

Persistent Private Information

Noah Williams*

Department of Economics, University of Wisconsin - Madison

E-mail: nmwilliams@wisc.edu

Revised December 22, 2010

This paper studies the design of optimal contracts in dynamic environments where agents have private information that is persistent. In particular, I focus on a continuous time version of a benchmark insurance problem where a risk averse agent would like to borrow from a risk neutral lender to stabilize his utility. The agent privately observes a persistent state variable, typically either income or a taste shock, and he makes reports to the principal. I give verifiable sufficient conditions showing that incentive compatible contracts can be written recursively, conditioning on the report and two additional state variables: the agent's promised utility and promised marginal utility of the private state. I then study two examples where the optimal contracts can be solved in closed form, showing how persistence alters the nature of the contract. Unlike the previous discrete time models with i.i.d. private information, the agent's consumption under the contract may grow over time. Furthermore, in my setting the efficiency losses due to private information increase with the persistence of the private information, and the distortions vanish as I approximate an i.i.d. environment.

1. INTRODUCTION

Private information is an important component of many economic environments, and in many cases that information is persistent. Incomes of individuals from one period to the next are highly correlated, and can be difficult for an outsider to verify. Similarly, a worker's skill at performing a task, a manager's ability in leading a firm, and a firm's efficiency in producing goods are all likely to be highly persistent, and contain at least some element of private information. While models of private information have become widespread in many areas of economics, with very few exceptions previous research has focused on situations where the private information has no persistence. But estimates suggest that idiosyncratic income streams and individual skills, which likely have large private information components, are highly persistent.¹ In this paper I analyze a benchmark dynamic contracting model with persistent private information. Working in continuous time, I develop general sufficient conditions which allow for a recursive representation of incentive compatible contracts. This allows me to develop a characterization of such contracts, which has been lacking, and makes a range of new and empirically relevant applications amenable for analysis. I use my characterization

* I thank Dilip Abreu, Katsuhiko Aiba, Narayana Kocherlakota, Rui Li, Chris Sleet, and Ted Temzelides for helpful comments. I also especially thank David Levine (the editor) and two anonymous referees for comments that greatly improved the paper. Financial support from the National Science Foundation is gratefully acknowledged.

¹See for example Storesletten, Telmer, and Yaron (2004) and Meghir and Pistaferri (2004).

to explicitly solve for optimal contracts in some key examples, showing how increasing the persistence of the private information adversely affects the degree of risk-sharing.

In particular, I focus on a continuous time dynamic contracting relationship between a risk neutral principal and a risk-averse agent, where the agent privately observes a payoff-relevant persistent state variable. I start with a relatively general setting, which maps into applications similar to classic papers in the literature. One version of the model is a classic insurance problem along the lines of the Green (1987) and Thomas and Worrall (1990), where a risk averse agent would like to borrow from a risk neutral lender to stabilize his income stream. The income stream is private information to the borrower and is persistent. Another version of the general model is an insurance problem similar to Atkeson and Lucas (1992), where the agent is hit with privately observed preference shocks which alter his marginal utility of consumption. My general characterization, which applies to both applications, adapts the continuous time methods developed in Williams (2009) for hidden action problems to handle models with private information.

One of the main contributions of the paper is to provide a set of necessary and sufficient conditions which ensure that a contract is incentive compatible. To do so, I apply a stochastic maximum principle due to Bismut (1973)-(1978) to derive the agent's optimality conditions facing a given contract, which in turn provide conditions that an incentive compatible contract must satisfy. I show that such contracts must be history dependent, but this dependence is encapsulated in two endogenous state variables: the agent's promised utility and promised marginal utility under the contract.

The use of promised utility as a state variable in contracting is now well known, but when there is persistent private information incentive compatibility requires more conditioning information.² The difficulty with persistent information is that deviations from truth-telling alter agents' private evaluations of future outcomes, or put differently, agents' continuation types are no longer public information. I show that in my setting incentive compatibility can be ensured by conditioning on a discounted expected marginal utility state variable, which captures the expected future shadow value of the private information. The discount rate in this state variable includes a measure of the persistence. More persistent information is longer-lived and so has a greater impact on future evaluations, while as the persistence vanishes this additional state variable becomes unnecessary.

The other main contribution of the paper is to explicitly solve for optimal contracts in some special cases of interest. In particular, I study a hidden endowment model where the agent has exponential utility, and a model with permanent private taste shocks where the agent has power utility over consumption. The hidden endowment model highlights the relationship between the persistence of the private information

²The use of promised utility as a state variable follows the work of Abreu, Pearce, and Stacchetti (1986)-(1990) and Spear and Srivastava (1987). It used by Green (1987) and Thomas and Worrall (1990) in settings like mine, and Sannikov (2008) in a continuous time model. Similar marginal utility states have been used by Kydland and Prescott (1980) and in related contexts by Werning (2001), Abraham and Pavoni (2008), and Kapicka (2006), as well as my earlier paper Williams (2009).

and the magnitude of the distortions this information causes. The largest distortions occur when the private information is permanent. In this case the optimal contracts entail no risk sharing, consisting only of a deterministic transfer which changes the time profile of the agent's consumption. However as information becomes less persistent, risk sharing increases and the agent's consumption becomes more stable. In particular in the i.i.d. limit I obtain efficiency and complete stabilization of consumption. This result differs from the discrete time i.i.d. models in the literature, which clearly have nonzero distortions. But I show that as the period length shrinks in a discrete time model, the distortions shrink as well. Thus the existing results are consistent with my findings. The private taste shock model has much of the same structure, but it becomes intractable analytically outside the case of permanent shocks.

My results differ in some fundamental ways from those in the literature. In particular, Rogerson (1985) and Golosov, Kocherlakota, and Tsyvinski (2003) have shown in discrete time that an “inverse Euler equation” governs consumption dynamics in private information models. Technically, this implies that the inverse of the agent's marginal utility of consumption is a martingale. Closely related are the “immiseration” results of Thomas and Worrall (1990) and Atkeson and Lucas (1992) which imply that the agent's promised utility tends to minus infinity under the optimal contract. In my examples these results fail – the agent's promised utility follows a martingale and consumption has a positive drift under the optimal contract.

As I discuss in more detail below, these differences rely at least partly on differences in the environments. In the discrete analogue of my model, when deciding what to report in the current period, the agent trades off current consumption and future promised utility. In my continuous time formulation, the agent's private state follows a process with continuous paths, and the principal knows this. Thus in the current period the agent only influences the *future increments* of the reported state.³ Thus current consumption is independent of the current report, and all that matters for the reporting choice is how future transfers are affected. As the reporting problem and hence the incentive constraints become fully forward-looking, optimal allocations no longer involve the balance of current and future distortions that the inverse Euler equation embodies. By contrast, in a moral hazard setting with hidden effort such as Williams (2009), an inverse Euler equation does hold. When effort is costly, deviations from the behavior specified by the contract have instantaneous utility effects which are absent in the hidden information model. Thus the inverse Euler equation is sensitive to the source of the information frictions.

Besides my own previous work, the closest paper in terms of technique is DeMarzo and Sannikov (2006), who also study a continuous time hidden information problem.⁴ My conversion of the reporting problem to a hidden action problem follows their approach. However they focus on the case where agents are risk neutral and there is no persistence in information. My paper is more general along these dimensions.

³As discussed below, this is a requirement of absolute continuity. If the agent were to report a process which jumped discretely at any instant, the principal would be able to detect that his report was a lie.

⁴Biais, Mariotti, Plantin, and Rochet (2007) study a closely related discrete time model and its continuous time limit.

One particular difference is that given the risk neutrality they can define cash flows, their source of private information, as *increments* of a Brownian motion with constant drift. Thus private information in their case is i.i.d. However in my environment it is more natural, and more consistent with the literature, to define the private state as the *level* of a diffusion process. As all diffusions are persistent, private information is always persistent in my environment.⁵

In addition, there are a few recent papers which study models with persistent private information. In a discrete time setting, Battaglini (2005) and Tchistyi (2006) characterize contracts between risk neutral agents when the unobserved types switch according to a Markov chain. Their results rely heavily on risk neutrality, and thus do not extend to the classic insurance issues which I address here. In their analysis of disability insurance, Golosov and Tsyvinski (2006) study a particularly simple form of persistent information, where the type enters an absorbing state. Again, their results heavily exploit this particular structure. More closely related to the present paper, Zhang (2009) studies a continuous time hidden information model with two types that switch according to a continuous time Markov chain. His model is a continuous time version of Fernandes and Phelan (2000), who consider a relatively general approach for dealing with history dependence in dynamic contracting models. Both Zhang (2009) and Fernandes and Phelan (2000) allow for persistence in the hidden information, at the cost of a state space which grows with the number of types. In contrast, my approach deals with a continuum of types, and requires two endogenous state variables. This is achieved by exploiting the continuity of the problem to introduce a state variable which captures the shadow value (in marginal utility terms) of the hidden information. Thus my approach is analogous to a first-order approach to contracting.⁶

Kapicka (2006) considers a related first-order approach in a discrete time model. As I discuss in more detail in section 8, we both use similar ideas to arrive at recursive representations of contracts, but our state variables differ. Kapicka (2006) takes the agent's promised marginal utility of *consumption* as his additional state variable, while in my approach the marginal utility of *the hidden state* (discounted by the persistence of the private information) is the relevant state variable. Although these state variables agree in some cases, I present an example below where they differ. In addition, my proof of the applicability of the first-order approach is more complete and direct, and I show how to verify my sufficient conditions in applied settings. Our implications for contracts differ as well. Kapicka (2006) shows how optimal contracts may significantly distort the agent's consumption-savings decision. However I shown that when shocks are permanent, this channel need not be distorted. In my example, the only way the principal can ensure truthful revelation is to make the transfer independent of the report. Such a contract does not distort the agent's intertemporal decisions, but it clearly does not insure the agent either.

The persistence of private information gives the contracting problem an adverse selection component as well. In a continuous time setting, Sannikov (2007) and Cvitanic

⁵As I discuss in section 2.2 below, the i.i.d. case can be approximated by considering a limit of a mean reverting process as the speed of mean reversion goes to infinity.

⁶See Williams (2009) for more discussion of the first-order approach to contracting in continuous time.

and Zhang (2007) study contracting problems with adverse selection. Both focus on cases where there are two possible fixed types of agents, and apart from the agent's type there is no persistence in private information. While I do not emphasize the screening component of contracts, one could apply my methods to screening models where agents' types vary over time.

The rest of the paper is organized as follows. The next section describes the model and discusses a key change of variables which is crucial in analyzing it. Section 3 analyzes the agent's reporting problem when facing a given contract, deriving necessary optimality conditions. The two endogenous state variables arise naturally here. Section 4 characterizes truthful reporting contracts, presenting sufficient conditions which ensure incentive compatibility. Section 5 presents a general approach to solve for optimal contracts. Then in section 6 I consider a hidden endowment example where the optimal contract can be determined in closed form. I show how persistence matters for the contract, and contrast our results with some of the others in the literature. Section 8 studies a related model with private taste shocks, which can also be solved explicitly with permanent shocks. Finally section 9 offers some brief concluding remarks.

