An agent caught misreporting forgoes the continuation welfare prescribed under truthful revelation: The *effective* punishment is given by the difference between the continuation welfare along the equilibrium path and the worst continuation welfare. The planner can make the effective punishment stronger, and thus ease the provision of incentives, by deferring consumption and welfare, which in turn demands increased savings.

Deferring consumption, though, has its downside. As Golosov, Kocherlakota and Tsyvinski (2003) pointed out, the consumption spread needed for incentive provision increases with promised consumption. In the standard model with no state verification, this is the only force shaping intertemporal allocations and thus the planner finds it optimal to discourage savings.

There are thus two counteracting forces regarding the intertemporal allocation of resources in the optimal contract. Which force dominates dictates whether the marginal benefit of investment is above or below the marginal cost, that is, whether the intertemporal wedge is positive or negative, respectively, and whether a capital tax or a subsidy would be needed to implement the constrained efficient allocations. We illustrate both effects in a two-period economy.

An implied subsidy on capital is not just a mere possibility. Typically the sign of the intertemporal wedge will change across periods and across agents. We first discuss conditions for a negative intertemporal wedge in finite horizon economies. We show that the marginal benefit of investment will be strictly above the marginal cost at some node, provided that the probability of state verification is in a certain range. We also show that, for any positive probability of state-verification, there exists a long enough (but finite) time horizon such that the intertemporal wedge is non positive.

The results are even starker for infinite horizon economies: All distortions vanish in the long run. We find that, in the long run, almost all agents receive full insurance with a high level of consumption; and there are no distortions on the labor decision. Surprisingly, the result holds for *any* (even arbitrarily small) probability of state verification. We show how the planner can promise a high level of future consumption—high enough to make the effective punishment deter any misreporting even if full insurance is provided. Loosely speaking, the economy can outgrow the constraints imposed by private information.³ However, doing so is quite costly because the economy needs to accumulate aggregate capital and incur in low consumption along the transition path. Nevertheless, convergence to full insurance occurs with probability one. The rationale for deferring consumption and welfare dominates in an infinite horizon economy because all agents will experience a (long enough) streak of high skill level realizations. Along with these positive shocks, it becomes easier and easier to provide incentives with future consumption: Eventually the effective punishment is strong enough to support full insurance, which becomes an absorbing state.

The long-run properties stand in sharp contrast with the standard dynamic Mirrleesian economy with no state verification: Golosov, Kocherlakota and Tsyvinski (2003) show how the marginal benefit is always above the marginal cost of investment, and consumption converges to a lower bound for almost all agents. Note that a positive probability of state verification allows the planner to front-load all distortions, which is not possible at all in the standard Mirrleesian economy. Our result is connected with the general principle in Albanesi and Armenter (2008): As long as it is possible to front-load all the distortions present in the economy, the constrained efficient allocation rules out permanent intertemporal distortions.

We also find that agents whose skill level is observable are rewarded with higher consumption and leisure as long as incentive-compatibility constraints are binding. Thus, it is optimal to introduce some consumption uncertainty linked to the possibility of state verification. In infinite horizon economies, however, the rewards are eventually phased out.

In contrast with the savings decisions, the prescriptions regarding the labor supply are robust to the introduction of the random state verification. As in the standard Mirrleesian framework, the labor decision of the agent with the top skill level is left undistorted, while the output of other agents is taxed. The exception is that the labor supply decision of agents whose skill level has been verified always remains undistorted.

We discuss possible foundations for a theory of optimal monitoring and punishment that would validate our main assumptions. In the paper, we work with strong assumptions regarding the state-verification and punishment technology. These assumptions allow us to build a clear intuition for a general dynamic private information economy. We demonstrate that under some

³The assumption of constant returns to capital plays an important role here, ensuring that high consumption is indeed feasible.

conditions the optimal design would deliver these properties embedded in our assumptions, namely, that the probability of state-verification is not zero and the punishment level does not scale up with continuation welfare. More specifically, we study optimal choices of punishment when the planner's monitoring technology produces false positives. Furthermore, we derive optimal probabilities of state-verification for various degrees of history-dependence.

Our work relates to the literature on the optimal taxation of savings. Early work on economies with linear factor taxation finds a zero tax rate on capital income, see Judd (1985) and Chamley (1986). Subsequent research has further explored the Mirrlees approach to dynamic optimal taxation and disability insurance. A necessarily incomplete list in a rapidly increasing literature is Albanesi and Sleet (2006); Kocherlakota (2005); Golosov and Tsyvinski (2006); Golosov, Tsyvinski and Werning (2006); Werning (2007); Farhi and Werning (2007); Mankiw and Weinzierl (2009); and Weinzierl (2008).

There is a literature on the optimal design of audit and law compliance mechanisms in private information economies, dating back to Stigler (1970). Border and Sobel (1987) characterize an optimal audit mechanism design for a static economy in which agents have private information on their endowment. Cremer and Gahvari (1995) analyze the implications of an audit scheme for static private information economies. More recently, Ravikumar and Zhang (2009) study incentive provision through tax audits in a dynamic hidden income model. A closely related strand of the literature explores the costly state verification model of Townsend (1979). Wang (2005), Monnet and Quintin (2005), and Popov (2007) expand the application of costly state verification to dynamic economies. We take the state verification technology as given and focus instead on its implications for a full-fledged dynamic private information economy.

The paper is structured as follows. Section 2 lays out the setup for the general economy. Section 3 presents a simple economy with just two periods that illustrates our results concerning the incentives to save and the allocation of consumption and output across agents. We start with our formal results on the allocation of consumption and output across agents in Section 4, and follow with our results for the incentives-to-save in Section 5. Our results concerning finite horizon economies are in Section 6. Section 7 documents our findings on the long-run properties of the constrained efficient allocation. Section 8 discusses the optimal design of state-verification and punishment. Section 9 concludes.

2 Setup and Constrained-Efficient Allocations

2.1 Setup

Time is discrete $t = 0, 1, \dots, T$ with T possibly being infinite. The economy is populated by a continuum of infinitely lived, ex-ante identical agents. Agents are subject to idiosyncratic shocks to their skill levels. At each date t , the agent's skill $\theta_t > 0$ takes on a value from a finite support Θ , according to the probability distribution $\{\pi(\theta) : \theta \in \Theta\}$. Let $\theta^t = \{\theta_0, \theta_1, \dots, \theta_t\}$ be the history of skill realizations and Θ^t the corresponding support for histories up-to-date t . The skill level of an agent determines the labor effort l_t needed to produce output $y_t = \theta_t l_t$. We assume that skill levels are independent across agents and that the law of large numbers applies. We denote with $\pi(\theta^t)$ both the fraction of agents with history θ^t and the probability of such an event.

All agents derive utility from consumption $c_t \in \mathbb{R}_+$ according to the utility function $u : \mathbb{R}_+ \rightarrow \mathbb{R}$ and disutility from labor $l_t \in \mathbb{R}_+$ according to $v : \mathbb{R}_+ \rightarrow \mathbb{R}$. We find it useful to express the per-period payoff in terms of output rather than labor. Let the consumption-output bundle for an agent be $x_t = \{c_t, y_t\} \in \mathbb{R}_+^2$, then his or her utility is given by

$$h(x_t; \theta_t) = u(c_t) - v\left(\frac{y_t}{\theta_t}\right)$$

where $h : \mathbb{R}_+^2 \times \Theta \rightarrow \mathbb{R}$.

Relative to the literature on Mirrleesian economies, the key innovation of this paper is the possibility that the skill level becomes observable to a planner. An exogenous process determines whether the skill level of an agent is only privately known or observable. Let a_t be an i.i.d. random variable that takes the value $a_t = 1$ with probability p and $a_t = 0$ otherwise. At every date and for every agent, variable $\alpha_t = a_t \theta_t$ is observable: However, if $\alpha_t = 0$, then the variable is uninformative about the skill level.⁴ Let a^t be the history of realizations and $A^t = \{0, 1\}^t$ the corresponding support. Note that only the skill level at the current date is revealed in a given period. In terms of agents' decisions, we assume that consumption and output are observable, but the exerted effort is not.

Each agent's type consists of the history of skills and state verifications $s^T = \{\theta^T, a^T\} \in \Theta^T \times A^T \equiv S^T$. From the probability process for skills and the probability of state verification, we can define a probability over the set of types S^t for $t = 0, 1, \dots, T$, which we denote with $\pi(s^t)$ in a slight abuse of notation.

⁴This specification is a convenient shortcut for notation but requires that the uninteresting case of $\theta \neq 0$ is ruled out. Equivalently, we can let α be a random variable with values in $\Theta \cup \{\emptyset\}$. For each agent, α_t equals his or her skill level θ_t with probability p , and it is uninformative, $\alpha_t = \{\emptyset\}$, with probability $1 - p$.

Let x denote a mapping that assigns each type in $S^T \rightarrow \mathbb{R}_+^{2T}$ consumption and labor allocations at every date. Agents have von Neumann-Morgenstern preferences and value a plan x according to

$$W_0(x) = \sum_{t=0}^T \sum_{s^t \in S^t} \beta^t \pi(s^t) h(x(s^t); \theta_t)$$

where $\beta \in (0, 1)$ is the time discount factor. As part of the setup, we assume that period utility cannot drop below a lower bound v , i.e., $h(x(s^j); \theta_j) \geq v$. It turns out to be more convenient to translate this lower bound on utility into a the lower bound on continuation welfare. For all dates and all type realizations $s^t \in S^t$,

$$\sum_{j=t}^T \sum_{s^j \in S^j} \beta^{j-t} \pi(s^j | s^t) h(x(s^j); \theta_j) \geq (1 - \beta^{T+1-t}) \Upsilon \quad (1)$$

where $\Upsilon = v(1 - \beta)^{-1}$ and $\pi(s^j | s^t)$ is the conditional probability of history s^j given s^t . The lower bound on welfare plays an important role in our analysis as it determines the maximum punishment a planner can prescribe for a deviant agent. Section 8 discusses several theories of punishment which are compatible with our analysis.

We list below all the assumptions on preferences. In addition to the standard convexity and differentiability properties, we assume a particular set of boundary conditions that allow us to focus on interior allocations. The boundary conditions also ensure that the lower bound on welfare (1) can be attained with an allocation in the positive quadrant.

Assumption 1 (Preferences)

Preferences are such that

1. *(Monotonicity) Both $u(c)$ and $v(l)$ are strictly increasing;*
2. *(Differentiability) Both $u(c)$ and $v(l)$ are twice differentiable;*
3. *(Convexity) Utility function $u(c)$ is strictly concave and $v(l)$ is strictly convex;*
4. *(Boundary conditions) Utility function $u(c)$ is unbounded above and below, and $v(l)$ is bounded below and satisfies $v'(0) = 0$.*

Assumption 1.4 ensures that equation (1) specifies the effective lower bound. The disutility of labor is bounded from below, and hence we can normalize it to $v(0) = 0$ without loss of generality.

An allocation is a pair $\{x, K\}$ where $K = \{K_{t+1}\}_{t=0}^T$ is a nonnegative sequence specifying aggregate capital at all dates $t = 1, 2, \dots, T$. The initial level of capital $K_0 > 0$ is taken as given.

A feasible allocation satisfies the lower bound on welfare (1) for all $s^t \in S^t$ and the resource constraint at all dates t ,

$$C_t + K_{t+1} \leq RK_t + Y_t \tag{2}$$

where $R = \beta^{-1}$ denotes the gross interest rate, and C_t and Y_t stand for aggregate consumption and output, respectively.⁵ The production function displays constant marginal returns to capital. We will return to this assumption later in the paper as it plays an important role for some of our long-run results.

2.2 Efficient and constrained-efficient allocations

Before discussing the implications of information frictions, we define the best allocation only subject to feasibility. We denote such an allocation as “efficient” or “first best” to distinguish it from “constrained efficient,” or “second best,” allocations to be defined later.

Definition 1 (Efficient allocation)

A feasible allocation $\{x, K\}$ is efficient if

$$W_0(x) \geq W_0(x')$$

for all feasible allocations $\{x', K'\}$.

Condition (1) implies a lower bound on W_0 . To guarantee that the objective function W_0 is also bounded from above, we need some additional conditions on preferences. We therefore adopt a standard assumption used for economies with constant marginal returns to capital. The assumption encompasses several utility function specifications, including logarithmic and constant elasticity of substitution (CES) utility functions.

Assumption 2

There exists scalars a, b , and $\gamma < 1$ such that $u(c) \leq a + bc^\gamma$ for all $c \geq 0$.

This assumption ensures that the objective function is finite for the set of feasible allocations. Consequently, we can characterize the efficient and constrained efficient allocations as optimal control problems.