2. THE MODEL

2.1. Overview

The model consists of a risk averse agent who privately observes a state variable which affects his utility, and he receives transfers from a risk-neutral principal in order to smooth his consumption. The privately observed state variable can be an endowment process as in Thomas and Worrall (1990), or a taste shock as in Atkeson and Lucas (1992). If the principal could observe the private state, then he could absorb the risk and fully stabilize the agent's utility. However with the state being the private information of the agent, the principal must rely on the agent's reports. Under the full information contract, the agent would have an incentive to lie, for example reporting that his income is lower than it really is. Thus the key problem is to design a contract which provides the agent the incentive to truthfully report his information. Relative to the literature, the my key innovation is to allow for persistence in the agent's privately observed state.

2.2. Basic Layout

I start by considering a finite horizon $[0, T]$, and later let $T \rightarrow \infty$. I use a plain letter to denote a whole path of a variable, thus for instance $b = (b_t)_{t=0}^T$. I suppose that the privately observed variable b is given by a Markov diffusion process defined on a probability space with a Brownian motion W , which evolves as:

$$db_t = \mu(b_t)dt + \sigma dW_t.$$

I assume that $\sigma > 0$ is constant, and that the drift is affine:

$$\mu(b) = \mu_0 - \lambda b,$$

with $\lambda \geq 0$. All the results in the paper could be easily extended to the more general case where the drift $\mu : \mathbb{R} \rightarrow \mathbb{R}$ is twice continuously differentiable, (weakly) decreasing, and (weakly) concave. While this extension may be useful in applications, it adds little conceptually to the issues at hand and only leads to more cumbersome notation.⁷

With $\lambda = 0$, b is a Brownian motion with drift, so its increments are i.i.d., and shocks have permanent effects.⁸ On the other hand with $\lambda > 0$, b is an Ornstein-Uhlenbeck process (see Karatzas and Shreve (1991)). This is a continuous-time version of a stationary Gaussssian autoregressive process, with the following properties:

$$\begin{aligned} E(b_t|b_0) &= \frac{\mu_0}{\lambda} + \left(b_0 - \frac{\mu_0}{\lambda}\right) e^{-\lambda t} \\ Cov(b_t, b_s|b_0) &= \frac{\sigma^2}{2\lambda} (e^{-\lambda|s-t|} - e^{-\lambda(s+t)}) \end{aligned}$$

Thus μ_0/λ gives the mean of the stationary distribution, while λ governs the rate of mean reversion and hence the persistence of the process. As mentioned above, the private state cannot be i.i.d., but we can approximate an i.i.d. process by setting $\sigma = \bar{\sigma}\sqrt{\lambda}$ and $\mu_0 = \bar{\mu}\lambda$ for some $\bar{\sigma}, \bar{\mu} > 0$ and letting $\lambda \rightarrow \infty$. The limit is an i.i.d. normal process with mean $\bar{\mu}$ and variance $\bar{\sigma}^2/2$.

The key issue in this model, of course, is that the agent alone observes b_t , while the principal only observes the agent's report of it, which I denote y_t . To simplify matters, I assume that the agent cannot over-report the true state, so $y_t \leq b_t$. In most settings the relevant incentive constraint guards against under-reporting, so such a constraint is not overly restrictive.⁹ When the private state variable is the agent's endowment, this assumption can be motivated by requiring the agent to deposit a portion of his endowment in an account which the principal observes. The good is non-storable and the agent cannot manufacture additional endowments, so the deposit (or report) is a verifiable statement of at most the entire endowment.

The fact that the agent's private state has continuous sample paths is key to our analysis. This continuity allows us to use local information to summarize histories, it gives structure to the types of reports the agent can make, and together with the concavity of utility it allows us to map local optimality conditions into global ones. Other types of processes which allow for jumps may be useful in some applications (such as Zhang (2009)). But allowing for jumps is more complex, as the contract would require more global information. Models with continuous information flows provide a useful benchmark, and are natural in many applications. When the private state is the hidden endowment, similar continuous specifications have been widely used in asset pricing following Breeden (1979). Although lumpiness may be important for some types of income, business or financial income may involve frequent flows which

⁷Note that I allow b to affect utility in a general way, so b may be a transformation of an underlying state variable.

⁸As noted above, if we were to follow DeMarzo and Sannikov (2006) b_t would be the *cumulative* state process, and hence i.i.d. increments would correspond to an i.i.d. state variable. However apart from the risk neutral case they consider, defining utility over the increments of a diffusion process is problematic.

⁹In our setting this restriction facilitates analyzing the incentive constraints. Removing the restriction at the outset is difficult due to the linearity of agent's reporting problem.

may be difficult for an intermediary to observe. Similarly, when the private state is a preference shock, we focus not on lumpiness due to say discrete changes in household composition, but rather on more frequent, smaller fluctuations in tastes or skills.

The agent's information can be summarized by the paths of W , which induce Wiener measure on the space $C[0, T]$ of continuous functions of time. The agent's reporting strategy is thus a predictable mapping $y : C[0, T] \rightarrow C[0, T]$. Denote this mapping $y(\omega)$ and its time t component $y_t(\omega)$ which is measurable with respect to \mathcal{B}_t , the Borel σ -algebra of $C[0, T]$ generated by $\{W_s : s \leq t\}$. The principal observes y only, and thus his information at date t can be represented via \mathcal{Y}_t , the Borel σ -algebra of $C[0, T]$ generated by $\{y_s : s \leq t\}$. I assume that the private state is initialized at a publicly known value b_0 , and that the principal knows the process which the state follows (i.e. he knows μ and σ), but he does not observe the realizations of it.

As the principal obtains continuous reports from the agent, the agent is not free to choose an arbitrary reporting strategy. In particular, based on the reports the principal can construct a process W_t^y which evolves as:

$$dW_t^y = \frac{dy_t - \mu(y_t)dt}{\sigma}. \quad (1)$$

Under a truthful reporting strategy $y_t = b_t$, clearly we have $W_t^y = W_t$. Thus the agent is restricted to reporting strategies which ensure that W_t^y is a Brownian motion with respect to the principal's information set. If he were to choose a reporting strategy which did not make W_t^y a Brownian motion, say for instance he reported a constant endowment for a strictly positive length of time (which has probability zero), then the principal would detect this lie and would be able to punish him.

Formally, the agent's report y must be absolutely continuous with respect to his true state b . Hence via the Girsanov theorem and related results (see Chapters 6 and 7 of Liptser and Shiryaev (2000), for example), this means that the agent's reporting process is equal to the true state process plus a drift:

$$dy_t = db_t + \Delta_t dt,$$

where Δ_t is a process adapted to the agent's information set. Since the agent can report (or deposit) at most his entire state, we must have $\Delta_t \leq 0$. Integrating this evolution, and using $y_0 = b_0$, we see that the report of the private state is equal to the truth plus the cumulative lies:

$$y_t = b_t + \int_0^t \Delta_s ds \equiv b_t + m_t,$$

where we define $m_t \leq 0$ as the "stock of lies". With this notation, we can then write the evolution of the agent's reporting and lying processes as:

$$dy_t = [\mu(y_t - m_t) + \Delta_t]dt + \sigma dW_t, \quad (2)$$

$$dm_t = \Delta_t dt, \quad (3)$$

with $y_0 = b_0$ and $m_0 = 0$. The principal observes y_t but cannot separate Δ_t from W_t , and thus cannot tell whether a low report was due to a lie or a poor shock realization. Moreover, the stock of lies m_t is a hidden state which is unobservable to the principal, but influences the evolution of the observable report state.

In our environment, a contract is a specification of payments from the principal to the agent conditional on the agent's reports. The principal makes payments to the agent throughout the period $s : C[0, T] \rightarrow C[0, T]$ which are adapted to his information $\{\mathcal{Y}_t\}$, as well as a terminal payment $S_T : C[0, T] \rightarrow \mathbb{R}$ which is \mathcal{Y}_T -measurable. This is a very general representation, allowing almost arbitrary history dependence in the contract. We will study the choice of the optimal contract s by the principal, finding a convenient representation of the report history.

2.3. A Change of Variables

As the contract is history dependent, the agent's utility at any given time will in general depend on the whole past history of reports. This makes direct analysis difficult, as standard dynamic programming methods are not applicable. This lack of a recursive structure is well-known in contracting models, and the seminal works of Spear and Srivastava (1987), Green (1987), and Abreu, Pearce, and Stacchetti (1990) show how to make the problem recursive by enlarging the state space. We use the agent's optimality conditions to derive such additional necessary state variables.

To do so, as in Williams (2009), I follow Bismut (1978) and change state variables to focus on the distribution of observed reports. Rather than directly analyzing the agent's choice of what to report at each date, it is easier to view the agent as choosing a probability distribution over the entire reporting process that the principal observes. Alternative reporting strategies thus change the distribution of observed outcomes.

A key simplification of our continuous time setting is the natural mapping between probability measures and Brownian motions. The following calculations rely on standard stochastic process results, as in Liptser and Shiryaev (2000). To begin, fix a probability measure P_0 on $C[0, T]$, and let $\{W_t^0\}$ be a Wiener process under this measure, with $\{\mathcal{F}_t\}$ the completion of the filtration generated by it. Under this measure, the agent reports that his private state follows a martingale:

$$dy_t = \sigma dW_t^0, \quad (4)$$

with y_0 given. Different reporting choices alter the distribution over observed outcomes in $C[0, T]$. In particular, for any feasible lying choice $\Delta = \{\Delta_t\}$, define the family of \mathcal{F}_t -predictable processes:

$$\Gamma_t(\Delta) = \exp \left(\int_0^t \frac{\mu(y_s - m_s) + \Delta_s}{\sigma} dW_s^0 - \frac{1}{2} \int_0^t \left(\frac{\mu(y_s - m_s) + \Delta_s}{\sigma} \right)^2 ds \right). \quad (5)$$

Γ_t is an \mathcal{F}_t -martingale with $E_0[\Gamma_T(\Delta)] = 1$, where E_0 represents expectation with respect to P_0 . Thus by the Girsanov theorem, we define a new measure P_Δ via:

$$\frac{dP_\Delta}{dP_0} = \Gamma_T(\Delta),$$

and the process W_t^Δ defined by:

$$W_t^\Delta = W_t^0 - \int_0^t \frac{\mu(y_s - m_s) + \Delta_s}{\sigma} ds \quad (6)$$

is a Brownian motion under P_Δ . For the truthful reporting strategy $\Delta_t^* \equiv 0$ we denote $P_{\Delta^*} = P^*$, with corresponding Brownian motion W_t^* . Hence each effort choice Δ results in a different Brownian motion, similar to what we constructed above in (1). However the Brownian motion W_t^y is based on the principal's observations of y , assuming truthful reporting. Here the Brownian motion W_t^Δ is based on the agent's (full) information, as a means of depicting the distribution P_Δ over observed reports.

The density process Γ_t captures the effect of the agent's lies on observed outcomes, and we take it rather than y_t as the key state variable. Although the contract depends on the entire history of reports, this density allows us to average over that history in a simple manner. Using (5), the density evolves as:

$$d\Gamma_t = \frac{\Gamma_t}{\sigma} [\mu(y_t - m_t) + \Delta_t] dW_t^0, \quad (7)$$

with $\Gamma_0 = 1$. The stock of lies m_t is the other relevant state variable for the agent, and it's convenient to analyze its evolution under the transformed probability measure. Thus we take $z_t = \Gamma_t m_t$, which captures covariation between reports and the stock of lies, as the relevant hidden state variable. Simple calculations show that it follows:

$$dz_t = \Gamma_t \Delta_t dt + \frac{z_t}{\sigma} [\mu(y_t - m_t) + \Delta_t] dW_t^0 \quad (8)$$

with $z_0 = 0$. By changing variables from (y, m) to (Γ, z) the agent's reporting problem is greatly simplified. Instead of depending on the entire past history of the state y_t , the problem only depends on contemporaneous values of the states Γ_t, z_t . The history y is treated as an element of the probability space. This leads to substantial simplifications, as I show below.

3. THE AGENT'S REPORTING PROBLEM

In this section I derive optimality conditions for an agent facing a given contract. The agent's preferences take a standard time additive form, with discount rate ρ , smooth concave flow utility $u(s, b)$, and terminal utility $U(s, b)$ defined over the payment and the private state. Below I focus on the special cases where b is the agent's endowment, so $u(s, b) = v(s + b)$, and where b is the log of a taste shock, so $u(s, b) = v(s) \exp(-b)$ where $v(s) \leq 0$. I suppose that the agent has the option at date zero of rejecting a contract and remaining in autarky. This gives a participation constraint that the utility the agent achieves under the contract must be greater than his utility under autarky, denoted $V^a(b_0)$. However after date zero, both the parties are committed to the contract and cannot leave it.

The agent's preferences for an arbitrary reporting strategy $\{\Delta_t\}$ can be written:

$$\begin{aligned} V(y; s) &= E_\Delta \left[\int_0^T e^{-\rho t} u(s_t(y), y_t - m_t) dt + e^{-\rho T} U(S_T(y), y_T - m_T) \right] \\ &= E_0 \left[\int_0^T e^{-\rho t} \Gamma_t u(s_t(y), y_t - m_t) dt + e^{-\rho T} \Gamma_T U(S_T(y), y_T - m_T) \right]. \end{aligned}$$

Here the first line uses the contract, substitutes $y - m$ for b , and takes the expectation with respect to the measure P_Δ over reporting outcomes. The second line uses the density process defined above. The agent takes the contract s and the evolution (7)-(8) as given and solves:

$$\sup_{\{\Delta_t \leq 0\}} V(y; s)$$

Under the change of variables, the agent's reporting problem is a control problem with random coefficients. As in Williams (2009), I apply a stochastic maximum principle from Bismut (1973)-(1978) to derive the agent's necessary optimality conditions. I first derive all conditions using the new variables (Γ, z) , then express everything in terms of the original states (y, m) . Analogous to the deterministic Pontryagin maximum principle, I define a (current-value) Hamiltonian function which the optimal control will maximize. As in the deterministic theory, associated with the state variables are co-state variables which have specified terminal conditions. However to respect the stochastic information flow and satisfy the terminal conditions, the co-states are now *pairs* of processes which satisfy backward stochastic differential equations. Thus I introduce (q, γ) as the co-states associated with the state Γ , and (p, Q) associated with the state z . In each pair, the first co-state multiplies the drift of the state while the second multiplies the diffusion term. Thus the Hamiltonian is given by:

$$\mathcal{H}(\Gamma, z) = \Gamma u(s(y), y - z/\Gamma) + (\Gamma\gamma + Qz) (\mu(y - z/\Gamma) + \Delta) + p\Gamma\Delta$$

The Hamiltonian captures the instantaneous utility flows to the agent. Here we see that a lie Δ affects the future increments of the likelihood of outcomes, Γ , as well as the hidden state variable z . Moreover these state variables matter for the agent's future utility evaluations, both through the direct dependence of the utility function on the hidden state and through the change in the likelihood of alternative reporting paths.

The agent's optimal choice of his report perturbation $\Delta \leq 0$ is given by maximizing the Hamiltonian, and the evolution of the co-state variables is given by differentiating the Hamiltonian. To simplify matters, I invoke the revelation principle and focus on contracts which induce truthful revelation. Thus we have $\Delta_t \equiv 0$, $m_t \equiv 0$, $y_t = b_t$. As the Hamiltonian \mathcal{H} is linear Δ and $\Gamma \geq 0$, in order to have a truthful current report $\Delta = 0$ be optimal we thus require:

$$\gamma + Qm + p \geq 0. \tag{9}$$

Moreover, given truthful reporting in the past (so $m = 0$) it must be optimal to report truthfully in the present, and thus we can strengthen this to:

$$\gamma + p \geq 0. \quad (10)$$

Under truthful revelation, the co-state variables evolve as follows:

$$\begin{aligned} dq_t &= \left[\rho q_t - \frac{\partial \mathcal{H}(\Gamma, z)}{\partial \Gamma} \right] dt + \gamma_t \sigma dW_t^0 \\ &= [\rho q_t - u(s_t, y_t)] dt + \gamma_t \sigma dW_t^*, \end{aligned} \quad (11)$$

$$q_T = U(s_T, y_T).$$

$$\begin{aligned} dp_t &= \left[\rho p_t - \frac{\partial \mathcal{H}(\Gamma, z)}{\partial z} \right] dt + Q_t \sigma dW_t^0 \\ &= [\rho p_t - \lambda \gamma_t + u_b(s_t, y_t)] dt + Q_t \sigma dW_t^*, \end{aligned} \quad (12)$$

$$p_T = -U_b(s_T, y_T).$$

Here (11) and (12) carry out the differentiation, evaluate the result under truthful revelation, and change the Brownian motion as in (6). Details of the derivations are in Appendix A.1. Notice that p_t and q_t solve backward stochastic differential equations, as they have specified terminal conditions but unknown initial conditions.

Below I show that these co-state processes encode the necessary history dependence that truthful revelation contracts require. In effect, the principal is able to tune the coefficients γ_t and Q_t in these state variables. That is, incentive compatible contracts can be represented via specifications of $s_t(y) = s(t, y_t, q_t, p_t)$, $\gamma_t = \gamma(t, y_t, q_t, p_t)$, and $Q_t = Q(t, y_t, q_t, p_t)$ for some functions s, γ, Q . Notice as well that the co-state evolution equations do not depend on Γ_t or z_t and thus the entire system consists of (2) with $\Delta_t = m_t = 0$, (11), and (12).

To help interpret the co-state equations, consider first the co-state (11). As can be easily verified, we can write its solution as:

$$q_t = E_* \left[\int_t^T e^{-\rho(\tau-t)} u(s_\tau, y_\tau) d\tau + e^{-\rho(T-t)} U(s_T, y_T) \middle| \mathcal{F}_t \right], \quad (13)$$

where E_* is the expectation with respect to P^* . Thus $q_0 = V(y)$, and the agent's optimal utility process becomes a state variable for the contracting problem. Using utility as a state variable is a well-known idea in the literature following Abreu, Pearce, and Stacchetti (1986) and Spear and Srivastava (1987), and has been widely used in contexts like ours following Green (1987) and Thomas and Worrall (1990).

In environments without persistent private information, the promised utility encapsulates the necessary history dependence. However with persistent private information, additional information is necessary. The agent's decisions at any date now may depend on both his previous history of reports as well as the true history of his private state. Put differently, with persistent information a lie in the current period affects both the principal's expectations of future realizations of the hidden state and

the agent's expectations of his own future payoffs. However, as I show below, when the agent's utility function is concave in the hidden state, local information suffices to capture the additional dependence. The variable p_t , which we call the promised marginal utility state, gives the marginal value in utility terms of the hidden state variable m_t evaluated at $m_t = 0$. Thus it captures the marginal cost of not lying.

In particular, suppose that (10) binds almost everywhere.¹⁰ Then (12) becomes:

$$dp_t = [(\rho + \lambda)p_t + u_b(s_t, y_t)] dt + Q_t \sigma dW_t^y.$$

It is easy to verify that the solution of this equation is:

$$p_t = -E_* \left[\int_t^T e^{-(\rho+\lambda)(\tau-t)} u_b(s_\tau, y_\tau) d\tau + e^{-(\rho+\lambda)(T-t)} U_b(s_T, y_T) \middle| \mathcal{F}_t \right]. \quad (14)$$

Thus p_t is the negative of the agent's optimal marginal utility (of the private state b_t) under the contract. As noted in the introduction, similar state variables have been used in related contexts by Werning (2001), Abraham and Pavoni (2008), Kapicka (2006), and Williams (2009).¹¹ Here the persistence of the endowment effectively acts as an additional discount, since larger λ means faster mean reversion which shortens the effective life of the private information. In the limit as $\lambda \rightarrow \infty$ we approximate an i.i.d. process, and p_t converges to an i.i.d. random variable. Thus in the i.i.d. case, a contract only need condition on the utility process, while persistent private information requires the marginal utility process as well.

Notice that γ_t , the loading of the utility process on the Brownian motion, is a key means of inducing truthful revelation. When b_t is the agent's endowment, a risk-neutral principal who directly observes b_t will fully stabilize the agent's consumption, and thus set $\gamma_t = 0$. (This should be clear intuitively, but will be shown explicitly below.) But when the agent has private information, the principal induces truthful revelation by making the agent's utility vary with his report. In particular (10) implies $\gamma_t \geq -p_t \geq 0$ so that promised utility q_t increases with larger realizations of the private state. With persistent private information, the principal also chooses Q_t , the loading of the promised marginal utility state on the endowment shocks. Thus when considering lying, the agent must balance the higher utility promise with the effects on his marginal utility promise.

4. TRUTHFUL REVELATION CONTRACTS

I now characterize the class of contracts which are individually rational and can induce the agent to truthfully report his private information. Instead of considering a contract as a general history-dependent specification of $s(y)$ as above, I now characterize a contract as processes for $\{s_t, \gamma_t, Q_t\}$ conditional on the states $\{y_t, q_t, p_t\}$. In

¹⁰Here we make this assumption only to help interpret the variable p_t , but it is not necessary for our results. We discuss this condition in more detail in Section 5.3 below.

¹¹Most of these references use either current or discounted future marginal utility of consumption as the additional state variable. The marginal utility of the private state is not necessarily the same as the marginal utility of consumption, as we show below.

this section I provide conditions on these processes that a truthful revelation contract must satisfy.