⁵Aggregate consumption and output are the sum over individual consumption and output weighted by the mass of people:

$$\begin{aligned} C_t &= \sum_{s^t} \pi(s^t) c(s^t) \\ Y_t &= \sum_{s^t} \pi(s^t) y(s^t). \end{aligned}$$

Let us now turn to analyzing the set of feasible allocations that can be implemented given the presence of private information. We start by being precise about how information is revealed and communicated to both planner and agents alike. At the beginning of each period t , the skill level θ_t becomes known only to the agent. The agent then submits a report to the planner, $\sigma_t(s^{t-1}, \theta_t) \in \Theta$. The state verification variable a_t is realized only after the report has been filed. The planner then dictates a consumption and output allocation given the report, the realization of state verification a_t , and, if the skill level is revealed, the true type of the agent. We allow the planner to specify different allocations for an agent, depending on whether the skill level is public information, even if the agent is reporting truthfully.

Formally, we define a reporting strategy as a mapping $\sigma_t : S^T \rightarrow \Theta^T$ such that for all dates $t \geq 0$, the report σ_t is measurable with respect to $\{s^{t-1}, \theta_t\}$. The truth-telling reporting strategy σ^* is given by $\sigma_t^*(s^{t-1}, \theta_t) = \theta_t$ for all realizations in S^T . Let Σ be the set of reporting strategies. A contract ξ specifies consumption and output at all dates given history s^{t-1} , report $\hat{\theta}_t$, and observable variable $\alpha_t = a_t\theta_t$.⁶ Let $\xi(\sigma)$ be the allocation plan spanned by the reporting strategy σ .

A feasible allocation $\{x, K\}$ is admissible if $\xi(\sigma^*) = x$ for some contract ξ such that

$$W_0(\xi(\sigma^*)) \geq W_0(\xi(\sigma))$$

for all $\sigma \in \Sigma$. That is, the truth-telling strategy is optimal for all agents. We now use this language to define a constrained efficient allocation.

Definition 2 (Constrained-efficient allocation)

An admissible allocation $\{x, K\}$ is constrained efficient if

$$W_0(x) \geq W_0(x')$$

for all admissible allocations $\{x', K'\}$.

2.3 Characterization of constrained-efficient allocations

We seek to characterize constrained-efficient allocations as a solution to an optimal control problem. The first step is to write the incentive-compatibility constraints in terms of the allocation plan. This task is simplified by two well-known results. First, we only need to consider one-step deviations from the truth-telling strategy. Second, we can write the incentive-compatibility constraint knowing that the planner prescribes the lowest continuation welfare possible whenever an agent is caught misreporting his or her skill level. The reason is that no agent will actually

⁶Note that a contract and allocation are different mappings because a contract must be measurable over the support for the report and observable variable α_t in addition to the actual skill level.

misreport on the equilibrium path: to prescribe the maximum punishment available in such events comes at no cost and improves incentives on the equilibrium path.

We simplify further the formulation of the incentive-compatibility constraints by keeping track of promised utility. Let $z(s^t)$ be the expected continuation utility for date $t + 1$ after history s^t . We state our promise-keeping constraint to make the definition of $z(s^t)$ explicit,

$$z(s^t) = \sum_{s^{t+1} \in S^{t+1}} \pi(s^{t+1}|s^t) \{h(x(s^{t+1}); \theta_{t+1}) + \beta z(s^{t+1})\} \quad (3)$$

for all histories $s^t \in S^t$, and dates $t \geq 0$. If T is finite, the terminal condition, $z(s^T) = 0$, applies to all histories $s^T \in S^T$, as part of the set of promise-keeping constraints. These constraints allow us to rewrite the set of admissible allocations in a different form. Using the promised utility, we can express truthful reporting in a single constraint for each type and period. A feasible allocation $\{x, K\}$ is admissible if and only if the following incentive compatibility constraint is satisfied

$$\sum_{a_t \in \{0,1\}} \pi(a_t) \{h(x(s^t); \theta_t) + \beta z(s^t)\} \geq p(1 - \beta^{T+1-t}) \Upsilon + (1 - p) \{h(x(\hat{s}^t); \theta_t) + \beta z(\hat{s}^t)\} \quad (4)$$

at all dates $t \geq 0$, for all $\{s^{t-1}, \theta_t\} \in S^{t-1} \times \Theta$ and all skill levels $\theta' \in \Theta$, where $\hat{s}^t = \{s^{t-1}, \theta', 0\}$. Note that there is no separate incentive-compatibility constraints for $a_t = 0$ and $a_t = 1$: The agent must submit the report before knowing the realization of a_t . Thus, the value of truthful reporting and misreporting is given by a lottery spanned by the realization of the state verification variable a_t .

We characterize constrained efficient allocations as solutions to the following optimal control problem

$$\sup_{\{x, K, z\}} W_0(x)$$

subject to the resource constraint (2) at all dates $t = 0, 1, \dots, T$, and to the promise-keeping and incentive-compatibility constraints (3) and (4) at all nodes $s^t \in S^t$, as well as nonnegativity constraints and the lower bound constraint on continuation welfare (1).

The state verification technology relaxes the incentive-compatibility constraints, as can be seen from equation (4), and thus enlarges the set of admissible allocations. Before we proceed further, we ask under which conditions the efficient allocation is admissible, rendering trivial the problem of constrained efficiency. From the incentive-compatibility constraint (4), the role of the lower bound on welfare, v , becomes clear. If welfare were unbounded below, the planner could prescribe an infinite punishment in the event of misreporting: the efficient allocation would then trivially satisfy the incentive-compatibility constraints. We shall treat the lower bound v as a

fixed parameter for the remainder of the paper.

The next assumption ensures that the lower bound is low enough such that the set of admissible allocations exists and is not a singleton. To ensure that the lower bound can be achieved with a feasible allocation, we also assume Υ is finite.

Assumption 3

The lower bound v is set in a way that the set of admissible allocations has a non-empty interior and $v > -\infty$.

Taken the punishment level as given, we now ask about conditions on the state verification probability such that the efficient allocation is admissible. Clearly, the efficient allocation is not implementable for a sufficiently low state verification probability p . For the case in which the probability of state verification is zero, we retrieve the standard dynamic Mirrlees formulation. In that case, the efficient allocation is never admissible. On the other end, if p is large enough, the efficient allocation satisfies the incentive-compatibility constraints (4). Indeed, there exists a threshold p^* such that the set of admissible allocations for $p < p^*$ does not contain the efficient allocation.

Lemma 1

There exists $p^ \in (0, 1)$ such that the efficient allocation is not admissible if and only if $p < p^*$.*

3 A Simple Two-Period Economy

We use a two-period economy to illustrate the mechanics of the setup. We are interested in how exogenous state verification shapes the incentives to save and work under the constrained efficient allocation.

Therefore, we use the setup of the previous section with only two periods as a special case. For illustrative purposes, we do not allow the planner to give out rewards for truth-telling. At time $t = 1$, each agent chooses how much of his or her initial endowment to consume, c_1 , leading to an amount of savings $s = y_1 - c_1$. All agents behave symmetrically in the first period since there is no source of heterogeneity.

In the second period, agents learn their idiosyncratic skill level. Half of the agents are able to produce one unit of output with $1/\theta$ units of effort where $\theta > 1$. The other half produces only one unit of output for each unit of effort. The resource constraint at date $t = 2$ is given by

$$\frac{1}{2} (c_2(\theta) - y_2(\theta)) + \frac{1}{2} (c_2(1) - y_2(1)) \leq s.$$

where $\{c_2(\theta), y_2(\theta)\}$ and $\{c_2(1), y_2(1)\}$ denote consumption and output produced at date $t = 2$

for agents with skill level θ and 1, respectively. For simplicity, we assume that the return to savings is 1. A feasible allocation x is a vector specifying consumption at date $t = 1$, savings, and consumption and output at both dates as a function of the skill level such that the resource constraint is satisfied at both dates.

3.1 Constrained-efficient allocation

The incentive-compatibility constraints do not reflect any promised utility since we only set up a two-period economy. Hence, we can write them as

$$u(c_2(\theta)) - v(y_2(\theta)/\theta) \geq pv + (1-p)(u(c_2(1)) - v(y_2(1)/\theta)) \quad (5)$$

where v stands for the amount of utility agents receive if caught misreporting. The incentive-compatibility constraint ensures that the high skill level reveals his or her type before knowing whether it will be observable. As in the standard Mirrlees problem, the corresponding incentive-compatibility constraint for the low skill agent is automatically satisfied.

The possibility of the skill level becoming public information relaxes the incentive compatibility constraint. The higher the probability of observing the true type and the more severe the punishment, the easier it is to provide incentives. Indeed, if v is low enough or p is sufficiently close to 1, the first best allocation becomes incentive compatible. On the other end, if skills are never verified, we recover the standard formulation of the dynamic Mirrleesian setup.

To arrive at the constrained efficient allocations, we maximize social welfare

$$u(c_1) - v(y_1) + \frac{1}{2} \left(u(c_2(\theta)) - v\left(\frac{y_2(\theta)}{\theta}\right) \right) + \frac{1}{2} (u(c_2(1)) - v(y_2(1)))$$

subject to the incentive compatibility constraint (5) and the resource constraint.

We are interested in the constrained efficient allocation, that is, the incentive-compatible allocation attaining the highest welfare. From the first-order conditions of the optimal control problem, we derive a relation between the consumption levels of the two different types in the second period:

$$u_2^c(\theta)(1 + \mu) = u_2^c(1)(1 - (1-p)\mu)$$

where μ is the Lagrangian multiplier associated with the incentive-compatibility constraint.⁷ As long as the first best allocations are not incentive compatible, that is $\mu > 0$, consumption is not equated across agents. More precisely, the agent of type θ has a higher level of consumption, as in the standard Mirrlees problem.

⁷We have normalized the Lagrangian multiplier μ by the measure of agents of type θ .

We start by exploring how private information and the state verification technology shape agents' incentives to work. We find that the prescriptions of the static Mirrlees problem are remarkably robust. For any state verification probability, the marginal rate of substitution is equated to the marginal rate of transformation for the agent with the highest skill.

$$u'(c_2(\theta)) = \frac{1}{\theta} v' \left(\frac{y_2(\theta)}{\theta} \right)$$

Thus, the implied marginal labor wedge is zero for the agent with high skills, as in Mirrlees (1971). As long as the first-best allocation is not incentive compatible, the allocations suggest a positive labor tax for low skill types.

$$v'(y_2(1)) \left(1 - (1-p)\mu \frac{v'(y_2(1)/\theta)}{\theta v'(y_2(1))} \right) = (1 - (1-p)\mu) u'(c_2(1))$$

Both results follow immediately from the first-order conditions for consumption and output.

The probability of state verification p influences the magnitude of the labor wedge of the constrained-efficient allocation. It is easier to provide incentives as p increases, and the optimal labor wedge is decreasing. Figure 1 plots the optimal labor wedge for the low skill agent as a function of p .⁸ Labor is undistorted only when the probability of state verification is high enough that the incentive-compatibility constraint is no longer binding. In other words, labor distortions only disappear when the first-best allocation is incentive compatible. We show these results for the general case with multiple types and many periods in section 4.

Let us now sign the intertemporal wedge for the optimal allocation. The question here is whether the intertemporal distortion suggests a tax or subsidy for capital. By manipulation of the first-order conditions derived from the maximization problem above, we arrive at an intertemporal optimality condition given by

$$E \left\{ \frac{1}{u_2^c} \right\} = \frac{1}{u_1^c} \left(1 + \frac{p\mu}{2} \right). \quad (6)$$

In the special case in which types are never verified, equation (6) encompasses the well-known *inverse Euler equation*,

$$u_1^c = \left[E \left\{ \frac{1}{u_2^c} \right\} \right]^{-1}.$$

In this case, there is always variation in consumption across types. Jensen's inequality then implies $u_1^c < E \{u_2^c\}$. The intertemporal wedge must be positive in order to implement the con-

⁸Agents value allocations according to preferences that satisfy the assumptions of the previous section. For our simple economy, we choose utility of consumption $u(c) = \log(c)$ and disutility of labor $v(l) = -l^4/4$ with parameter choices $\theta = 2$ and $v = -1/4$.




Figure 1: Implied labor tax rate for various probabilities of state verification.

strained efficient allocations. This result carries on for infinite horizon economies with arbitrary stochastic processes for the agents' type.

Now consider the possibility of exogenous state verification, $p > 0$. As long as $p\mu > 0$ in equation (6), we cannot guarantee the above result of a positive implied savings tax.

Fortunately, it is still possible to sign the term $u_1^c - E\{u_2^c\}$. Using the first-order conditions for $c_2(\theta)$ and $c_2(1)$, we can solve for $E\{u_2^c\}$ in terms of μ and p , which in turn implies $u_1^c \geq E\{u_2^c\}$ whenever

$$\mu \leq \frac{1}{2} \frac{p}{1-p}. \quad (7)$$

Moreover, $u_1^c = E\{u_2^c\}$ if and only if $\mu = 0$ or (7) holds with strict equality. It is possible to express condition (7) in terms of a threshold \hat{p} for the probability of state verification at which the equality is satisfied. Heuristically, the Lagrangian multiplier μ is decreasing in p as the incentive-compatibility constraints become relaxed, and the right-hand side of (7) is increasing in p . We derive the general version of this inverse Euler equation for our setup in Section 5.