In standard dynamic contracting problems, like Thomas and Worrall (1990), contracts must satisfy participation, promise-keeping, and incentive constraints. Similar conditions must hold in our setting. I have already discussed the participation constraint, which must ensure that the agent's utility under the contract is greater than under autarky. This simply puts a lower bound on the initial utility promise:

$$q_0 \geq V^a(y_0), \tag{15}$$

where we use that $y_0 = b_0$. The promise-keeping constraints in our environment simply state that the contract is consistent with the evolution of the utility and marginal utility state variables (11)-(12). The incentive constraint is the instantaneous truth-telling condition (10) which ensures that if the agent has not lied in the past, he will not do so in the current instant. However to be sure that a contract is fully incentive compatible, we must consider whether the agent may gain by lying now and in the future. Such “double deviations” can be difficult to deal with in general (see Kocherlakota (2004) and Williams (2009) for example), but I show that in our setting they can be ruled out by some further restrictions on Q_t under the contract.

I now provide the key theoretical results characterizing the class of truthful revelation contracts. The theorem adapts my previous results from Williams (2009), which in turn build on Schattler and Sung (1993) and Zhou (1996). However the setting here leads to a more direct proof, which is contained in Appendix A.2. The idea of the proof is to use the representations of the agent's utility and marginal utility processes under the contract, along with the concavity assumptions on the primitives to bound the gain from deviating from truth-telling.

THEOREM 4.1. *Assume that the agent's utility functions u and U are twice differentiable, increasing, and concave in b , and that $\lambda \geq 0$. Then we have:*

1. *Any truthful revelation contract $\{s_t, \gamma_t, Q_t\}$ satisfies (i) the participation constraint (15), (ii) the promise keeping constraints (11)-(12), and (iii) the incentive constraint (10).*

2. *Suppose $\lambda \neq 0$ and u is three times differentiable in b with $u_{bbb} \geq 0$. Then any contract satisfying conditions (i)-(iii) and insuring that the following conditions hold:*

$$Q_t \leq -\frac{u_{bb}(s_t, y_t)}{2\lambda} \tag{16}$$

$$Q_t \leq -E_* \left[\int_t^T e^{-\rho(\tau-t)} [u_{bb}(s_\tau, y_\tau) + 2\lambda Q_\tau] d\tau \middle| \mathcal{F}_t \right] \tag{17}$$

for all t is a truthful revelation contract.

3. *Suppose $\lambda = 0$ and u is three times differentiable in b with $u_{bbb} \geq 0$. Then any contract satisfying conditions (i)-(iii) and insuring that the following condition holds:*

$$Q_t \leq -E_* \left[\int_t^T e^{-\rho(\tau-t)} u_{bb}(s_\tau, y_\tau) d\tau \middle| \mathcal{F}_t \right] \tag{18}$$

for all t is a truthful revelation contract.

Since u is concave in b , the right side of the inequalities (16) and (18) are positive, and (16) implies that the right side of (17) is positive as well. Thus $Q_t \leq 0$ for all t is a stronger, but simpler, sufficient condition. In particular, the incentive constraint (10) and $Q_t \leq 0$ together are sufficient to ensure that the agent's optimality condition (9) holds, since $m_t \leq 0$. However the sufficient conditions (16)-(17) or (18) in the theorem are weaker in that they allow some positive Q_t settings. We show below that these conditions, but not the stronger $Q_t \leq 0$ condition, are satisfied in our examples.

Intuitively, the sufficient conditions trade off level and timing changes in the agent's consumption profile under the contract. As discussed above, the contract links the volatility of the agent's promised utility to his report, spreading out continuation utility in order to insure truthful revelation. However the promised marginal utility state allows the contract to condition the spread of continuation utilities, making the volatility of future continuation utility depend on the current report. In particular, if an agent were to lie and set $\Delta_t < 0$ for some increment $\delta > 0$ of time, his promised future utility $q_{t+\delta}$ would be lowered as $\delta\gamma_t\Delta_t < 0$. If the stronger sufficient condition $Q_t < 0$ holds, then the lie would also increase (i.e. make less negative) the promised marginal utility state $p_{t+\delta}$ as $\delta Q_t\Delta_t > 0$. Since this state variable gives the negative of marginal utility, this is associated with a reduction in expected marginal utility. This roughly means that after a lie an agent could expect lower lifetime consumption overall (from the fall in q), but consumption would grow in the future (from the reduction in future marginal utility). Moreover, from the incentive constraint (10), we see that the current lie would lead to smaller utility dispersion in the future as $\gamma_{t+\delta}$ is now bounded by the smaller quantity $-p_{t+\delta}$. By using the concavity of utility and the continuity of sample paths, we show that contracts which appropriately balance these effects are incentive compatible.

In some cases the sufficient conditions (16)-(18) may be overly stringent or difficult to verify. To be sure that the contract does indeed ensure truthful revelation, we then must re-solve the agent's problem facing the given contract. This is similar to the ex-post incentive compatibility checks that Werning (2001) and Abraham and Pavoni (2008) conduct. Re-solving the agent's problem typically would require numerical methods, as indeed would finding the contract itself. In the examples below we verify that the sufficient conditions do hold, but we also illustrate how to analytically re-solve the agent's problem.

5. OPTIMAL CONTRACTS ON AN INFINITE HORIZON

I now turn to the principal's problem of optimal contract design over an infinite horizon. Formally, I take limits as $T \rightarrow \infty$ in the analysis above. Thus we no longer have the terminal conditions for the co-states in (11) and (12) above, but instead we have the transversality conditions $\lim_{T \rightarrow \infty} e^{-\rho T} q_T = 0$ and $\lim_{T \rightarrow \infty} e^{-\rho T} p_T = 0$. For the purposes of optimal contract design, we can effectively treat the backward

equations governing the dynamics of promised utility and promised marginal utility as forward equations under the control of the principal. Thus the principal's contract choice is a standard stochastic control problem.

5.1. Basic Layout

As is standard, I suppose that the principal's objective function is the expected discounted value of the transfers:

$$J = E_y \left[\int_0^\infty e^{-\rho t} s_t dt \right],$$

where he discounts at the same rate as the agent. The principal's problem is to choose a contract to minimize J subject to satisfying: (i) the participation constraint (15), (ii) the promise keeping constraints (11)-(12), and (iii) the incentive constraint (10). We focus on the relaxed problem, and do not impose the sufficient conditions from Theorem 4.1, but instead check ex-post whether they are satisfied (and if not, we directly check if the contract is incentive compatible). While the participation constraint bounds q_0 , p_0 is effectively free and so we treat it as a choice variable of the principal. Thus a contract consists of $\{s_t, \gamma_t, Q_t\}$ and values for q_0 and p_0 .

I focus on dynamic programming approach to the principal's problem.¹² Abusing notation slightly, denote the principal's value function $J(y, q, p)$, and let $J_y(y, q, p)$ and so on be its partial derivatives. Via standard arguments (see e.g. Yong and Zhou (1999)), the value function satisfies the Hamilton-Jacobi-Bellman equation:

$$\begin{aligned} \rho J = & \min_{\{s, \gamma \geq -p, Q\}} \left\{ s + J_y \mu(y) + J_q [\rho q - u(s, y)] + J_p [\rho p + \gamma \mu'(y) + u_b(s, y)] \right. \\ & \left. + \frac{\sigma^2}{2} [J_{yy} + J_{qq} \gamma^2 + J_{pp} Q^2 + 2(J_{yq} \gamma + J_{yp} Q + J_{pq} \gamma Q)] \right\} \end{aligned} \quad (19)$$

where we suppress the arguments of J and its derivatives. Given the solution to (19), the initial value p_0 is chosen to maximize $J(y_0, q_0, p_0)$.

5.2. Full Information

As a benchmark, we consider first the full information case where the principal observes all state variables. His problem in this case is to maximize J subject only to the participation constraint. It's possible to include this as a constraint at date zero, but to make the analysis more comparable with the previous discussion, we include q_t as a state variable. The principal's costs depend on y_t and q_t , which governs the amount of utility he must deliver to satisfy the participation constraint, but there are no direct costs associated with the marginal utility state p_t .

¹²Cvitanic, Wan, and Zhang (2009) use a stochastic maximum principle approach to characterize contracts in a moral hazard setting. In an earlier draft, I'd used the maximum principle to provide some partial characterizations of the dynamics of optimal contracts.

Denoting the full information value function $J^*(y, q)$, we see that it solves an HJB equation similar to (19):

$$\rho J^* = \min_{\{s, \gamma\}} \left\{ s + J_y^* \mu(y) + J_q^* [\rho q - u(s, y)] + \frac{\sigma^2}{2} [J_{yy}^* + J_{qq}^* \gamma^2 + 2J_{yq}^* \gamma] \right\}$$

The optimality conditions are:

$$J_q^* u_s(s, y) = 1, \quad \gamma = -J_{yq}^* / J_{qq}^*. \quad (20)$$

The first condition balances the instantaneous costs of payments with their impact on lessening future transfers. Since γ governs both the volatility of the utility promise q as well as its covariation with the report y , the second condition balances these effects. Letting $s(y, q)$ be the solution of the first optimality condition, the HJB becomes:

$$\rho J^* = s(y, q) + J_y^* \mu(y) + J_q^* [\rho q - u(s(y, q), y)] + \frac{\sigma^2}{2} \left[J_{yy}^* - \frac{(J_{yq}^*)^2}{J_{qq}^*} \right]. \quad (21)$$

In general the agent's consumption will vary with the state y . However in the hidden endowment case, we obtain the standard result that a risk neutral principal absorbs all the risk, completely stabilizing the agent's consumption. In particular, when $u(s, b) = v(s + b)$, let $\bar{c}(q) = v^{-1}(\rho q)$ be the constant consumption consistent with promised utility q . Then it is straightforward to verify that the solution to (21) is given by:

$$J^*(y, q) = \frac{\bar{c}(q)}{\rho} + j(y) \quad (22)$$

where $j(y)$ solves the second order ODE:

$$\rho j(y) = -y + j'(y) \mu(y) + \frac{1}{2} j''(y) \sigma^2. \quad (23)$$

Below we provide an explicit solution of this ODE for a particular parameterization. Thus the optimal full information contract indeed calls for complete stabilization, setting $s_t = \bar{c}(q_t) - y_t$, and $\gamma_t = 0$. Together these imply that consumption and promised utility are constant under the contract: $c_t = \bar{c}(q_0)$ and $q_t = q_0$ for all t .

5.3. Hidden Information

We now return to the hidden information problem, where it is difficult to characterize the general case in much detail. We now assume that the incentive constraint (10) binds, so that $\gamma = -p$. We relax this condition in an example below and verify that it holds. However we also conjecture that it holds more generally. Under full information, the principal need not worry about incentives, and can provide constant promised utility. With hidden information, utility must vary to provide incentives. But because the agent is risk averse, the principal will typically find it optimal to induce as little volatility in the agent's utility as possible in order to induce truthful revelation, and thus will set γ at the lowest feasible level.

The γ choice is constrained, while the other optimality conditions from (19) are:

$$J_q u_s(s, y) - J_p u_{bs}(s, y) = 1, \quad (24)$$

$$J_{pp} Q + J_{yp} - p J_{qp} = 0. \quad (25)$$

Relative to the full information case, the first order condition for the payment (24) has an additional term coming from the effect of consumption on the marginal utility state p_t . That is, payments affect both the promised utility and promised marginal utility, and hence impact on the agent's future valuations. The condition (25) for Q_t balances the effects of the variability of p_t with the covariations of p_t with y_t and q_t .

We can use these optimality conditions and the HJB equation to solve for a candidate optimal contract. Then we can check whether the sufficient conditions from Theorem 4.1 hold. If not, we would directly solve the agent's reporting problem given the contract to check incentive compatibility. We follow this procedure below.

6. A CLASS OF HIDDEN ENDOWMENT EXAMPLES

In this section, we study a class of examples which allows for explicit solutions. We suppose b is the agent's endowment, and that the agent has exponential utility:

$$u(s, b) = -\exp(-\theta(s + b)).$$

As is well known, exponential utility with linear evolution often leads to explicit solutions, and this is once again the case here. We fully solve for the optimal contract and verify that it is indeed incentive compatible. Then we discuss how persistence affects the contract, and compare our results to those in the literature. Some of the intermediate calculations for this section are in Appendix A.3.

6.1. Full Information

As a benchmark, we first consider the full information case from section 5.2. Inverting the agent's utility function, consumption under the contract is:

$$\bar{c}(q) = -\frac{\log(-\rho q)}{\theta}.$$

It is easy to verify that the solution of the ODE (23) for the principal's cost is:

$$j(y) = -\frac{\mu_0}{\rho(\lambda + \rho)} - \frac{y}{\lambda + \rho}.$$

We then have that the full information cost function is:

$$J^*(y, q) = -\frac{\log(-\rho q)}{\rho\theta} - \frac{\mu_0}{\rho(\lambda + \rho)} - \frac{y}{\lambda + \rho}.$$

Thus the principal's full information cost is linear and decreasing in the endowment y as well as being linear in the consumption equivalent of the promised utility q , which here implies that the cost is logarithmic in q .

6.2. Persistent Endowment

We now suppose that the agent's endowment is private information and is persistent. We show that the contract is similar to the full information case, but now it may depend on the promised marginal utility p_t . However when the endowment is permanent, the ratio p_t/q_t is necessarily constant. This follows from the proportionality of utility and marginal utility with exponential preferences, since (13) and (14) with $\lambda = 0$ imply $p_t = \theta q_t$. Thus we can dispense with p_t as a separate state variable. When $\lambda > 0$, the p_t/q_t ratio may vary over time. However we show that under the optimal contract it is indeed constant, at a level which depends on λ . This is clearly a very special case of our general results, as typically both p_t and q_t would be necessary state variables. However this special case facilitates explicit calculations.

We also show that the degree of risk sharing increases as the endowment becomes less persistent. There is no risk sharing at all in the permanent case, while in the i.i.d. limit the contract allocation converges to full-information with complete consumption stabilization. Thus the distortions inherent in the discrete time i.i.d. environments of Thomas and Worrall (1990) and others vanish in continuous time, as we discuss below.

6.2.1. The Optimal Contract with an Arbitrary Initial Condition

We first consider contracts with an arbitrary initial condition for the marginal utility state p_0 . Later, we consider the choice of the initial value. Thus first consider the principal's problem given the states (y, q, p) . We show in Appendix A.3.1 that the principal's cost function can be written:

$$J(y, q, p) = j_0 + j_1 y - j_2 \log(-q) + h(p/q), \quad (26)$$

for some constants (j_0, j_1, j_2) and some function h which solves a second-order ODE given in the Appendix. Letting $k = p/q$, in the permanent endowment case $k_t = \theta$ for all t , but when $\lambda > 0$ then k_t becomes a state variable. We also show in Appendix A.3.1 that the agent's consumption under the contract c and the diffusion coefficient Q on marginal utility can be written:

$$c(q, k) = -\frac{\log(-q\hat{c}(k))}{\theta}, \quad Q(q, k) = -q\hat{Q}(k)$$

for functions $\hat{c}(k)$ and $\hat{Q}(k)$ which we provide there. The consumption and cost functions are similar to the full information case, but they depend on the ratio of the utility and marginal utility promises k , which in general will vary over time.

Using the previous results, the dynamics of the co-states can be written:

$$\begin{aligned} dq_t &= [\rho - \hat{c}(k_t)]q_t dt - \sigma p_t dW_t, \\ dp_t &= [(\rho + \lambda)p_t - \theta \hat{c}(k_t)q_t] dt - \sigma q_t \hat{Q}(k_t) dW_t. \end{aligned} \quad (27)$$

Here we presume truthful revelation, so the principal effectively observes the true shocks $W_t^y = W_t$. Applying Ito's lemma gives the dynamics of the ratio:

$$dk_t = \left[\hat{c}(k_t)(k_t - \theta) + \lambda k_t + \sigma^2 k_t (k_t^2 - \hat{Q}(k_t)) \right] dt + \sigma (k_t^2 - \hat{Q}(k_t)) dW_t. \quad (28)$$

When $\lambda = 0$ this ratio is necessarily constant, but now it may evolve stochastically. But we show next that if the ratio is initialized optimally, it remains constant.

6.2.2. The Optimal Initial Condition

Recall that the marginal utility state p_t is an endogenous, backward stochastic differential equation. Thus its initial condition is not specified, but is instead free to be chosen by the principal. While the optimal choice of k_0 for $\lambda > 0$ is difficult to establish analytically, in Appendix A.3.1 we verify numerically that the optimal initial condition is $k_0^* = \frac{\rho\theta}{\rho+\lambda}$. This result is intuitive, as it is proportional to the ratio of the discount rates in the utility and marginal utility state variables. Note that although the choice of initial condition required numerical methods, conditional on it, the rest of our results are analytic. Given k_0^* the expressions in the appendix then imply $\hat{c}(k_0^*) = \rho$ and $\hat{Q}(k_0^*) = (k_0^*)^2$, which in turn imply:

$$c(q, k_0^*) = -\frac{\log(-\rho q)}{\theta} = \bar{c}(q), \quad Q(q, k_0^*) = -q(k_0^*)^2.$$

Moreover, from (28) we see that the promised utility and marginal utility states remain proportional throughout the contract: $k_t = k_0^*$ for all t .

The principal's cost function is nearly the same as the full information case, with the only change being an additional additive constant term. In particular, using the expressions in Appendix A.3.1 and letting $p_0^* = k_0^*q_0$, we have:

$$J(y_0, q_0, p_0^*) = J^*(y_0, q_0) + \frac{\sigma^2\theta}{2(\rho + \lambda)^2}.$$

The additional cost of not observing the endowment is increasing in the local variance of the endowment, which is a measure of the “size” of the private information. The cost is also increasing in the agent's risk aversion parameter θ , reflecting the additional expected payments which must be made to compensate for risk, and decreasing in rate of time preference, reflecting the effective horizon of the payments. Finally, the additional cost declines as λ increases and the endowment becomes less persistent.

The optimal contract has striking implications for the dynamics of consumption and utility. In particular, consumption is the same function of promised utility whether the agent's endowment is private information or not. However in the full information case promised utility is constant, while with private information it varies. In fact, using $\hat{c} = \rho$ and $p_t = k_0^*q_t$ in (27) we have:

$$dq_t = -\frac{\sigma\rho\theta}{\rho + \lambda}q_t dW_t, \tag{29}$$

and thus promised utility follows a martingale. This equation can be solved explicitly, leading to an explicit solution for consumption as well:

$$\begin{aligned} q_t &= q_0 \exp\left(-\frac{\sigma^2\rho^2\theta^2}{2(\rho + \lambda)^2}t - \frac{\sigma\rho\theta}{\rho + \lambda}W_t\right) \\ c_t &= \bar{c}(q_0) + \frac{\sigma^2\theta\rho^2}{2(\rho + \lambda)^2}t + \frac{\sigma\rho}{\rho + \lambda}W_t \end{aligned}$$

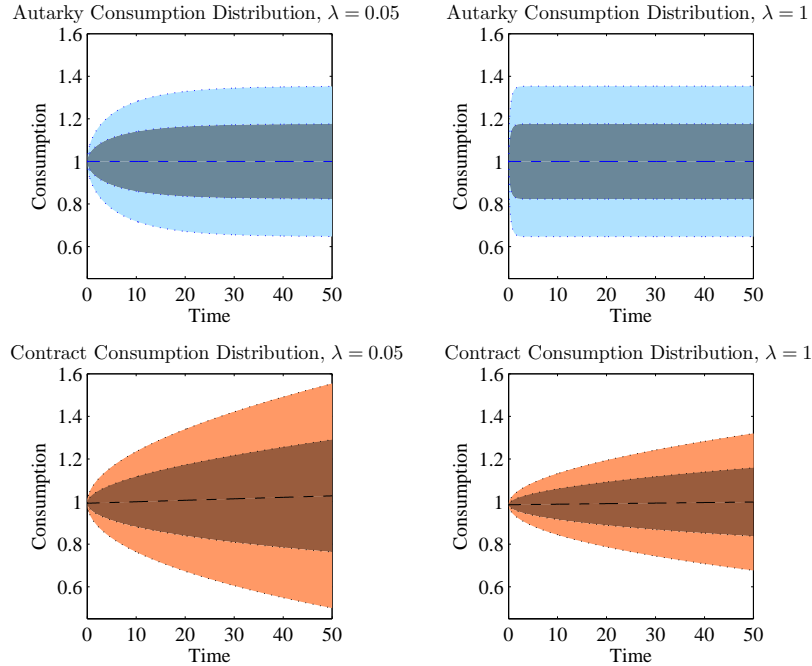


FIGURE 1. Distributions of consumption under the optimal contract and under autarky for two different values of λ .

Therefore consumption follows a Brownian motion with drift, growing over time in order to provide risk compensation for the increasing variability. As consumption tends to grow over time, promised utility tends toward its upper bound of zero.

The volatility of consumption provides one measure of the degree of risk sharing that the contract provides. When $\lambda = 0$ the endowment and consumption are both Brownian motions with drift, having the same local variance σ^2 . Thus when the endowment is permanent, the contract provides no risk sharing. It only alters the time path of consumption, providing more consumption up front in exchange for a lower average growth rate. However as information becomes less persistent, the local variance of consumption falls and the amount of risk sharing increases. In particular, in the limit approaching an i.i.d. endowment (with $\sigma = \bar{\sigma}\sqrt{\lambda}$ and $\lambda \rightarrow \infty$) the optimal contract with private information converges to the efficient, full information allocation.