We thus find an implied subsidy if it is likely enough to observe the skill level, $p \geq \hat{p}$. Figure 2 displays the implied savings tax needed to implement the constrained efficient allocation for different state-verification probabilities. The comparative statics for the implied tax are not monotone and lead to a sign-switch: around $p = .05$, the optimal saving tax turns into a subsidy. The exact value of the threshold depends on the primitives of the economy—the curvature of the preferences, the dispersion in skills, as well as the lower bound on welfare. Importantly though, the first-best allocations are not incentive compatible at $p = \hat{p}$.

Before interpreting our results, let us first briefly revisit why the dynamic Mirrlees problem

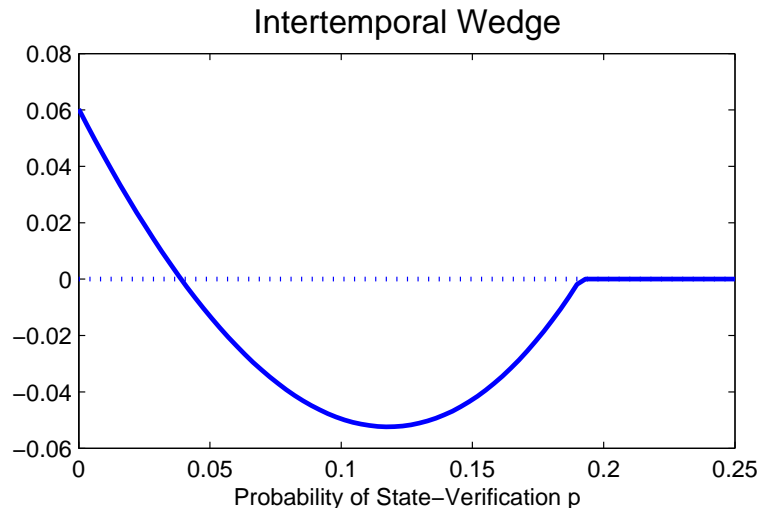


Figure 2: Implied savings tax rate for various probabilities of state verification.

calls for savings to be taxed. The incentive-compatibility constraints require some spread between the consumption of high and low skill level agents. The more resources that are available in the economy, the larger the spread must be. In other words, savings make the provision of incentives more difficult. As a result, consumption and welfare are frontloaded under the constrained-efficient allocation.

The possibility of observing the agent's type ex-post opens a new channel to improve incentives. Being able to verify the state allows to prescribe a punishment for any agent caught misreporting the skill level. Since this punishment is never administered on the equilibrium path, the planner equates the welfare of any misreporting agent to the lower bound on welfare. Everything else constant, the higher expected consumption for truthful reporting, the sterner the punishment for misreporting, and thus the more effective is state verification in relaxing the incentive-compatibility constraints. Hence, additional resources, and thus savings, can improve the provision of incentives by raising overall welfare.

In short, savings can help the provision of incentives in the presence of state verification. Whether savings are taxed or subsidized depends, ultimately, on which effect dominates. Indeed, there is a probability of state verification such that the implied tax on savings is negative. We discuss the incentives to save in more detail in Section 5.

We show the existence of an implied subsidy of capital in Section 6 for the finite horizon case. There, we also discuss its economic interpretation. In the infinite horizon case, we can demonstrate even stronger predictions: All economic agents will eventually have a nonpositive intertemporal wedge (see Section 7).

4 Output and consumption distribution

We start our analysis of the general setting with the allocation of consumption and output across agents. First, we want to know whether agents allocate between consumption and leisure optimally. In particular, we explore how allocations compare between an agent whose skill level becomes public information and an agent with identical history and skill level but the latter remains private information.

The derivations of the first-order necessary conditions can be found in the Appendix, where we also prove that constrained-efficient allocations are interior.⁹

Let us first consider the case in which the agent's skill level remains private information, $a_t = 0$. We find that the results of the classic, static Mirrlees economy carry over. An agent with skill θ_t will face a wedge between the marginal rate of substitution and the marginal rate of transformation as long as some other agent θ' incentive-compatibility constraint to report the type θ_t is binding, that is, the Lagrangian multiplier $\mu(s^{t-1}, \theta'; \theta_t)$ is positive. By encouraging leisure, the allocation prescribed for skill level θ_t looks less attractive to agents with higher skill levels. Again as in Mirrlees (1971), no agent has incentives to report the top skill level, and thus the marginal rate of substitution and the marginal rate of transformation are equated for agents with top skill. Note that the presence of positive marginal labor taxes is assured as long as private information constrains the set of admissible allocations.¹⁰ We summarize these findings in the following proposition.

Proposition 1 (Mirrleesian results)

For any date $t = 0, 1, \dots, T$ and node $s^{t-1} \in S^{t-1}$ with $a_t = 0$,

$$h^y(s^t; \theta_t) \leq u'(s^t) F_t^y$$

with strict inequality iff $\mu(s^{t-1}, \theta'; \theta_t) > 0$ for some $\theta' \in \Theta$.

Proof: In the Appendix. □

We now turn our attention to allocation whenever the skill level becomes public information, $a_t = 1$. Allocations conditional on $a_t = 1$ are not part of the payoff of misreporting for any agent. It is thus possible to relax the incentive-compatibility constraints by prescribing different allocations in the event $a_t = 1$. Indeed we find that the constrained efficient allocation does so.

First, there is no need to distort the labor supply decision of agents whose type is public information. The only reason for the labor supply to be distorted for some agents is to ease

⁹See Proposition 6.

¹⁰Our assumptions on preferences rule out the possibility of “bunching,” in which some agents with skill level $\theta_t < \bar{\theta}$ face a zero marginal tax. See Tuomala (1990) and Salanie (2003) for a discussion and references herein.

the incentives to misreport. But there is no need to do so for the allocation with observable skills. In this case, delivering the efficient allocation between consumption and leisure *relaxes* the incentive-compatibility constraints (4). Any tax system that would decentralize the constrained efficient allocation would need to include a provision for a zero marginal tax rate—a tax break—whenever the skill level of an agent is public information.

We go one step further by ranking consumption and labor allocations between agents with and without state verification. We find that for two agents with identical histories s^{t-1} and identical skill level θ_t , consumption is weakly higher for the agent whose type is public information,

$$c(\{s^{t-1}, \theta_t, 0\}) \leq c(\{s^{t-1}, \theta_t, 1\})$$

with strict inequality as long as the incentive-compatibility constraint is binding. The possibility of state verification thus introduces consumption variation *along the equilibrium path*. Similarly, the agent with observable skill enjoys not only more consumption but also more leisure:

$$y(\{s^{t-1}, \theta_t, 0\}) \geq y(\{s^{t-1}, \theta_t, 1\}).$$

Furthermore, the agent with observable skills will also continue to receive higher utility than the agent whose skill level remained private information. We infer this observation from a higher Lagrangian multiplier on the constraint on promised utility.

Agents are thus “rewarded” whenever their skill level is observed by having more consumption and more leisure, as well as more continuation welfare. This result may appear surprising at first: the planner could prescribe full redistribution among all agents whose skill is observed, since there is no chance to misreport. Agents, though, must report their type before they know whether it will be observable or not. It is thus important to consider the ex-ante effect on incentives. By delivering higher welfare in the event of state verification, say, for agent θ , the expected value of truth-telling increases for agent θ without increasing the payoff of misreporting for any other agent $\theta' \neq \theta$ — because if agent θ' is observed and found to be misreporting, she would be punished, not rewarded. The variation in consumption, however, is costly, since agents are risk averse. Risk aversion thus puts a bound on the optimal spread in consumption levels between agents with observed and unobserved type.

Proposition 2 formalizes these results.

Proposition 2 (Rewards for state verification)

For any date $t = 0, 1, \dots, T$ and node $s^{t-1} \in S^{t-1}$,

$$\begin{aligned} c(\{s^{t-1}, \theta_t, 0\}) &\leq c(\{s^{t-1}, \theta_t, 1\}) \\ y(\{s^{t-1}, \theta_t, 0\}) &\geq y(\{s^{t-1}, \theta_t, 1\}) \end{aligned}$$

for all $\theta_t \in \Theta$, and with strict inequality iff $\mu(s^{t-1}, \theta'; \theta_t) > 0$ for some $\theta' \in \Theta$. Moreover, if $a_t = 1$,

$$h^y(s^t; \theta_t) = u'(s^t) F_t^y.$$

Proof: In the Appendix. □

Finally, we turn our attention to the agent with the highest skill level $\bar{\theta} = \max\{\theta \in \Theta\}$. Since our preferences satisfy the single-crossing condition, the incentive-compatibility constraints for agents with $\theta_t < \bar{\theta}$ to report $\bar{\theta}$ are not binding. It follows from Proposition 2 that the allocations of consumption, output, and promised utility for the agent with top skill level are identical, regardless of whether the skill level is made public information or remains private information. The reason is that no agent with $\theta_t < \bar{\theta}$ has incentives to misreport their type as $\bar{\theta}$. Thus, prescribing distinct allocations for $a_t = 0$ and $a_t = 1$ does not relax any incentive-compatibility constraint. Since it does not improve incentives, providing the same allocation whether the state is verified or not delivers the same welfare to the top skill agent with less resources.

5 The incentives to save

This section is concerned with the allocation of resources across periods. This topic has been the focus of most of the literature on dynamic Mirrleesian economies, starting with the finding in Golosov, Kocherlakota and Tsyvinski (2003) that the marginal benefit of investment is above the marginal cost. The standard dynamic Mirrleesian economy is nested in our setup when the probability of state verification equals zero.

Since aggregate capital does not directly affect incentives, the first-order condition for K_t equates the shadow cost of resources λ_t across periods as the efficient allocation would do. Of course, the shadow cost of resources will generally not be the same for efficient and constrained-efficient allocations, and thus the optimal paths for aggregate capital will differ as well. Without further assumptions, it is difficult to say much about the transition path for capital.

Fortunately we can characterize the incentives to save for private agents. Proposition 6 in the Appendix shows that consumption and output are strictly positive in the constrained-efficient allocation so that the first-order conditions hold with strict equality. Given this result, we follow the same steps involved in the derivation of the inverse-Euler equation—see Rogerson (1985) and Golosov, Kocherlakota and Tsyvinski (2003). The resulting optimality condition for intertemporal allocations reads

$$u'(s^t) = \left(E_t \left\{ \frac{1}{u'(s^{t+1})} \right\} \right)^{-1} \left(1 + p \frac{\mu^*(s^t)}{1 + \eta(s^t)} \right) \quad (8)$$

where

$$\mu^*(s^t) = \sum_{\theta, \theta' \in \Theta} \{\pi(\theta)\mu(s^t, \theta; \theta')\}$$

is a weighted average of Lagrangian multipliers $\mu(s^t, \theta; \theta')$ on the incentive-compatibility constraints.

If there is no chance that the state is verified, $p = 0$, we recover the familiar inverse Euler equation from (8). Following Golosov, Kocherlakota and Tsyvinski (2003), we can establish the existence of a positive wedge between the marginal benefit and the marginal cost of investment, as

$$u'(s^t) = \left(E_t \left\{ \frac{1}{u'(s^{t+1})} \right\} \right)^{-1} < E_t \{u'(s^{t+1})\}$$

where the last step follows from Jensen's inequality. If we were to decentralize the constrained-efficient allocation in the case $p = 0$, we would find that we need a positive capital tax. In short, reducing savings helps in improving the provision of incentives in future periods because it becomes easier to provide the necessary spread in the marginal utility of consumption.

On the other end of the spectrum where all the skill types are observed, information is no longer private. Hence, the first-best allocation becomes feasible for the planner. The interesting results of this paper thus deal with probabilities in between these two extrema.

Once we consider the possibility of state verification, however, we can no longer establish unambiguously that the marginal benefit of investment is above the marginal cost, or vice versa. The sign of the intertemporal wedge $E_t \{u'(s^{t+1})\} - u'(s^t)$ depends on the size of Lagrangian multipliers and probabilities of state verification.¹¹

Random state verification opens a new channel to improve incentives for future periods. Recall that an agent caught misreporting his or her type is assigned continuation welfare equal to Υ . Thus, the amount of welfare forgone once caught misreporting is $z(s^t) - \Upsilon$, what we call the effective punishment. The harsher the effective punishment is, the better the incentive provision. Everything else equal, the planner can improve allocations by promising higher continuation welfare. In turn, higher continuation welfare requires accumulating more capital in order to generate the necessary additional resources.

However, the opposite force of standard private information economies remains: Incentives require a larger consumption spread when there are more savings. Whether the constrained efficient allocation prescribes a positive or negative intertemporal wedge (8) depends on which channel is stronger. In other words, savings present the planner with a trade-off between more effective punishment and a larger consumption spread.