The results are illustrated in Figure 1, which plots the distributions of consumption under autarky and under the contract with two different values of λ . In each case we show the mean of the distribution along with one and two standard deviation bands. We scale the drift and diffusion parameters with λ to maintain a constant unconditional mean and variance.¹³ The top two panels plot the distributions of the endowment, and hence consumption under autarky. The endowment is stationary, so after being initialized at a common point (the unconditional mean) the variance of the distribution grows until it reaches the stationary distribution. Clearly with less persistence (larger λ) this convergence is much faster. The bottom two panels plot the consumption

¹³In particular, we choose $\theta = 1$, $\rho = 0.1$, $\bar{\sigma} = 0.25$, and $\bar{\mu} = 1$.

distribution under the optimal contract. Here we see that early on the distribution of consumption under the contract is tighter than under autarky, reflecting the risk sharing properties of the contract. Clearly with less persistence there is more risk sharing, so the consumption distribution is even more compressed. However unlike the autarky case, consumption is nonstationary under the contract. Consumption has an upward trend (which is difficult to see in the figure) and its distribution fans out over time.

6.3. Verifying Incentive Compatibility

In Appendix A.3.2 we verify that the sufficient conditions (16) and (17) from Theorem 4.1 hold. Thus we know the contract derived above is indeed the optimal incentive-compatible contract. However, directly verifying that the contract is incentive compatible helps to understand how the contract provides incentives for revelation. When the $\lambda = 0$, this is straightforward. If the agent were to deviate from truth-telling and report a lower endowment, then he would receive a permanently lower promise of future utility and hence future consumption. This would exactly balance the increased consumption he would gain due to his lower transfer payment to the principal.

Since we allow the possibility of lying, we distinguish between the true endowment shock W_t and the principal's prediction of it W_t^y . Under the optimal contract and an arbitrary reporting strategy we have:

$$\begin{aligned}
 c_t &= s(q_t, y_t) + y_t - m_t \\
 &= -\log(-\rho q_t)/\theta - y_t + y_t - m_t \\
 &= \bar{c}(q_0) + \frac{\sigma^2\theta}{2}t + \sigma W_t^y - m_t \\
 &= \bar{c}(q_0) + \frac{\sigma^2\theta}{2}t + \int_0^t [dy_\tau - \mu_0 d\tau] - m_t \\
 &= \bar{c}(q_0) + \frac{\sigma^2\theta}{2}t + \int_0^t [db_\tau + (\Delta_\tau - \mu_0)d\tau] - m_t \\
 &= \bar{c}(q_0) + \frac{\sigma^2\theta}{2}t + \sigma W_t,
 \end{aligned}$$

where we've used $b_t = b_0 + \mu_0 t + \sigma W_t$ and the definition of m_t . Thus the agent's consumption is independent of his reporting process y_t , depending only on his initial utility promise, a deterministic transfer, and the underlying shocks to his endowment. Put differently, the contract ensures that the agent receives the same consumption whatever he reports. Thus the contract is consistent with truthful revelation, although reporting truthfully is by no means the unique best response of the agent.

When $\lambda > 0$, directly verifying incentive compatibility requires fully solving the agent's problem facing the contract. These more complicated calculations are carried out in Appendix A.3.2.

7. RELATIONSHIP WITH PREVIOUS LITERATURE

7.1. General Discussion

My results in the previous section are in sharp contrast to the discrete time models with i.i.d. private information. As discussed above, Thomas and Worrall (1990) show that the optimal contract leads to immiseration, with the agent's promised utility tending toward minus infinity. But under the optimal contract here consumption grows on average over time. In both our setting and theirs, the distribution of promised utility fans out over time in order to provide incentives. However in our setting, the Brownian motion shocks driving the endowment process have a distribution which naturally fans out over time. Therefore by linking the agent's promised utility (and hence consumption) to the reported endowment shocks, the principal is able to achieve the fanning out of the utility distribution which is required to provide incentives, yet still have consumption increase over time. The speed at which the utility distribution fans out depends on the persistence of the information. As the endowment becomes less persistent, the effective life of the private information is shorter, and smaller increases in the utility dispersion are required to ensure truthful reporting.

To provide a bit more detail, note that we can define the cumulative endowment $Y_t = \int_0^t y_u du$ and the cumulative transfers $S_t = \int_0^t s_u du$ as ordinary (not stochastic) integrals. By requiring Y_t and S_t to thus be absolutely continuous we are able to define utility over their flows. If Y_t were of unbounded variation, say by specifying $Y_t = \int_0^t y_u du + \int_0^t \sigma_u^Y dW_u$ then it would be unclear how to define utility over the flow, except in the risk neutral case studied by DeMarzo and Sannikov (2006). While this technical difference may seem relatively minor, it implies that as we approach the i.i.d. limit the information frictions vanish. In particular, our limit result is driven by fact the diffusion term in (29) vanishes as $\lambda \rightarrow \infty$ since σ increases with $\sqrt{\lambda}$. If σ increased with λ then there would be nonzero frictions in the limit, but the variance of the endowment process would explode.

The efficient limit result is also implicit in Thomas and Worrall (1990), a point we make explicit in the next section. In the course of establishing that as the discount factor tends to unity the contract converges to the first best (their Proposition 4), they show that the deviations from efficiency are tied to the cost of inducing the efficient actions for one period. Rather than letting discount rates decline as in their limit, our results effectively let the period length shrink to approach continuous time. But the consequence is the same – the inefficiency vanishes. Thus as periods become shorter, deviations from efficiency are sustainable if either the cumulative endowment has unbounded variation, which causes conceptual problems, or the endowment is persistent, which is quite natural. We illustrate the efficiency of the continuous time limit in Thomas and Worrall (1990) in the next section.

As a final comparison, note that although the inverse Euler equation does not hold in our environment, it does hold in the continuous time moral hazard model in Williams (2009). In the moral hazard (hidden effort) setting, the contract specifies a consumption payment to the agent conditional on realized output. However effort is costly to the agent, and thus the incentive constraints reflect the instantaneous effect of the information friction. By contrast, in the hidden information model lying has no instantaneous effect on the agent – it only affects his future expected transfers. Thus the incentive constraints are fully forward looking, implying constraints on the evolution of promised utility. Thus whether the inverse Euler equation holds does not

depend on whether time is continuous or discrete, but it is sensitive to the source of the information friction.

7.2. Efficiency in a Limit of Thomas and Worrall (1990)

In this section we demonstrate that as the period length shrinks to zero in the discrete time model of Thomas and Worrall (1990), the contract converges to the efficient allocation. As discussed previously, such a limit is implicit in their paper, here we make the argument explicit for the case of exponential preferences and binary endowment realizations.

In particular, suppose now that time is discrete with interval ε between periods. The agent's endowment realization b_t each period is i.i.d., and can take on values b_1 with probability 0.5 and $b_2 > b_1$ with probability 0.5. The restriction to equal probabilities of high and low realizations helps preserve the link with our continuous time model, where over short intervals upward and downward movements are equally likely. The agent has exponential preferences with parameter θ as above and discount factor $\alpha = \exp(-\rho\varepsilon)$. Each period, upon receiving a report of b_i , the contract specifies a transfer s_i and a utility promise (now relative to autarky) of q_i for the next period. To simplify some notation, let $\bar{u} = 0.5[u(b_1) + u(b_2)]$, and define the value of autarky as $A = \bar{u}\varepsilon/(1 - \alpha)$. In their Proposition 6, Thomas and Worrall (1990) show that in this setting the optimal contract can be represented as a set of numbers $0 \leq a_2 \leq a_1$ and $d_1 \geq d_2 \geq 0$ which satisfy:

$$\begin{aligned} \exp(-\theta s_i) &= -a_i(q + A), \\ q_i &= d_i q + (d_i - 1)A \\ \frac{0.5}{a_1} + \frac{0.5}{a_2} &= -A \\ \frac{0.5}{d_1} + \frac{0.5}{d_2} &= 1, \end{aligned}$$

along with the (binding) downward incentive constraint:

$$u(s_2 + b_2)\varepsilon + \alpha q_2 = u(s_1 + b_2)\varepsilon + \alpha q_1$$

and the promise-keeping constraint:

$$0.5[u(s_1 + b_1)\varepsilon - u(b_1)\varepsilon + \alpha q_1] + 0.5[u(s_2 + b_2)\varepsilon - u(b_2)\varepsilon + \alpha q_2] = q.$$

To solve for the optimal contract, we first substitute for s_i and q_i in terms of a_i and d_i , then use the normalizations of a_i and d_i to eliminate a_2 and d_2 . Finally, we combine the incentive and promise-keeping constraints to obtain a single equation in a_1 :

$$-\exp(-\theta b_2) \left(\frac{0.5}{\frac{\bar{u}}{1-\alpha} + \frac{0.5}{a_1\varepsilon}} + a_1\varepsilon \right) + \frac{0.5\alpha}{1 - \frac{0.5\alpha}{1 + \bar{u}a_1\varepsilon}} - (1 + \bar{u}a_1\varepsilon) = 0. \quad (30)$$

The solution of equation (30), which is a cubic equation in a_1 , determines the optimal contract. After some simplification, it can be written:

$$0.5 \frac{(1-\alpha)^2}{\varepsilon^2} + [\exp(-\theta b_2)(1-0.5\alpha) + (1+.5(2-\alpha))\bar{u}] \frac{(1-\alpha)}{\varepsilon} a_1 + \bar{u} [\exp(-\theta b_2)(2-1.5\alpha) + \bar{u}(2.5-1.5\alpha)] a_1^2 + \bar{u}^2 [\exp(-\theta b_2) + \bar{u}] \varepsilon a_1^3 = 0$$

We now find the limiting solution as the period length shrinks. We use the facts that $\lim_{\varepsilon \rightarrow 0} \frac{(1-\alpha)}{\varepsilon} = \rho$, $\lim_{\varepsilon \rightarrow 0} \alpha = 1$, and note that the cubic term is of order ε so it vanishes in the limit. Therefore the limit contract solves the quadratic equation:

$$0.5\rho^2 + [0.5 \exp(-\theta b_2) + 1.5\bar{u}] \rho a_1 + [1.5 \exp(-\theta b_2) + \bar{u}] a_1^2 = 0.$$

The cost-minimizing root of this equation is then:

$$a_1 = \rho \exp(\theta b_1).$$

Then note that utility in state 1 under the contract is:

$$u(s_1 + b_1) = -\exp(-\theta b_1) \exp(\theta s_1) = \exp(-\theta b_1) a_1 (q + A) = \rho(q + A),$$

with the same result for state 2. So consumption and utility are completely smoothed under the contract, and promised utility remains constant ($d_1 = d_2 = 1$). Thus the continuous time limit of the discrete time i.i.d. model of Thomas and Worrall (1990) agrees with the i.i.d. limit of our continuous time persistent information model above, with both implying efficiency in the limit.