¹¹The exception is when $\mu^* = 0$. In this case, no incentive-compatibility constraint is binding at nodes $s^{t+1} \in S^{t+1}|s^t$, and the allocation features full insurance with respect to date $t + 1$ shocks.

The random state verification channel has a key property: The effective punishment can be large enough to overcome all the incentive problems. That is, if the continuation welfare is high enough, then no agent risks losing it by misreporting, even if the allocations prescribe full insurance and no labor distortions. For constrained-efficient allocations, we can say even more: Full insurance and no labor distortions are necessary properties if the continuation welfare is above a certain threshold. Proposition 3 formalizes the first key result of the paper.

Proposition 3

Let $p > 0$. There exists threshold $z^* < \infty$ such that if $\{x, K, z\}$ is a constrained-efficient allocation and

$$z(s^{t-1}) \geq z^*,$$

then for all dates $j \geq t$ and continuation nodes $s^j \in S^j|s^{t-1}$, there are no labor distortions:

$$v'(y(s^j)/\theta_j) = u'(c(s^j))\theta_j,$$

and there is full insurance:

$$c(s^j) = c(s^t)$$

for all $s^t \in S^t|s^{t-1}$.

Proof: In the Appendix. □

The proof of Proposition 3 is quite straightforward. The welfare gains for misreporting the skill type and going undetected are bounded above when the efficient allocations are prescribed. We thus only need a large enough effective punishment to enforce truthful reporting.

Proposition 3 does not imply that the upper bound will ever be reached. Whether we will reach the range of full insurance is a matter of discussion in the following sections. However, it already tells us that we cannot have very high levels of promised utility and still have the same prediction of a positive intertemporal wedge as for standard dynamic Mirrleesian economies.

Next we highlight a key implication of Proposition 3 that is featured prominently in our results for finite and infinite horizon economies.

Proposition 4

Let $p > 0$ and $\{x, K\}$ be a constrained-efficient allocation. If $\{x, K\}$ is not efficient, then

$$c(s^t) \leq \bar{c} < \infty$$

for all nodes $s^t \in S^t$ and dates $t = 0, 1, \dots, T$.

Proof: In the Appendix. □

Note this stands in contrast with the standard Mirrleesian economy with $p = 0$: Individual consumption is usually unbounded above.

Proposition 3 states that the constrained-efficient allocation will have some surprising properties if the continuation welfare is large enough. However, we still do not know whether it is part of the constrained-efficient allocation to back-load welfare until the threshold in Proposition 3 is achieved. Doing so is not costless: It requires low initial consumption and, in general, a less smooth path for consumption. In the next sections we will characterize the conditions under which the random state verification channel dominates and the marginal return to investment is below its cost in constrained-efficient allocations.

6 Finite Horizon

In this section, we turn to finite horizon economies with an arbitrary time horizon. We already know the results for a special case. Section 3 discussed a two-period economy. In that section, we saw the size of intra- and intertemporal wedges for a particular calibration.

Besides the special case, we gathered a fair bit of information on the finite horizon economy. We know that the distortions for the labor supply decision have the same qualitative properties as in a standard dynamic Mirrleesian economy, as Section 4 shows. There, we also discussed the ranking of consumption and output allocations. Section 5 discussed the optimality condition for the intertemporal allocation. Compared with the standard Mirrleesian setting, the intertemporal wedge has an extra term that prevents us from determining the sign of the wedge directly.

This section sheds more light on under which conditions the intertemporal wedge is negative or positive. Therefore, we discuss two sets of results. First, for *any* finite horizon economy, the implied savings tax is negative if the probability of state verification is in a certain range. And second, for a given (potentially tiny) probability of state verification, there exists a time horizon T such that the implied savings tax is negative for some agents in the economy.

Let us turn to the first result. As we have seen in previous sections, a higher probability of state verification relaxes the incentive-compatibility constraints. We can get intuition by studying two extreme cases. On one end of the spectrum, no skill level gets observed. In this case, we retrieved the standard dynamic Mirrleesian setup for which we know that the intertemporal wedge is positive, and, hence, the implicit tax on savings is positive. At the other end of the spectrum, all skill levels become public information before allocations get assigned. Since there is no private information, no incentive-compatibility constraints have to be satisfied, and we recover the first-best allocation in which full risk sharing prevails.

For the two-period economy, however, we showed that the intertemporal wedge does not

decline monotonically. It rather turns into a subsidy region before it enters the range in which the first-best allocation can be delivered. The rationale for the subsidy region is given by the fact that for a higher probability of state verification, more savings relax incentive constraints for two reasons. First, the effective punishment in case of misreporting increases. But second, with higher probabilities of the skill becoming public information, every extra unit of effective punishment has a bigger effect on the expected punishment.

We now show that the existence of a subsidy is a general result that holds for any finite horizon economy of our setup. We therefore prove the following proposition.

Proposition 5

For any economy with time horizon of length $T < \infty$, there exists a $\hat{p} < p^$ such that the intertemporal wedge is negative for some node $s^t \in S^t$,*

$$u'(s^t) > E_t \{u'(s^{t+1})\}.$$

Proof: In the Appendix. □

The key to this result lies in the incentive compatibility constraints. Imagine an economy with the lowest probability of state verification that still supports the first-best outcome. The fact that the full consumption sharing can be implemented means that the effect of private information has been overcome. In essence, the effective punishment is strong enough to allow full risk sharing. That means, however, that the reason for front-loading also has been eliminated. If the probability of state verification is, however, just a slight bit smaller, the effective punishment channel calls for a subsidy to push agents toward the region of a first-best allocation. Since the effect of front-loading consumption is marginal, the intertemporal wedge becomes negative.

But is a subsidy an option for *any* given probability of state verification? Surprisingly, the answer is yes — provided that the time horizon is long enough. The next theorem shows that any economy with positive probability of skills becoming observable features a subsidy.

Theorem 1

Let $p > 0$. There exists $T < \infty$ such that the constrained-efficient allocation $\{x, K\}$ satisfies

$$u'(s^t) \geq E_t \{u'(s^{t+1})\}$$

for some node $s^t \in S^t, t \leq T - 1$.

Proof: In the Appendix. □

Intuitively, we know that for a given probability p , there is a level of savings such that you can give agents first-best allocations. Now the question remains of whether agents will ever

reach that level. Therefore, consider the relevant trade-offs in a given period. To induce agents to work, they need to be given consumption in return. However, this consumption does not have to be delivered in the current period. It can merely be promised to be delivered in future periods. The promise keeping constraint ensures that eventually consumption will be provided. But therefore, the economy has to build up enough savings, which in turn raises the effective punishment level. If the time horizon is long enough, a sufficient amount of savings guarantees that some agents have a nonpositive intertemporal wedge.

7 Infinite horizon economies

This section answers the question regarding which sign the intertemporal wedge takes on the long run. We find that *all* agents (but a set of mass zero) converge to full insurance and no intertemporal wedge. Even more interestingly, this result does not depend on the probability of state verification. Even a tiny chance of skill levels being revealed leads agents to converge to a nonpositive intertemporal wedge.

Let us be clear about the precise meaning of the concept of convergence. Our results concerning the long run are stated in terms of convergence in probability.¹² For example, we will say that consumption converges to some level \hat{c} , denoted $c(s^t) \rightarrow \hat{c}$, if for any $\delta > 0$

$$\lim_{t \rightarrow \infty} \sum_{s^t \in S^t} \pi(s^t) \chi [|c(s^t) - \hat{c}| \geq \delta] = 0$$

where χ is the indicator function. The long-run properties will thus hold for all agents with the possible exception of a subset with measure zero.

The following theorem summarizes the main result of this section.

Theorem 2

Let $p > 0$. The economy converges to full insurance and no labor distortions:

$$c(s^t) = c(s^{t+1})$$

and

$$v'(y(s^t)/\theta_t) = u'(c(s^t)) \theta_t$$

for all $\theta_t \in \Theta$.

Proof: In the Appendix. □

¹²See Stokey and Lucas (1996) for a formal treatment.

The rest of this section is concerned with shedding some more light on our result. We first establish the existence of an absorbing upper bound on promised welfare. Then we show that there is a non-absorbing lower bound, which means that agents close to that bound can reach higher promised welfare after good shock realizations. To establish the result of the theorem, the remaining piece is to show that the Lagrangian multiplier on the promised welfare constraint follows a submartingale. We then apply the submartingale convergence theorem to prove the result of Theorem 2. Note that the Lagrangian multiplier associated with the resource constraint (2) is constant due to constant marginal returns to capital. This assumption becomes important in the infinite horizon case because it allows the economy to converge to an arbitrarily large level of aggregate capital and thus defer the delivery of utility to the future.

As demonstrated in Proposition 3, there exists an upper threshold on promised utility such that full insurance without labor distortions can be implemented. The intuition for this upper bound goes as follows. The larger the punishment associated with misreporting, the more the state verification technology can relax the incentive-compatibility constraints. The continuation welfare prescribed in the event of misreporting is exogenously fixed at Υ . However, the right measure of punishment is the amount of “forgone” welfare, that is, the difference between the promised continuation welfare and the punishment. There is a level of continuation welfare z^* for which the threat of punishment is so large that the effects of private information are dominated and full insurance can be implemented.

Proposition 3 actually goes one step further and shows that full insurance with no labor distortions is a *necessary* property of the constrained-efficient allocation whenever promised utility is at or above the threshold for full insurance. Intuitively, for any level of promised utility, the cost-efficient solution to the promise keeping constraint is precisely to provide full insurance with no labor distortions. If $z(s^t) \geq z^*$, the incentive-compatibility constraints are no longer an impediment to do so. An immediate corollary of Proposition 3 is that $z(s^t) \geq z^*$ is an absorbing state since the promised continuation welfare will be constant under full insurance.

Corollary 1

Let $p > 0$. If $z(s^t) \geq z^*$, then $z(s^j) = z(s^t)$ for all $\{s^j : s^t \in s^j, j \geq t\}$.

So far, we have established that there exists an upper bound on continuation welfare and that it is an absorbing state. We now turn to showing that there is also a non-absorbing lower bound. By assumption, there is a lower bound on welfare (1). We briefly discuss how continuation welfare behaves close to the lower bound. We find that the lower bound on continuation welfare (1) may be binding with positive probability. However, *promised* continuation welfare is always strictly higher than the lower bound.

Lemma 2

Let $\{x, K\}$ be a constrained-efficient allocation and z be the associated plan of continuation

welfare. Then $z(s^t) > \Upsilon$.

Proof: In the Appendix. □

The key implication of Lemma 2 is that the lower bound is not absorbing. Instead, an agent hitting the lower bound suffers a very low per-period welfare but is promised continuation welfare above the lower bound.¹³

We now link the results above with the stochastic process for the multiplier $\eta(s^t)$ associated with the promise keeping constraint (3). The multiplier effectively governs the process for continuation welfare. An upper bound η^* can be derived from Proposition 4, and characterizes an absorbing state. There is a lower bound as well that is transient as shown in Lemma 2.

The key equation is the characterization of the one-step ahead expectation that is given by

$$E_t \{ \eta(s^{t+1}) \} = \eta(s^t) + p\mu^*(s^t).$$

As long as $p > 0$ and some incentive-compatibility constraint is binding, the multiplier $\eta(s^t)$ follows a submartingale process with a strictly positive drift. The process is bounded above and below as established by proposition 3 and Lemma 2. There is an absorbing state, as established by Corollary 1. These conditions allow us to apply the submartingale convergence theorem to show that the multiplier $\eta(s^t)$ converges to the upper bound, as does continuation welfare. Hence, in the long run, the economy features full insurance and no labor distortion.

Distortions introduced by private information are thus temporary in the sense that they vanish in the long run for all agents but a subset of measure zero. We saw in Proposition 3 that the economy can “outgrow” the incentive-compatibility constraints by promising higher continuation welfare. However, to sustain the welfare promises, it is necessary to accumulate aggregate capital, and this transition is costly.

Why then does the economy converge to full insurance independent of the initial conditions, the discount rate, and the state-verification probability? The key insight is that, under state verification, *back-loading* welfare can help to ease incentives. If we need to provide incentives to a high skill agent at date t to reveal her type, we can increase consumption allocated to high types at date t ; or we can increase the promised continuation welfare from date $t + 1$ onward. The latter is more efficient because it also improves the incentives at date $t + 1$ by increasing the amount of welfare forgone in case an agent being caught misreporting. Yet, there are other forces working in the opposite direction. First, concavity favors smooth consumption paths. Second, the rationale for front-loading welfare in the standard private information economies remains: To deliver higher promised welfare at date $t + 1$, the planner must prescribe a larger

¹³Our results closely mirror those in Atkeson and Lucas (1993).

consumption spread between skill levels to ensure that the high skill type truthfully reveals the agent's type.

As promised utility becomes larger, the punishment channel becomes more effective, and the private information channel becomes weaker because the dispersion of consumption will be lower. Hence, for very high promised utilities, there is essentially no trade-off.

Hence, in the long run, the benefits of back-loading welfare eventually dominate. The key feature of an infinite horizon economy is that every agent receives a series of high skill realizations of arbitrary (but finite) length with probability one. During this "streak," promised welfare is naturally increasing, and eventually the state verification effect dominates.

The result in Theorem 2 stands in stark contrast with the results in dynamic Mirrleesian economies without state verification. In these economies, aggregate consumption is decreasing. No agent ever achieves full insurance, and labor distortions are always present. It may seem puzzling that even an arbitrarily small state verification probability can lead to such different results. However, these are long-run properties: The welfare improvement at date $t = 0$ associated with a small state-verification probability may be small.

8 Extensions

8.1 Decreasing marginal returns to capital

The paper focused on the characterization of the constrained-efficient allocations for a linear production function. However, our results concerning the consumption and output distribution, as well as the finite horizon properties, hold for general constant returns to scale production functions.

The assumption of constant marginal returns to capital does play a key role in the long-run results in Theorem 2. For full insurance to be incentive compatible, the effective punishment must be high enough to induce truth-telling. This, in turn, requires that the economy must be able to accumulate enough capital in the long run to sustain a possibly large level of aggregate consumption.

In economies with decreasing marginal returns to capital, this may not be feasible. Consider a more general constant returns to scale production function and let $f(K_t)$ be output per effective unit of labor net of capital depreciation. Decreasing marginal returns to capital imply that f is a strictly concave function. If the marginal product of capital tends to zero as K_t grows large, then the steady-state aggregate consumption is maximized at capital level $K_t = \bar{K}$ with $f'(\bar{K}) = 1$. Aggregate consumption is bounded above but, yet, full insurance may still be feasible. Theorem 2 still applies if \bar{K} is high enough.

It is harder to venture a conjecture for the long-run properties when the upper bound \bar{K} is not high enough to sustain full insurance. In particular, we did not find any guarantee that there exists a stationary consumption distribution. No agent will enjoy full insurance for all continuing dates with probability one. However, there will be a subset of agents enjoying close to full insurance at any given date. This suggests that, if there exists a stationary distribution, it will be skewed to the left, with low consumption households facing higher consumption volatility than households with high average consumption.

8.2 Endogenous punishment level

In our setup, incentives improve with stronger punishments. We took the punishment level as exogenous to focus on the implications of the random state-verification technology. Here we briefly discuss possible foundations for a theory of punishment and their implications for our findings.

The easiest route to an endogenous punishment level is to specify an outside option for agents. One possibility is that an agent can always opt out from all trading and savings, as in the limited commitment models of Kehoe and Levine (1993) and Kocherlakota (1996). Thus the planner cannot prescribe a lower level of continuation utility than what the agent would get under autarky. The punishment level for an agent with type s^t is then given by

$$\Upsilon(s^t) = \max_{y(s^j): s^j \in S^j | s^t} \sum_{j=t}^{\infty} \sum_{\theta_j \in \Theta} \beta^{j-t} \pi(\theta_j | \theta_t) (u(y(s^j)) - v(y(s^j)/\theta_j)).$$

The punishment level given by autarky does not depend on the full history: Only the present realization of skill level θ_t matters. Importantly, none of the allocations prescribed on the equilibrium path change the value of the punishment. In this case, the first-order conditions we derived are literally unchanged, and all of our results remain valid.

It is indeed important whether the threat of punishment depends on equilibrium-path allocations, or it is purely an off-equilibrium threat. Our results rely on the ability of the planner to strengthen the effective punishment by increasing continuation welfare. If the punishment level scales up with continuation welfare, our results may not hold.¹⁴ For example, if we prescribe autarky as punishment but we allow the agent to hold onto a share of aggregate savings, then the effective punishment may *decrease* with savings, reversing our arguments. Similarly, if the punishment is modeled as fixed drop in utility or is limited to a monetary amount as in Ravikumar and Zhang (2009), the effective punishment becomes constant or decreasing with savings. However, realistic specifications of punishment such as a certain time in prison will have the

¹⁴For this to be the case, on-equilibrium allocations must also influence off-equilibrium threats, which is the key channel shut down by our assumption of an exogenous punishment level.

property that threats increase with savings. We are confident that a satisfactory theory would allow for the effective punishment to increase with continuation welfare.

Imperfect monitoring provides an alternative theory of punishment. In this environment, the planner observes the agent's type imperfectly and might mistakenly take an agent for being of a different type. To provide incentives, the planner prescribes some punishment for agents suspected of misreporting, even though only truth-telling strategies prevail along the equilibrium path. The planner now may want to soften the punishment to provide some insurance against monitoring errors.¹⁵ Clearly, full insurance in infinite horizon economies is no longer possible, as all agents will be eventually punished with probability one.

In a two-period economy with imperfect monitoring, we found that the prescribed punishment is arbitrarily large if the equivocation error is small. This suggests that if period payoffs are bounded below, the lower bound would be binding, in which case the imperfect monitoring economy is equivalent to the economy we studied.

8.3 Towards a theory of optimal monitoring

Our previous analysis takes the probability of state verification as exogenously given, and thus out of the planner's control. The specification allows us to study a full-fledged dynamic private information economy and compare our results with those in the literature. However tax authorities and insurance companies regularly choose who and how often to monitor. Thus we would ideally want to study allocations under the *optimal* state-verification policy. Unfortunately the joint characterization of optimal monitoring probabilities and an allocations is very difficult and an exhaustive analysis beyond the scope of this paper.¹⁶

Nevertheless, we are able to show some key properties of the optimal state-verification policy. Since our previous results hold for any small but strictly positive state-verification probability, we would like to know whether it is ever optimal to stop monitoring agents whose incentive compatibility constraint is binding.

We briefly describe the setting for an optimal state-verification policy. An agent with history s^{t-1} who reports type θ_t will be monitored with probability $p(s^{t-1}, \theta_t)$. Let

$$\mathbf{p} = \{p(s^{t-1}, \theta_t) \in [0, 1] : (s^{t-1}, \theta_t) \in S^{t-1} \times \Theta\}$$

be a state-verification policy and P the associated set of state-verification policies. We specify

¹⁵The results in Abreu et al. (1990) suggest that the optimal allocation will still rely on the maximum punishment available, perhaps using lotteries.

¹⁶Recently several researchers have made significant progress by focusing on different environments. For example Ravikumar and Zhang (2009) can characterize the optimal audit schedule assuming an absorbing state for the income process.

a simple linear cost function for state-verification,

$$M(\mathbf{p}) = \kappa \sum_{t=0}^T \beta^t \sum_{s^{t-1}, \theta_t} \pi(s^{t-1}, \theta_t) p(s^{t-1}, \theta_t),$$

where $\kappa > 0$. The cost is denominated in units of the final good and thus enters the resource constraint directly. If monitoring is very costly, i.e., if κ is sufficiently large, no monitoring will take place, i.e. $p_t = 0$ at all dates. Instead we are interested in whether the monitoring probability is bounded away from zero for low but positive monitoring costs.¹⁷

We turn to the trade-offs the planner faces when choosing the optimal probability of state-verification. To gather some intuition for our results, consider the marginal benefit of increasing the monitoring probability for an agent with history s^{t-1} and reporting type θ^t ,

$$\Gamma(s^{t-1}, \theta^t) = \sum_{\theta'} \mu(s^{t-1}, \theta'; \theta_t) \frac{\pi(\theta')}{\pi(\theta_t)} (h(x(s^t); \theta') + z(s^t) - (1 - \beta^{T+1-t})\Upsilon) \quad (9)$$

where $s^t = \{s^{t-1}, 0, \theta_t\}$. Function $\Gamma(s^{t-1}, \theta^t)$ does not include the terms arising due to the rewards assigned to monitored agents—in the Appendix we show that these additional terms are always positive and thus $\Gamma(s^{t-1}, \theta^t)$ is a lower bound on the marginal benefit.

The terms in $\Gamma(s^{t-1}, \theta^t)$ have a straightforward interpretation. Increasing the monitoring probability decreases the expected gain of misreporting its own type as θ_t . Thus the marginal benefit sums up the effect on the incentive-compatibility constraints that ensure that no type θ' reports θ_t . Of course if any of them is not binding, $\mu(s^{t-1}, \theta'; \theta_t) = 0$, there is no marginal benefit for that particular constraint. The term $h(x(s^t); \theta') + z(s^t) - (1 - \beta^{T+1-t})\Upsilon$ is the difference between getting away with misreporting and the punishment. Since Υ effectively bounds period utility from below, the term is always positive.

Let us now turn to the question of whether the optimal monitoring probability is positive. The results from our previous analysis hold up whenever we find conditions in which the probability exceeds zero. We analyze several cases by restricting the set of types that the planner can discriminate over.

We start with the simplest possible case, a constant monitoring probability $p(s^{t-1}, \theta_t) = p$. We evaluate $\Gamma(s^{t-1}, \theta^t)$ at the optimal allocations corresponding to $p = 0$ which results in

$$\sum_{t=0}^{\infty} \beta^t \sum_{s^{t-1}, \theta_t} \pi(s^{t-1}, \theta_t) \Gamma(s^{t-1}, \theta^t) > 0$$

¹⁷For the sake of exposition we do not include a full definition of an optimal monitoring probability and the associated second best allocations.

since the first-best allocation is never admissible with $p = 0$ and thus the incentive-compatibility constraints are binding. Thus there exists a low enough $\kappa > 0$ such that a positive probability, $p > 0$, is optimal.

Consider next the case of time-dependent monitoring probabilities, $p(s^{t-1}, \theta_t) = p_t$. Importantly, we can rule out that the probability of monitoring will ever eventually be zero. If at any date d the monitoring probability was zero for all periods thereafter then the economy would be identical to a standard dynamic Mirrleesian economy from then onwards. Incentive-compatibility constraints will be binding for all agents but those with the highest skill. Thus

$$\sum_{s^{t-1}, \theta_t} \pi(s^{t-1}, \theta_t) \Gamma(s^{t-1}, \theta^t) > 0$$

at all periods $t \geq d$. Hence it is not optimal to set the monitoring probability to zero in finite time.¹⁸

Finally we consider the most general case of history-dependent monitoring probabilities, $p(s^{t-1}, \theta_t)$. We highlight two properties. First, the probabilities will depend on the reported type. From $\Gamma(s^{t-1}, \theta^t)$ it is clear that there is no need to monitor any agent that reports to have the highest skill level, $\bar{\theta}$, since no agent would want to pretend as having the highest skill. In other words, $\Gamma(s^{t-1}, \bar{\theta}) = 0$ since $\mu(s^{t-1}, \theta'; \bar{\theta}) = 0$ for all $\theta' \in \Theta$. Fortunately our results go through if we exclude agents of type $\bar{\theta}$ from monitoring. The second property of notice is history-dependence: probabilities will generally depend on the entire history of types. It is difficult to provide a general characterization for this case. We instead compute the optimal history-dependent probability for a three-period economy. First, if the planner can condition on past skill realizations only, we find almost exactly constant probabilities of state verification across types and time¹⁹. Again, only the agent reporting low skills are being monitored. When the planner can condition on the full type, i.e., skill and whether the agent has been monitored in the past, we find that past monitoring leads to lower future probabilities of state verification.

9 Conclusions

We set out to take a first step toward a more realistic information framework for dynamic Mirrleesian economies. From observing the implementation of social insurance schemes and the prevention of fraud, we introduced an exogenous state verification technology. The skill level of any agent becomes public information with some given probability. The planner can specify

¹⁸Unfortunately we cannot say much more about the sequence $\{p_t\}$ of optimal monitoring probabilities. In particular, the optimal monitoring probability may not converge at all.

¹⁹Our preferred calibration uses logarithmic utility over consumption, quadratic disutility of labor, $\theta_i = i$ for $i = 1, \dots, 3$, and $\Upsilon = -5$.

allocations as a function of the reported skill level, whether the skill level is observable and, if so, the actual skill level.

Our results establish that state verification can radically shape the constrained efficient allocation. In particular in the infinite horizon economy with constant returns to capital, we discover very strong predictions. We find that even an arbitrarily small probability of state verification implies that the economy converges to full insurance, and all labor distortions vanish in the long run. In finite horizon economies we show that different agents at different times will generally have different incentives to save, and a decentralization would require that savings are sometimes taxed and sometimes subsidized.