8. A PRIVATE TASTE SHOCK EXAMPLE

We now turn to an example where the agent has private information about a preference shock which affects his marginal utility of consumption. Similar discrete time models have been studied by Atkeson and Lucas (1992) with i.i.d. shocks, and Kapicka (2006) with persistent private information. We now suppose that b is the logarithm of a private taste shock, and the agent has power utility over consumption:

$$u(s, b) = \exp(-b) \frac{s^{1-\theta}}{1-\theta},$$

where the coefficient of relative risk aversion is $\theta > 1$.¹⁴ Clearly this specification satisfies all of our assumptions, as $u_b > 0$, $u_{bb} < 0$, and $u_{bbb} > 0$.

This example is similar to the previous one, with some differences that we discuss below. For simplicity, we focus here on the permanent shock case as this allows for explicit solutions. Since $u_b = -u$ in this model, (11)-(12) imply that $p = q$, and thus we can dispense with p as a separate state variable. We also assume that there is no trend in the taste shocks and $E(\exp(-b)) = 1$, which implies that $\mu_0 = \frac{\sigma^2}{2}$.

¹⁴We restrict attention to $\theta > 1$ so that $u < 0$. Our results hold more generally.

8.1. Full Information

As before, we begin with the full information benchmark. Using the parameterization above, we show in Appendix A.3.3 that the optimality conditions (20) and HJB equation (21) can be solved explicitly:

$$\begin{aligned} J^*(y, q) &= A \exp\left(\frac{1}{1-\theta}y\right) ((1-\theta)q)^{\frac{1}{1-\theta}}, \\ s(y, q) &= A^{\frac{1}{\theta}} \exp\left(\frac{1}{1-\theta}y\right) ((1-\theta)q)^{\frac{1}{1-\theta}}, \\ \text{where: } A &\equiv \left(\rho + \frac{\sigma^2(\theta-1)}{2\theta^2}\right)^{\frac{\theta}{1-\theta}}. \end{aligned}$$

The flow utility is then proportional to promised utility q and independent of y :

$$u(s(y, q), y) = \frac{\exp(-y)}{1-\theta} \left(A^{\frac{1}{\theta}} \exp\left(\frac{1}{1-\theta}y\right) ((1-\theta)q)^{\frac{1}{1-\theta}} \right)^{1-\theta} = \left(\rho + \frac{\sigma^2(\theta-1)}{2\theta^2} \right) q.$$

Instead of fully stabilizing consumption and utility, as the full-information contract did in the hidden endowment example, here the optimal contract offsets the impact of the taste shocks on the current utility flow. Since the taste shocks are stochastic, promised utility is then no longer constant. In particular, from (20) we have:

$$\gamma = -\frac{J_{yq}^*}{J_{qq}^*} = -\frac{1}{\theta}q \quad (31)$$

Therefore the evolution of promised utility (11) is:

$$dq_t = -\frac{\sigma^2(\theta-1)}{2\theta^2}q_t dt - \frac{\sigma}{\theta}q_t dW_t.$$

Thus since $q_t < 0$ and $\theta > 1$, promised utility has a positive drift and thus tends to increase over time. The taste shocks get translated into fluctuations in promised utility, and the larger is the degree of risk aversion θ , the smaller is the volatility of promised utility.

8.2. Private Information

When the agent's taste shocks are private information, the principal must provide incentives for truthful revelation. He does this by making promised utility respond more to new observations than it does under full information. In particular, the incentive constraint (10) here becomes: $\gamma = -p = -q$, again assuming the constraint binds. Since $\theta > 1$, comparison with (31) shows that an incentive-compatible contract clearly provides less insurance against taste shocks than the first best contract.

As we show in Appendix A.3.3, the solution of the HJB equation (19) is very similar to the full information case, with a different leading constant:

$$\begin{aligned} J(y, q) &= B \exp\left(\frac{1}{1-\theta}y\right) ((1-\theta)q)^{\frac{1}{1-\theta}}, \\ s(y, q) &= B^{\frac{1}{\theta}} \exp\left(\frac{1}{1-\theta}y\right) ((1-\theta)q)^{\frac{1}{1-\theta}}, \end{aligned} \quad (32)$$

where: $B \equiv (\rho)^{\frac{\theta}{1-\theta}}$.

Moreover, since $\theta > 1$ and $\frac{\theta}{1-\theta} < 0$ we see that $B > A$. Thus, as one would expect, the principal's costs $J(y, q)$ of delivering a given level of promised utility q given a truthful report of y are higher under private information than with full information. In addition, the agent obtains a larger transfer s under private information, as he must be compensated for bearing additional risk. Once again, the flow utility is proportional to promised utility q , but now with proportionality factor ρ . This implies that under the optimal contract, promised utility is a martingale:

$$dq_t = -\sigma q_t dW_t. \quad (33)$$

This is just like the hidden endowment case (29) with $\lambda = 0$. Once again, with permanent shocks there is a complete lack of risk sharing, as the full magnitude of the shocks to the private state get translated into shocks to promised utility.

In fact, the contract can be implemented via a constant payment $s_t = \bar{s}$, with \bar{s} chosen to deliver the initial promised utility q_0 . In turn, this insures that:

$$q_t = \frac{\bar{s}^{1-\theta}}{(1-\theta)\rho} \exp(-y_t),$$

which satisfies (33), and when substituting in (32) gives indeed that $s(y, q) = \bar{s}$. Thus the agent's consumption is independent of his reports y_t , and there is no risk-sharing. As above, we could verify the sufficient condition (18) ensuring implementability, but the constancy of the payment s_t yields it directly. As the payment is independent of the report, it is clearly (weakly) consistent with truthful revelation.

Although direct comparisons are a little difficult given the differences in the setups, my results differ substantially from Kapicka (2006). We both use similar ideas, using a first-order approach to contracting to justify a recursive representation of a contract with additional endogenous state variables. However, our state variables differ. Kapicka (2006) takes the agent's promised marginal utility of *consumption* as his additional state variable, while in my approach the marginal utility of *the hidden state* is the relevant state variable. With the multiplicative permanent taste shocks in this example, the promised marginal utility of the hidden state is equal to the level of promised utility, and hence is superfluous. More generally, my state variable captures the lifespan of the private information, as the promised marginal utility is discounted by the persistence of the hidden state.

Our implications for contracts differ as well. In a numerical example similar to the one in this section, Kapicka (2006) shows how the contract may significantly

distort the evolution of the agent’s marginal utility of consumption. This leads to a sizeable “intertemporal wedge” distorting the savings margin of the type discussed by Golosov, Kocherlakota, and Tsyvinski (2003) among others. But in our case, where we obtain explicit analytic solutions for the optimal contract, consumption is deterministic. This implies that the marginal utility of consumption is a constant multiple of promised utility, and so is a martingale. Thus we obtain a standard Euler equation, with no intertemporal wedge. The reasons for these differences are largely the same as in the hidden endowment example above. Kapicka’s results are driven by the interaction between current transfers and future promises, with the response of current consumption to a report playing a key role. But as we’ve seen, a lie cannot instantaneously affect the agent’s consumption in continuous time, so all incentive provision is loaded on promised utility. Since the shocks are permanent, a current lie does not affect the evolution of future reports. In this example, this means that a lie effectively has a multiplicative, scale effect on the reported taste shock process, and hence on utility. With the preferences in this example, future behavior is unaffected by a permanent multiplicative scale effect on utility. Thus the only way the principal can ensure truthful revelation is to make the transfer independent of the report. Such a contract does not lead to an intertemporal wedge, but it clearly does not insure the agent against the taste shocks either.

Clearly some of my results are special to the permanent shock case, as with persistent but not permanent shocks a current lie would influence future reports. Kapicka (2006) does consider intermediate cases of persistent but not permanent shocks, and shows numerically how persistence affects the results. While we were able to do so analytically in the hidden endowment example above, such results in this setting would require numerical methods.

9. CONCLUSION

In this paper I have developed some relatively general methods to study contracting problems with persistent private information. I have shown how the methods can lead to explicit solutions in examples. By casting the model in continuous time, I was able to use powerful tools from stochastic control. These allowed me to deduce that an optimal contract must condition on two additional endogenous state variables, the agent’s promised utility and promised marginal utility under the contract. While the use of promised utility as a state variable is now widely applied, the marginal utility state variable is more novel, although it does have some precedents in the literature.

My main results use the representation of the contract derived from analyzing the agent’s necessary conditions for optimality. The key theoretical results in the paper provide some sufficient conditions which guarantee that the contract does indeed provide incentives for truthful revelation. In my examples, I show that these conditions are verified. But in other settings, they may be more difficult to establish. In such cases, my main results could still be used, although one may then have to verify directly incentive compatibility. Apart from special cases like my examples, such solutions would require numerical methods. This is not surprising as typical dynamic decision problems, not to mention the more complex contracting problems we consider here, require numerical solution methods. In any case, there are an array of well developed

numerical methods for solving partial differential equations like those which result from our analysis.

In addition to laying out a framework, I have established some substantive results on the nature of continuous time contracts with persistent private information. The examples highlight the close link between the persistence of the private information and the size of efficiency losses this information causes. In particular when shocks are permanent distortions are the largest and risk sharing may break down, while in the i.i.d. limit we obtain efficiency. Although this efficient limit differs from the discrete time i.i.d. models in the literature, it is inherent in those models as period length shrinks. Correspondingly, I have shown that the “inverse Euler equation” of Rogerson (1985) and Golosov, Kocherlakota, and Tsyvinski (2003) need not hold, and the related immiseration results of Thomas and Worrall (1990) may fail. Both of these are sensitive to the details of the contracting problem, depending on how deviations from the contracted behavior affect utility and the state evolution.

APPENDIX

A.1. DERIVATION OF THE CO-STATE EVOLUTION

Given the change of variables, the evolution of the co-states follows from the maximum principle in Bismut (1973)-(1978). The maximum principle also requires some smoothness and regularity conditions, and a linear growth condition on μ :

$$|\mu(y)| \leq K(1 + |y|), \text{ for some } K.$$

All of these conditions hold under the affine specification that we focus on.

The text spells out the general co-state evolution in the lines preceding (11)-(12). Carrying out the differentiation, we have (suppressing arguments of functions):

$$\begin{aligned} \frac{\partial \mathcal{H}(\Gamma, z)}{\partial \Gamma} &= u + \gamma(\mu + \Delta) + p\Delta + u_b \frac{z}{\Gamma} + (\Gamma\gamma + Qz)\mu' \frac{z}{\Gamma^2} \\ \frac{\partial \mathcal{H}(\Gamma, z)}{\partial z} &= -u_b + Q(\mu + \Delta) - \left(\gamma + Q\frac{z}{\Gamma}\right)\mu'. \end{aligned}$$

Thus under truthful revelation $\Gamma = 1, z = \Delta = 0$ we have:

$$\begin{aligned} \frac{\partial \mathcal{H}(\Gamma, z)}{\partial \Gamma} &= u + \gamma\mu \\ \frac{\partial \mathcal{H}(\Gamma, z)}{\partial z} &= -u_b + Q\mu - \gamma\mu'. \end{aligned}$$

Substituting these above evolution and using the change of measure (6) gives:

$$\begin{aligned} dq_t &= [\rho q_t - u(s_t, y_t) - \gamma_t \mu(y_t)] dt + \gamma_t \sigma dW_t^0 \\ &= [\rho q_t - u(s_t, y_t)] dt + \gamma_t \sigma dW_t^*, \\ dp_t &= [\rho p_t + u_b(s_t, y_t) - Q\mu(y_t) + \gamma_t \mu'(y_t)] dt + Q_t \sigma dW_t^0 \\ &= [\rho p_t + u_b(s_t, y_t) + \gamma_t \mu'(y_t)] dt + Q_t \sigma dW_t^*, \end{aligned}$$

then using $\mu' = -\lambda$ we have (11)-(12) in the text.