References

- Abreu, Dilip, David Pearce, and Ennio Stacchetti**, “Toward a Theory of Discounted Repeated Games with Imperfect Monitoring,” *Econometrica*, September 1990, 58 (5), 1041–1063.
- Albanesi, Stefania and Christopher Sleet**, “Dynamic Optimal Taxation with Private Information,” *The Review of Economic Studies*, 2006, 73, 1–30.
- and **Roc Armenter**, “Intertemporal Distortions in the Second Best,” *Working Paper*, Columbia University, 2008.
- Atkeson, Andrew and Robert E. Lucas**, “On Efficient Distribution with Private Information,” *The Review of Economic Studies*, 1993, 59 (3), 427–453.
- Border, Kim C. and Joel Sobel**, “Samurai Accountant: A Theory of Auditing and Plunder,” *The Review of Economic Studies*, 1987, 54 (4), 525–540.
- Buckland, William Warwick**, *A Textbook of Roman Law from Augustus to Justinian*, Cambridge University Press, 1975.
- Chamley, Christophe**, “Optimal Taxation of Capital Income in General Equilibrium with Infinite Lives,” *Econometrica*, 1986, 54 (3), 607–622.
- Cremer, Helmuth and Firouz Gahvari**, “Tax Evasion and the Optimum General Income Tax,” *Journal of Public Economics*, 1995, 60, 235–249.
- Farhi, Emmanuel and Ivan Werning**, “Capital Taxation: Quantitative explorations of the inverse Euler equation,” *MIT Working Paper 06-15*, 2007.
- Golosov, Mikhail, Aleh Tsyvinski, and Ivan Werning**, “New Dynamic Public Finance: A User’s Guide,” *NBER Macroeconomics Annual*, 2006.
- and – , “Designing Optimal Disability Insurance: A Case for Asset Testing,” *Journal of Political Economy*, 2006, 114 (2), 257–279.
- , **Narayana Kocherlakota, and Aleh Tsyvinski**, “Optimal Indirect and Capital Taxation,” *The Review of Economic Studies*, 2003, 70 (3), 569–587.
- Judd, Kenneth L.**, “Redistributive Taxation in a Perfect Foresight Model,” *Journal of Public Economics*, 1985, 28, 59–84.
- Kehoe, Timothy J. and David K. Levine**, “Debt-Constrained Asset Markets,” *Review of Economic Studies*, 1993, 60(4), 865–88.

- Kocherlakota, Narayana**, “Implications of Efficient Risk Sharing without Commitment,” *Review of Economic Studies*, 1996, *63* (4), 595–609.
- , “Expected Wealth Taxes: A Mirrlees Approach to Dynamic Optimal Taxation,” *Econometrica*, 2005, *73* (5), 1587–1621.
- Mankiw, N. Gregory and Matthew C. Weinzierl**, “The Optimal Taxation of Height: A Case Study of Utilitarian Income Redistribution,” *NBER Working Paper*, w14976, 2009.
- Mirrlees, James**, “An Exploration in the Theory of Optimum Income Taxation,” *The Review of Economic Studies*, 1971, *38*, 175–208.
- Monnet, Cyril and Erwin Quintin**, “Optimal Contracts in a Dynamic Costly State Verification Model,” *Economic Theory*, 2005, *26*, 867–885.
- Popov, L.**, “Stochastic Costly State Verification and Dynamic Contracts,” *Working paper*, University of Iowa, 2007.
- Ravikumar, B. and Yuzhe Zhang**, “Optimal Auditing and Taxation in a Dynamic Model of Tax Compliance,” *Working Paper*, University of Iowa, 2009.
- Rogerson, William**, “Repeated moral hazard,” *Econometrica*, 1985, *53*, 69–76.
- Salanie, Bernard**, *The economics of Taxation*, The MIT Press, Cambridge, MA, 2003.
- Stigler, G.**, “The Optimum Enforcement of Laws,” *Journal of Political Economy*, 1970, *78*, 526–536.
- Stokey, Nancy and Robert Lucas**, *Recursive Methods in Economic Dynamics*, Harvard University Press, Cambridge, MA, 1996.
- Townsend, Rob**, “Optimal Contracts and Competitive Markets with Costly State Verification,” *Journal of Economic Theory*, 1979, *21*, 265–293.
- Tuomala, Matti**, *Optimal Income Tax and Redistribution*, Oxford University Press, 1990.
- Wang, C.**, “Dynamic Costly State Verification,” *Economic Theory*, 2005, *25*, 88–916.
- Weinzierl, Matthew C.**, “The Surprising Power of Age-Dependent Taxation,” *mimeo*, Harvard University, 2008.
- Werning, Ivan**, “Optimal Fiscal Policy with Redistribution,” *The Quarterly Journal of Economics*, 2007, *122* (3), 925–967.

A Proofs for Section 4

We repeat the optimal control problem of the main text. We maximize

$$\sup_{\{x, K, z\}} W_0(x)$$

subject to the incentive-compatibility and promise keeping constraints (4) and (3) at all nodes $s^t \in S^t$, and the resource constraint (2) at all dates $t = 0, 1, \dots, T$, as well as nonnegativity constraints and the lower bound constraint on continuation welfare (1).

By Assumption 1, $h(x(s^t); \theta_t)$ is twice differentiable in \mathbb{R}_+^2 for all $\theta_t \in \Theta$. We can thus characterize the constrained-efficient allocation with the Lagrangian

$$\begin{aligned} \mathcal{L}(x, K, z; \lambda, \eta, \mu) = & W_0(x) - \sum_{t=0}^T \beta^t \lambda_t \{C_t + K_{t+1} - F(K_t, Y_t)\} \\ & - \sum_{t=0}^{T-1} \sum_{s^t \in S^t} \beta^t \pi(s^t) \eta(s^t) \left\{ z(s^t) - \sum_{s^{t+1} \in S^{t+1}} \pi(s^{t+1} | s^t) \{h(x(s^{t+1}); \theta_{t+1}) + \beta z(s^{t+1})\} \right\} \\ & - \sum_{t=0}^T \sum_{s^t \in S^t} \beta^t \pi(s^t) \zeta(s^t) \{(1 - \beta^{T+1-t}) \Upsilon - h(x(s^t); \theta_t) - \beta z(s^t)\} \\ & - \sum_{t=0}^T \sum_{s^{t-1} \in S^{t-1}} \sum_{\theta_t, \theta' \in \Theta} \beta^t \pi(s^{t-1}, \theta_t) \mu(s^{t-1}, \theta_t; \theta') \left\{ \sum_{a_t \in \{0,1\}} \pi(a_t) \{h(x(s^t); \theta_t) + \beta z(s^t)\} \right. \\ & \left. - p(1 - \beta^{T+1-t}) \Upsilon - (1-p) \{h(x(\hat{s}^t); \theta') + \beta z(\hat{s}^t)\} \right\}. \end{aligned}$$

We have omitted from the Lagrangian the nonnegativity constraints on consumption and output, as well as the terminal condition on $z(s^T)$ for finite horizon economies.

Next, we state the necessary first-order conditions. For nodes s^t when the state remain

private information, $a_t = 0$, we have:

$$u'(s^t) \left(1 + \eta(s^{t-1}) + \sum_{\theta' \in \Theta} \left\{ \mu(s^{t-1}, \theta_t; \theta') - \frac{\pi(\theta')}{\pi(\theta_t)} \mu(s^{t-1}, \theta'; \theta_t) \right\} \right) \leq \lambda_t, \quad (10)$$

$$h^y(s^t; \theta_t) \left(1 + \eta(s^{t-1}) + \sum_{\theta' \in \Theta} \left\{ \mu(s^{t-1}, \theta_t; \theta') - \frac{\pi(\theta')}{\pi(\theta_t)} \mu(s^{t-1}, \theta'; \theta_t) \frac{h^y(s^t; \theta')}{h^y(s^t; \theta_t)} \right\} \right) \geq \lambda_t, \quad (11)$$

$$\eta(s^{t-1}) + \sum_{\theta' \in \Theta} \left\{ \mu(s^{t-1}, \theta_t; \theta') - \frac{\pi(\theta')}{\pi(\theta_t)} \mu(s^{t-1}, \theta'; \theta_t) \right\} \leq \eta(s^t). \quad (12)$$

For nodes when the skill level is observable, $a_t = 1$, we have:

$$u'(s^t) \left(1 + \eta(s^{t-1}) + \sum_{\theta' \in \Theta} \mu(s^{t-1}, \theta_t; \theta') \right) \leq \lambda_t, \quad (13)$$

$$h^y(s^t; \theta_t) \left(1 + \eta(s^{t-1}) + \sum_{\theta' \in \Theta} \mu(s^{t-1}, \theta_t; \theta') \right) \geq \lambda_t, \quad (14)$$

$$\eta(s^{t-1}) + \sum_{\theta' \in \Theta} \mu(s^{t-1}, \theta_t; \theta') \leq \eta(s^t). \quad (15)$$

In addition f.o.n.c. on aggregate capital is:

$$\lambda_t = \lambda_{t+1} \beta R, \quad (16)$$

for all dates $t < T$. Given our assumption that $\beta R = 1$, we have that $\lambda_t = \lambda_{t+1}$ for all dates $t < T$.

The next Proposition guarantees that the constrained-efficient allocation is interior at all nodes.

Proposition 6

Let $\{x, K\}$ be a constrained-efficient allocation. Then for all $s^t \in S^t, t = 0, 1, \dots, T$

$$\begin{aligned} c(s^t) &\geq c^b > 0 \\ y(s^t) &> 0 \end{aligned}$$

Proof: By (1) any admissible allocation must satisfy

$$u(c(s^t)) \geq u(c(s^t)) - v \left(\frac{y(s^t)}{\theta_t} \right) \geq v.$$

Assumptions 1 and 3 guarantee there exists $c^b > 0$ that achieves this lower bound, and thus for all admissible allocations $c(s^t) \geq c^b > 0$. Next we show $y(s^t) > 0$ if $\{x, K\}$ is a constrained efficient allocation. Assume $y(s^t) = 0$ for any node. By Assumption 1, $v'(0) = 0$. The necessary first order conditions associated with the Lagrangian imply that $0 \geq \lambda_t F_t^y$, and thus we reach a contradiction. \square

Proof: [Proof of Proposition 1] Our preferences satisfy the single-crossing condition and thus $\mu(s^{t-1}, \theta'; \theta_t) = 0$ for $\theta' < \theta_t$. Using the strict convexity of h ,

$$\mu(s^{t-1}, \theta'; \theta_t) h^y(s^t; \theta') \leq \mu(s^{t-1}, \theta'; \theta_t) h^y(s^t; \theta_t).$$

Combining (10)-(11) and using the inequality above we obtain the result. \square

Proof: [Proof of Proposition 2] By Proposition 6 the system of f.o.n.c. (10)-(11) and (13)-(14) holds with strict equality. Comparing (10) and (13), it is clear that

$$u'(s^{t-1}, \theta_t, 0) \geq u'(s^{t-1}, \theta_t, 1)$$

with strict equality iff $\mu(s^{t-1}, \theta'; \theta_t) = 0$ for all $\theta' \in \Theta$. Since u is strictly concave, the result on consumption follows. A similar argument applies for output, now comparing (11) and (14). To show that $h^y(s^t; \theta_t) = u'(s^t) F_t^y$ in case $a_t = 1$ combine (13)-(14). \square

B Proofs for Section 5

Proof: [Proof of Proposition 3]

The proof proceeds as follows. First, we construct a candidate threshold z^* . Then we show that, for any promised utility $z(s^{t-1}) \geq z^*$, there exists an allocation attaining the promised welfare $z(s^{t-1})$ that satisfies all the incentive-compatibility constraints and features full insurance and no labor distortions. Finally, we note that the proposed allocation solves an unconstrained optimization problem over continuation nodes and thus must be part of the constrained-efficient allocation.

For any $c^e > 0$, let the system of functions $\{y^e(c^e; \theta) : \theta \in \Theta\}$ solve

$$v' \left(\frac{y^e(c^e; \theta)}{\theta} \right) = u^c(c^e) \theta.$$

By Assumption 1 output $y^e(c^e; \theta)$ is strictly decreasing with c^e . Let

$$h^e(c^e; \theta, \theta') = u(c^e) - v\left(\frac{y^e(c^e; \theta')}{\theta}\right).$$

Clearly $h^e(c^e; \theta, \theta)$ is increasing in c^e as well for any θ . Consider now the static payoff of misreporting θ' when the true type is θ , $h^e(c^e; \theta, \theta') - h^e(c^e; \theta, \theta)$. Since consumption is constant, the payoff is bounded above,

$$h^e(c^e; \theta, \theta') - h^e(c^e; \theta, \theta) \leq v\left(\frac{y^e(c^e; \theta)}{\theta}\right).$$

The upper bound $v\left(\frac{y^e(c^e; \theta)}{\theta}\right)$ is decreasing in c^e . Hence, there exists c^* such that for all $c^e \geq c^*$,

$$ph^e(c^e; \theta, \theta) \geq p\Upsilon + (1-p)(h^e(c^e; \theta, \theta') - h^e(c^e; \theta, \theta)) \quad (17)$$

for all θ, θ' . The candidate threshold is

$$z^* = \frac{1}{1-\beta} \sum_{\theta \in \Theta} \pi(\theta) h^e(c^*; \theta, \theta).$$

For any $z(s^{t-1})$ there exists a consumption level c^e such that

$$z(s^{t-1}) = \frac{1 - \beta^{T+1-t}}{1 - \beta} \sum_{\theta \in \Theta} \pi(\theta) h^e(c^e; \theta, \theta),$$

since h^e is increasing in c^e . Hence, if $z(s^{t-1}) \geq z^*$, then the solution above satisfies $c^e \geq c^*$. Consider an allocation over the continuation nodes $\{s^j \in S^j : s^t \in s^{t-1}, j = t, t+1, \dots, T\}$ with $c(s^t) = c^e$ and output levels as given by y^e . Note that the allocation prescribes the same continuation welfare for all types at a given date, $d = t, t+1, \dots, T$,

$$z_d^e = \frac{1 - \beta^{T+1-d}}{1 - \beta} \sum_{\theta \in \Theta} \pi(\theta) h^e(c^e; \theta, \theta).$$

Now we show that such an allocation satisfies the incentive-compatibility constraints over all continuation nodes. Re-arrange terms in (17) for any pair θ, θ' and date $d = t, t+1, \dots, T$ to obtain

$$\begin{aligned} h^e(c^e; \theta, \theta) &\geq p\Upsilon + (1-p)h^e(c^e; \theta, \theta') \\ &\geq p\left(1 - \beta^{T+1-d}\right)\Upsilon + (1-p)h^e(c^e; \theta, \theta'). \end{aligned}$$

Adding βz_{d+1}^e to both sides results in the incentive-compatibility constraint for pair θ, θ' and

date d .

Finally, note that the proposed allocation over continuation nodes is the solution to the unconstrained problem:

$$\min_{\{\hat{c}, \hat{y}\}} \sum_{j=t}^T \sum_{s^j \in S^j | s^{t-1}} R^{-(j-t)} \pi(s^j) \{ \hat{c}(s^j) - \hat{y}(s^j) \}$$

subject to

$$\sum_{j=t}^T \sum_{s^j \in S^j | s^{t-1}} \beta^{(j-t)} \pi(s^j) h(\hat{c}(s^j), \hat{y}(s^j); \theta_j) \geq z(s^{t-1})$$

and nonnegativity constraints. Thus, the constrained-efficient allocation must be equal to the proposed allocation over the continuation nodes

□

Proof: [Proof of Proposition 4] For any $c^e > 0$, let the system of functions $\{y^e(c^e; \theta) : \theta \in \Theta\}$ solve

$$v' \left(\frac{y^e(c^e; \theta)}{\theta} \right) = u^c(c^e) \theta,$$

and let

$$h^e(c^e; \theta) = u(c^e) - v \left(\frac{y^e(c^e; \theta)}{\theta} \right).$$

Clearly $h^e(c^e; \theta)$ is strictly increasing in c^e for all θ . Pick \bar{c} such that

$$\beta \pi(\bar{\theta}) h^e(\bar{c}; \bar{\theta}) = z^* - (1 + \beta) \Upsilon, \tag{18}$$

where z^* is the threshold value given by Proposition 3.

Let $\{x, K\}$ be a constrained-efficient allocation and assume that there exists s^t such that $c(s^t) > \bar{c}$. We take $\theta_t = \bar{\theta}$ without loss of generality as $c(s^{t-1}, \bar{\theta}) \geq c(s^t)$ by the single-crossing condition. If $t \leq 1$, then (18) implies that $W_0(x) \geq z^*$ by construction. Proposition 3 then concludes that $\{x, K\}$ is efficient. If $t > 1$, then (18) implies that $z(s^{t-2}) \geq z^*$ by construction. Proposition 3 then implies that $c(s^{t-1}) = c(s^t) \geq \bar{c}$. By induction, we conclude that $\{x, K\}$ is efficient. Thus if $\{x, K\}$ is not efficient then we must have $c(s^t) \leq \bar{c}$. □

Propositions 6 and 4, together with the characterization for consumption

$$u'(s^t) (1 + \eta(s^t)) = \lambda_t, \tag{19}$$

imply that the Lagrangian $\eta(s^t)$ is bounded above and below. We state this result as a simple

Lemma that comes in use in Theorems 1 and 2.

Lemma 3

Let $p > 0$ and $\{x, K\}$ be a constrained-efficient allocation and $\{\eta(s^t) : s^t \in S^t, t \geq 0\}$ the associated Lagrangian multiplier on the promise keeping constraint (3). Then

$$-1 < B_0 \leq \eta(s^t) \leq B_1 < \infty$$

for all $s^t \in S^t, t \geq 0$.

Proof: The lower bound follows from Proposition 6. Since $c(s^t) \geq c^b > 0$, f.o.n.c. (19) implies

$$\eta(s^t) = \frac{\lambda}{u'(c(s^t))} - 1 \geq \frac{\lambda}{u'(c^b)} - 1 > -1.$$

For the upper bound, we distinguish two cases. If $\{x, K\}$ is efficient, then consumption is constant at all nodes, and by (19) it follows that $\eta(s^t)$ is also constant and finite. If $\{x, K\}$ is not efficient, then $c(s^t) \leq \bar{c} < \infty$ by Proposition 4. Thus by (19)

$$\eta(s^t) = \frac{\lambda}{u'(c(s^t))} - 1 \leq \frac{\lambda}{u'(\bar{c})} < \infty.$$

□

C Proofs for Section 6

Proof: [Proof of Proposition 5] We lay out a constructive proof for the value of \hat{p} . For a given node s^t , there exists the lowest level for the probability of state verification $p^* \in (0, 1]$ such that the first-best allocation can be implemented from then on for all values $p \in [p^*, 1]$. At the cut-off value p^* , all Lagrangian multipliers on the incentive-compatibility conditions $\mu(s^t, \theta, \theta')$ equal zero for all θ and θ' . This result follows directly from the fact that at p^* , the maximization problem of the social planner with and without incentive-compatibility constraints leads to the same outcome.

We write the intertemporal wedge as an implicit tax (i.e. wedge) in the intertemporal Euler equation

$$\tau = 1 - \frac{1 + \frac{p\mu^*(s^t)}{1+\eta(s^t)}}{E[u'(s^{t+1})]E[\frac{1}{u'(s^{t+1})}]} = 1 - \frac{u'(s^t)}{E[u'(s^{t+1})]} = 1 - \frac{\frac{\lambda_t}{1+\eta(s^t)}}{E[u'(s^{t+1})]}$$

and compute the derivative with respect to the Lagrangian multiplier $\mu(s^t, \theta, \theta')$ to get

$$\frac{\partial \tau}{\partial \mu(s^t, \theta, \theta')} = \frac{\frac{\lambda_t}{1+\eta(s^t)}}{E[u'(s^{t+1})]^2} \cdot \frac{\partial E[u'(s^{t+1})]}{\partial \mu(s^t, \theta, \theta')}.$$

As we see from the first-order conditions directly, expected marginal utility is given by

$$E[u'(s^{t+1})] = \lambda_{t+1} \left[p \sum_{\theta} \frac{\pi(\theta)}{1 + \eta(s^t) + \sum_{\theta'} \mu(s^t, \theta, \theta')} \right. \\ \left. + (1-p) \sum_{\theta} \frac{\pi(\theta)}{1 + \eta(s^t) + \sum_{\theta'} \left\{ \mu(s^t, \theta, \theta') - \frac{\pi(\theta')}{\pi(\theta)} \mu(s^t, \theta', \theta) \right\}} \right]$$

We can compute the derivative as

$$\frac{\partial E[u'(s^{t+1})]}{\partial \mu(s^t, \theta, \theta')} = \lambda_{t+1} \left[-p \frac{\pi(\theta)}{(1 + \eta(s^t) + \sum_{\theta'} \mu(s^t, \theta, \theta'))^2} \right. \\ \left. + (1-p) \left[-\frac{\pi(\theta)}{\left(1 + \eta(s^t) + \sum_{\theta'} \left\{ \mu(s^t, \theta, \theta') - \frac{\pi(\theta')}{\pi(\theta)} \mu(s^t, \theta', \theta) \right\}\right)^2} \right. \right. \\ \left. \left. + \frac{\pi(\theta') \frac{\pi(\theta)}{\pi(\theta')}}{\left(1 + \eta(s^t) + \sum_{\theta'} \left\{ \mu(s^t, \theta, \theta') - \frac{\pi(\theta')}{\pi(\theta)} \mu(s^t, \theta', \theta) \right\}\right)^2} \right] \right]$$

Hence, expected marginal utility decreases with higher Lagrangian multipliers $\mu(s^t, \theta, \theta')$

$$\left. \frac{\partial E[u'(s^{t+1})]}{\partial \mu(s^t, \theta, \theta')} \right|_{\vec{\mu}=0} = -\lambda_{t+1} p \frac{\pi(\theta)}{(1 + \eta(s^t))^2}.$$

Furthermore, we need to make sure that the effect of the subsidy does not vanish with higher η . Therefore, we compute

$$\frac{\partial \tau}{\partial \eta(s^t)} = \frac{\frac{\lambda_t}{(1+\eta(s^t))^2}}{(E[u'(s^{t+1})])^2} \left(E[u'(s^{t+1})] + (1 + \eta(s^t)) \frac{\partial E[u'(s^{t+1})]}{\partial \eta(s^t)} \right)$$

where

$$\frac{\partial E[u'(s^{t+1})]}{\partial \eta(s^t)} = \lambda_{t+1} \left[-p \sum_{\theta} \frac{\pi(\theta)}{(1 + \eta(s^t) + \sum_{\theta'} \mu(s^t, \theta, \theta'))^2} \right. \\ \left. - (1-p) \sum_{\theta} \frac{\pi(\theta)}{(1 + \eta(s^t) + \sum_{\theta'} \left\{ \mu(s^t, \theta, \theta') - \frac{\pi(\theta')}{\pi(\theta)} \mu(s^t, \theta', \theta) \right\})^2} \right]$$

Evaluating this expression at $\vec{\mu} = 0$, we get

$$\left. \frac{\partial E[u'(s^{t+1})]}{\partial \eta(s^t)} \right|_{\vec{\mu}=0} = \frac{\frac{\lambda_t}{(1+\eta(s^t))^2}}{E[u'(s^{t+1})]^2} \left[E[u'(s^{t+1})] - \frac{\lambda_{t+1}}{1 + \eta(s^t)} \right] = 0$$

The intertemporal wedge turns negative once $\mu(s^t, \theta, \theta')$ becomes slightly positive:

$$\frac{\partial \tau}{\partial \mu(s^t, \theta, \theta')} \leq -p \frac{\lambda_t}{E[u'(s^{t+1})]} \frac{\pi(\theta)}{(1 + B_1)^2} < 0$$

where B_1 denotes the upper bound of the Lagrangian multiplier given by Lemma 3.

We can now construct a probability at which the intertemporal wedge becomes negative. Since p^* is the lowest level of probabilities that implements the first-best allocation, at least some multiplier $\mu(s^t, \theta, \theta')$ must be strictly positive for a probability of $\hat{p} = p^* - \varepsilon$. From the above reasoning, it then follows that for this ε the intertemporal wedge $\tau(\hat{p})$ becomes negative. \square

Proof: [Proof of Theorem 1] The proof has two parts. First, we show that if

$$u'(s^t) < E_t \{u'(s^{t+1})\}$$

at some node s^t , then for some pair θ, θ' we have $\mu(s^t, \theta, \theta') \geq \bar{\mu} > 0$. For the second part we proceed by contradiction: If there does not exist a finite T such that $u'(s^t) \geq E_t \{u'(s^{t+1})\}$ for some node, then $\eta(s^t)$ would violate the upper bound imposed by Lemma 3.

Let the intertemporal wedge be given by

$$\tau(s^t) = 1 - \frac{u'(s^t)}{E[u'(s^{t+1})]}.$$

Using the f.o.n.c. we have

$$\begin{aligned} \tau(s^t) &= 1 - \frac{1 + \frac{p\mu^*(s^t)}{1+\eta(s^t)}}{E[u'(s^{t+1})]E[\frac{1}{u'(s^{t+1})}]} \\ &= 1 - \frac{\frac{\lambda_t}{1+\eta(s^t)}}{E[u'(s^{t+1})]}. \end{aligned}$$

We differentiate the wedge with respect to some $\mu(s^t, \theta, \theta')$:

$$\frac{\partial \tau}{\partial \mu(s^t, \theta, \theta')} = \frac{\frac{\lambda_t}{1+\eta(s^t)}}{\beta F_{t+1}^K (E[u'(s^{t+1})])^2} \cdot \frac{\partial E[u'(s^{t+1})]}{\partial \mu(s^t, \theta, \theta')}$$

Using the f.o.n.c, we write the expected marginal utility as

$$\begin{aligned} \frac{E[u'(s^{t+1})]}{\lambda_{t+1}} &= p \sum_{\theta} \frac{\pi(\theta)}{1 + \eta(s^t) + \sum_{\theta'} \mu(s^t, \theta, \theta')} \\ &+ (1-p) \sum_{\theta} \frac{\pi(\theta)}{1 + \eta(s^t) + \sum_{\theta'} \left\{ \mu(s^t, \theta, \theta') - \frac{\pi(\theta')}{\pi(\theta)} \mu(s^t, \theta', \theta) \right\}}. \end{aligned}$$

For any node s^t with $t < T$, and pair of types θ_i, θ_j , we can compute the derivative as

$$\begin{aligned} \frac{\partial E[u'(s^{t+1})]}{\partial \mu(s^t, \theta_i, \theta_j)} &= -p \frac{\lambda_{t+1} \pi(\theta_i)}{(1 + \eta(s^t) + \sum_{\theta'} \mu(s^t, \theta_i, \theta'))^2} \\ &- \frac{(1-p) \lambda_{t+1} \pi(\theta_i)}{(1 + \eta(s^t) + \sum_{\theta'} \left\{ \mu(s^t, \theta_i, \theta') - \frac{\pi(\theta')}{\pi(\theta_i)} \mu(s^t, \theta', \theta_i) \right\})^2} \\ &+ \frac{(1-p) \lambda_{t+1} \pi(\theta_i)}{(1 + \eta(s^t) + \sum_{\theta'} \left\{ \mu(s^t, \theta_j, \theta') - \frac{\pi(\theta')}{\pi(\theta_j)} \mu(s^t, \theta', \theta_j) \right\})^2}. \end{aligned}$$

We evaluate $\frac{\partial E[u'(s^{t+1})]}{\partial \mu(s^t, \theta_i, \theta_j)}$ at $\bar{\mu} = 0$ to obtain

$$\left. \frac{\partial E[u'(s^{t+1})]}{\partial \mu(s^t, \theta_i, \theta_j)} \right|_{\bar{\mu}=0} = -p \frac{\lambda_{t+1} \pi(\theta_i)}{(1 + \eta(s^t))^2}.$$

By Lemma 3 multiplier $\eta(s^t)$ is bounded above by $B_\eta < \infty$. Hence

$$\left. \frac{\partial E[u'(s^{t+1})]}{\partial \mu(s^t, \theta_i, \theta_j)} \right|_{\bar{\mu}=0} \leq -p \frac{\lambda \pi(\theta_i)}{(1 + B_\eta)^2} < 0.$$

It follows that the intertemporal wedge is negative $\tau(s^t) < 0$ if for all pairs θ_i, θ_j we have $\mu(s^t, \theta_i, \theta_j) < \bar{\mu}$, where $\bar{\mu}$ is strictly positive and constant across nodes and valid for all horizons T .

For a given horizon T and constrained-efficient allocation, assume that at all nodes for dates $t < T$ we have $\tau(s^t) > 0$. We can then construct a type \hat{s}^T by selecting such pair and $a_t = 1$ for all dates, that is, for any node \hat{s}^t select $\hat{s}^{t+1} = \{\hat{s}^t, \theta, 1\}$ where θ is such that $\mu(s^t, \theta, \theta') \geq \bar{\mu}$ for some θ' . If such θ does not exist, then $\tau(s^t) \leq 0$ contradicting the premise.

By f.o.n.c. 15 we have that

$$\eta(\hat{s}^T) \geq T \bar{\mu}. \tag{20}$$

If there does not exist a finite T such that $\tau(s^t) \leq 0$ for some node, then (20) implies that there exists a finite T^* with $T^* \bar{\mu} > B_\eta$, contradicting Lemma 3.

□

D Proofs for Section 7

Proof: [Proof of Lemma 2] Let $\bar{\theta} = \max\{\theta \in \Theta\}$ and x a constrained-efficient allocation with z the associated plan of continuation welfare. We show that for all nodes s^t ,

$$h(x(s^t); \bar{\theta}) + \beta z(s^t) > \Upsilon. \quad (21)$$

As previously shown, allocations for the top skill agent do not depend on the realization of a_t , and we can write the system of incentive-compatibility constraints for such agent as

$$h(x(s^t); \bar{\theta}) + \beta z(s^t) \geq p\Upsilon + (1-p)(h(x(\hat{s}^t); \bar{\theta}) + \beta z(\hat{s}^t))$$

for any $\theta' \neq \bar{\theta}$, $\hat{s}^t = \{s^{t-1}, 0, \theta'\}$. By Proposition 6, $y(\hat{s}^t) > 0$ and thus

$$h(x(\hat{s}^t); \bar{\theta}) > h(x(\hat{s}^t); \theta').$$

The lower bound on welfare (1) implies

$$h(x(\hat{s}^t); \theta') + \beta z(\hat{s}^t) \geq \Upsilon,$$

and thus (21) follows. Since $\pi(\theta_{t+1} = \bar{\theta} | s^t) > 0$ it is immediate that $z(s^t) > \Upsilon$. □

Proof: [Proof of Theorem 2] Define

$$\tilde{\eta}_t \equiv E\{\eta(s^t)\} = \sum_{s^t \in S^t} \pi(s^t) \eta(s^t).$$

By the law of iterated expectations and the conditional expectation of η_{t+1} , we have that

$$\tilde{\eta}_{t+1} = \tilde{\eta}_t + pE\{\mu^*(s^t)\} \quad (22)$$

Since $\mu^*(s^t) \geq 0$ the sequence $\{\tilde{\eta}_t\}$ is increasing by (22). Because the sequence is bounded above, it necessarily has to converge. Thus $\mu^*(s^t)$ converges in mean to 0,

$$\lim_{t \rightarrow \infty} \sum_{s^t \in S^t} \pi(s^t) \mu^*(s^t) = 0.$$

Since all terms are positive,

$$\lim_{t \rightarrow \infty} \sum_{s^t \in S^t} \pi(s^t) \mu(s^t, \theta; \theta') = 0$$

for all $\theta, \theta' \in \Theta$. Convergence-in-mean implies convergence-in-probability, so for any $\varepsilon > 0$, the probability $\Pr(\mu(s^t, \theta; \theta') > \varepsilon)$ converges to 0 for all $\theta, \theta' \in \Theta$.

Next, we show that if $|c(s^t) - c(s^{t+1})| \geq \delta$ for some $\delta > 0$ then $\mu(s^t, \theta; \theta') \geq \varepsilon_\delta > 0$ for some pair $\{\theta, \theta'\}$. Since we have shown the Lagrangian multiplier on promised utility η_t to be bounded, $c(s^t)$ is bounded above and below, and $\kappa = \inf\{|u''(s^t)/(u'(s^t))^2| : s^t \in S^t, t \geq 0\}$ is finite. Using a Taylor expansion we have that

$$\begin{aligned} \left| \frac{\lambda}{u'(s^t)} - \frac{\lambda}{u'(s^{t+1})} \right| &\geq \kappa |c(s^t) - c(s^{t+1})| \\ &\geq \kappa \delta = \delta'. \end{aligned}$$

If $a_{t+1} = 0$ then (10) at s^{t+1} combined with (19) at s^t imply

$$\left| \sum_{\theta' \in \Theta} \left\{ \mu(s^t, \theta_{t+1}; \theta') - \frac{\pi(\theta')}{\pi(\theta_{t+1})} \mu(s^t, \theta'; \theta_{t+1}) \right\} \right| \geq \delta$$

and thus

$$\sum_{\theta' \in \Theta} \left\{ |\mu(s^t, \theta_{t+1}; \theta')| + \frac{1}{\pi(\theta_{t+1})} |\mu(s^t, \theta'; \theta_{t+1})| \right\} \geq \delta'. \quad (23)$$

Set

$$\varepsilon_\delta = \frac{\delta' \min\{\pi(\theta) : \theta \in \Theta\}}{2N_\Theta}$$

where N_Θ is the cardinality of set Θ . By construction, if $\mu(s^t, \theta; \theta') < \varepsilon_\delta$ for all $\{\theta, \theta'\}$ then we contradict (23).

If $a_{t+1} = 1$ then (13) at s^{t+1} combined with (19) at s^t imply

$$\left| \sum_{\theta' \in \Theta} \mu(s^t, \theta_{t+1}; \theta') \right| \geq \delta',$$

and thus $\mu(s^t, \theta; \theta') > \varepsilon_\delta$ for some $\{\theta, \theta'\}$.

Thus for any $\delta > 0$

$$\lim_{t \rightarrow \infty} \Pr(|c(s^t) - c(s^{t+1})| \geq \delta) \leq \lim_{t \rightarrow \infty} \Pr(\mu(s^t, \theta; \theta') \geq \varepsilon_\delta) = 0.$$

Next, we show that there are no labor distortions in the limit following very similar steps.

If $|u'(s^t) - h^y(s^t; \theta_t)| \geq \delta$ then by manipulation of the f.o.n.c. (11) we have that

$$|u'(s^t) - h^y(s^t; \theta_t)| = \frac{\sum_{\theta' \in \Theta} \frac{\pi(\theta')}{\pi(\theta_t)} \mu(s^{t-1}, \theta'; \theta_t) h^y(s^t; \theta')}{1 + \eta(s^t)}.$$

For all binding IC constraints, $h^y(s^t; \theta') \leq h^y(s^t; \theta_t)$, and with the first-order conditions for the top skill agent and 2, $h^y(s^t; \theta_t) \leq u'(s^t)$. Furthermore, Lemma 3 implies $u'(s^t) \leq \kappa < \infty$ and $-1 < B_0 \leq \eta(s^t)$, so

$$|u'(s^t) - h^y(s^t; \theta_t)| \leq \frac{1}{(1 + B_0)\pi(\theta_t)} \sum_{\theta' \in \Theta} \mu(s^{t-1}, \theta'; \theta_t) \kappa.$$

By setting $\varepsilon_\delta = \delta(1 + B_0) \min\{\pi(\theta) : \theta \in \Theta\} / (\kappa N_\Theta)$ we can show that if $|u'(s^t) - h^y(s^t; \theta_t)| \geq \delta$ then $\mu(s^t, \theta; \theta') > \varepsilon_\delta$ for some $\{\theta, \theta'\}$ and $\varepsilon_\delta > 0$. Hence

$$\lim_{t \rightarrow \infty} \Pr(|u'(s^t) - h^y(s^t; \theta_t)| \geq \delta) = 0$$

for any $\delta > 0$.

E Proofs for Section 8

We show that $\Gamma(s^{t-1}, \theta^t)$ is indeed a lower bound on the marginal benefit of increasing the optimal monitoring probability. To simplify notation we set $s_1^t = \{s^{t-1}, \theta_t, 1\}$ and $s_0^t = \{s^{t-1}, \theta_t, 0\}$. For any optimal allocation given $\{p_t : t \geq 0\}$ with $p_t > 0$, the derivative of the Lagrangian with respect to the monitoring probability of agent $\{s^{t-1}, \theta_t\}$ is

$$\begin{aligned} & (h(x(s_1^t); \theta_t) + z(s_1^t) - h(x(s_0^t); \theta_t) - z(s_0^t)) \left(1 + \eta(s^{t-1}) + \sum_{\theta'} \mu(s^{t-1}, \theta_t; \theta') \right) \\ & - \lambda_t (c(s_1^t) - c(s_0^t) + y(s_0^t) - y(s_1^t)) \\ & + \sum_{\theta'} \mu(s^{t-1}, \theta'; \theta_t) \frac{\pi(\theta')}{\pi(\theta_t)} (h(x(s_0^t); \theta_t) + z(s_0^t) - (1 - \beta^{T+1-t})\Upsilon). \end{aligned}$$

By concavity,

$$h(x(s_0^t); \theta_t) \leq h(x(s_1^t); \theta_t) + D_x h(x(s_1^t); \theta_t) (x(s_1^t) - x(s_0^t)).$$

The necessary first-order conditions associated with $x(s_0^t)$ imply

$$\begin{aligned} u'(s^t) \left(1 + \eta(s^{t-1}) + \sum_{\theta' \in \Theta} \mu(s^{t-1}, \theta_t; \theta') \right) &= \lambda_t \\ h^y(s^t; \theta_t) \left(1 + \eta(s^{t-1}) + \sum_{\theta' \in \Theta} \mu(s^{t-1}, \theta_t; \theta') \right) &= \lambda_t, \end{aligned}$$

we can sign the terms

$$\begin{aligned} (h(x(s_1^t); \theta_t) + z(s_1^t) - h(x(s_0^t); \theta_t) - z(s_0^t)) \left(1 + \eta(s^{t-1}) + \sum_{\theta'} \mu(s^{t-1}, \theta_t; \theta') \right) \\ - \lambda_t (c(s_1^t) - c(s_0^t) + y(s_0^t) - y(s_1^t)) \end{aligned}$$

as positive since by Proposition 2 $c(s_1^t) \geq c(s_0^t)$, $y(s_1^t) \leq y(s_0^t)$, and $z(s_1^t) \geq z(s_0^t)$.

□