A.2. PROOF OF THEOREM 4.1

The necessity of the conditions is obvious, as they were derived from the agent's necessary optimality conditions. For the converse, we first use the representations of the agent's utility and marginal utility under the contract. Then we evaluate the potential gain from deviating from truth-telling, which involves calculating expected utility under an alternative reporting strategy. Using the concavity of u and U , we bound the utility difference by a linear approximation. We then use the incentive constraint to further bound the utility gain. The last steps of the proof hold fixed the alternative report and show that even if the agent had lied in the past, under the the stated sufficient conditions his utility gain from lying in the future is negative.

Although it is natural for discussion and computation to work with the discounted utility and marginal utility processes $\{q_t, p_t\}$ defined in (11)-(12), it is easier here to work with the undiscounted processes: $\tilde{q}_t = e^{-\rho t}q_t$, $\tilde{p}_t = e^{-\rho t}p_t$. Then making this substitution into (11) and integrating gives:

$$e^{-\rho T}U(S_T, y_T) = \tilde{q}_T = q_0 - \int_0^T e^{-\rho t}u(s_t, y_t)dt + \int_0^T e^{-\rho t}\gamma_t\sigma dW_t^*. \quad (\text{A.1})$$

Further, using (12) and (3) along with the substitution for \tilde{p}_t gives:

$$\tilde{p}_T m_T = \int_0^T e^{-\rho t} [p_t \Delta_t - \lambda m_t \gamma_t + m_t u_b(s_t, y_t)] dt + \int_0^T e^{-\rho t} Q_t m_t \sigma dW_t^*. \quad (\text{A.2})$$

For an arbitrary reporting policy Δ resulting in $\hat{y} = y - m$, from (6) we have:

$$\sigma dW_t^* = \sigma dW_t^\Delta + [\mu(y_t - m_t) + \Delta_t - \mu(y_t)] dt. \quad (\text{A.3})$$

Now for this arbitrary reporting policy, we wish to compute the gain of deviating from the truthful reporting strategy y (which gives promised utility q_0):

$$\begin{aligned} V(\hat{y}) - q_0 &= E_\Delta \left[\int_0^T e^{-\rho t} [u(s_t, y_t - m_t) - u(s_t, y_t)] dt + \int_0^T e^{-\rho t} \gamma_t \sigma dW_t^* \right] \\ &\quad + E_\Delta [e^{-\rho T} [U(S_T, y_T - m_T) - U(S_T, y_T)]], \end{aligned} \quad (\text{A.4})$$

where we've used (A.1). Now by the concavity of U we have:

$$\begin{aligned} E_\Delta [e^{-\rho T} [U(S_T, y_T - m_T) - U(S_T, y_T)]] &\leq -E_\Delta [e^{-\rho T} U_b(S_T, y_T) m_T] \\ &= E_\Delta [\tilde{p}_T m_T] \end{aligned} \quad (\text{A.5})$$

Hence combining (A.5) and (A.2) with (A.4) we get:

$$\begin{aligned}
V(\hat{y}) - q_0 &\leq E_\Delta \left[\int_0^T e^{-\rho t} [u(s_t, y_t - m_t) - u(s_t, y_t) + m_t u_b(s_t, y_t)] dt \right] \\
&+ E_\Delta \left[\int_0^T e^{-\rho t} [p_t \Delta_t - \lambda m_t \gamma_t] dt + \int_0^T e^{-\rho t} \sigma [\gamma_t + Q_t m_t] dW_t^* \right] \\
&= E_\Delta \left[\int_0^T e^{-\rho t} [u(s_t, y_t - m_t) - u(s_t, y_t) + m_t u_b(s_t, y_t)] dt \right] \\
&+ E_\Delta \left[\int_0^T e^{-\rho t} [p_t \Delta_t - \lambda m_t \gamma_t + (\gamma_t + Q_t m_t) (\mu(y_t - m_t) + \Delta_t - \mu(y_t))] dt \right] \quad (\text{A.6})
\end{aligned}$$

Here the equality uses the change of variables from (A.3) and the fact that the stochastic integral with respect to W^Δ has expectation zero with respect to P_Δ .

Next we use the incentive constraint (10) to eliminate the $p_t \Delta_t$ and $\gamma_t \Delta_t$ terms from (A.6), as their sum is bounded by zero. In addition, the functional form for μ gives:

$$\gamma_t (\mu(y_t - m_t) - \mu(y_t) - \lambda m_t) = 0, \quad \text{and:} \quad Q_t m_t (\mu(y_t - m_t) - \mu(y_t)) = \lambda Q_t m_t^2$$

Using these results and regrouping terms in (A.6) we have:

$$\begin{aligned}
V(\hat{y}) - q_0 &\leq E_\Delta \left[\int_0^T e^{-\rho t} [u(s_t, y_t - m_t) - u(s_t, y_t) + m_t u_b(s_t, y_t) + \lambda Q_t m_t^2 + Q_t m_t \Delta_t] dt \right] \\
&= E_\Delta \left[\int_0^T e^{-\rho t} [u(s_t, \hat{y}_t) - u(s_t, \hat{y}_t + m_t) + m_t u_b(s_t, \hat{y}_t + m_t)] dt \right] \\
&+ E_\Delta \left[\int_0^T e^{-\rho t} (\lambda Q_t m_t^2 + Q_t m_t \Delta_t) dt \right], \quad (\text{A.7})
\end{aligned}$$

where the equality follows from the definition of \hat{y} . By the concavity of u , we know that the sum of the first three terms in (A.7) is negative. But to bound the entire expression, the terms in m_t on the last line cause some additional complications. However when $Q_t \leq 0$ then we have both $Q_t m_t^2 \leq 0$ and $Q_t m_t \Delta_t \leq 0$. Thus the utility gain from lying is bounded by zero and the contract ensures truthful revelation.

We now develop the weaker sufficient conditions stated in the theorem which allow $Q_t > 0$. Since the expectation in (A.7) is calculated with respect to P_Δ , we can treat \hat{y} as fixed and vary m . To make clear that we are fixing the measure over \hat{y} (but now varying Δ) we denote $P_{\hat{y}} = P_\Delta$ for the particular Δ resulting in \hat{y} . In other words, we have changed the measure in calculating utility under the reporting strategy \hat{y} , and we now fix that measure but vary m . Thus we are implicitly varying the truth y , but we have eliminated the dependence of the problem on y . That is, we now solve:

$$\sup_{\{\Delta_t\}} E_{\hat{y}} \left[\int_0^T e^{-\rho t} [u(s_t, \hat{y}_t) - u(s_t, \hat{y}_t + m_t) + m_t u_b(s_t, \hat{y}_t + m_t) + \lambda Q_t m_t^2 + Q_t m_t \Delta_t] dt \right]$$

subject to (3) and $m_0 = 0$. By (A.7) we know that this maximized value gives an upper bound on $V(\hat{y}) - q_0$.

The Hamiltonian for this relatively standard stochastic control problem, with co-state ξ associated with m , is:

$$H^*(m, \Delta) = u(s, \hat{y}) - u(s, \hat{y} + m) + mu_b(s, \hat{y} + m) + \lambda Qm^2 + Qm\Delta + \xi\Delta. \quad (\text{A.8})$$

Thus the optimality condition for truth-telling ($\Delta_t = 0$) is:

$$Q_t m_t + \xi_t \geq 0. \quad (\text{A.9})$$

Unlike the incentive constraint (10), we want to be sure that (A.9) holds for all m_t . That is, we provide conditions to ensure that even if the agent had lied in the past, he will not want to do so in the future.

From Haussmann (1986), we know that (A.9) is sufficient for optimality when the maximized Hamiltonian $H^*(m, 0)$ is concave in m . Note that we have:

$$\begin{aligned} H_{mm}^*(m, 0) &= u_{bb}(s, \hat{y} + m) + mu_{bbb}(s, \hat{y} + m) + 2\lambda Q \\ &= u_{bb}(s, y) + mu_{bbb}(s, y) + 2\lambda Q. \end{aligned}$$

Now using the assumption $u_{bbb} \geq 0$ we know $mu_{bbb} \leq 0$. Thus we have:

$$H_{mm}^*(m, 0) \leq u_{bb}(s, y) + 2\lambda Q.$$

Then the required condition (16) in part 2 of the theorem guarantees that $H^*(m, 0)$ is concave. When $\lambda = 0$, $H_{mm}^*(m, 0)$ is bounded above by $u_{bb}(s, y)$. So for part 3, the concavity of u assures that $H^*(m, 0)$ is concave without any further restrictions.

We find the evolution of the co-state ξ by differentiating the Hamiltonian:

$$\begin{aligned} d\xi_t &= [\rho\xi_t - H_m^*(m, 0)]dt + \eta_t dW_t^{\hat{y}} \\ &= [\rho\xi_t - m_t(u_{bb}(s_t, \hat{y}_t + m_t) + 2\lambda Q_t)]dt + \eta_t dW_t^{\hat{y}} \\ \xi_T &= 0. \end{aligned}$$

As in (13) and (14), the solution of this can be written:

$$\xi_t = E_{\hat{y}} \left[\int_t^T e^{-\rho(\tau-t)} m_s (u_{bb}(s_\tau, \hat{y}_\tau + m_\tau) + 2\lambda Q_\tau) d\tau \middle| \mathcal{F}_t \right].$$

But the optimality condition (A.9) ensures that $\Delta_\tau = 0$ and thus $m_\tau = m_t$ for $\tau \geq t$. So then we have:

$$\xi_t = m_t E_{\hat{y}} \left[\int_t^T e^{-\rho(\tau-t)} (u_{bb}(s_\tau, \hat{y}_\tau + m_t) + 2\lambda Q_\tau) d\tau \middle| \mathcal{F}_t \right].$$

Substituting this into (A.9) gives:

$$Q_t m_t + m_t E_{\hat{y}} \left[\int_t^T e^{-\rho(\tau-t)} (u_{bb}(s_\tau, \hat{y}_\tau + m_t) + 2\lambda Q_\tau) d\tau \middle| \mathcal{F}_t \right] \geq 0.$$

When $m_t = 0$ this condition clearly holds, which is to say that if the agent never lied he will not find it optimal to do so. But when $m_t < 0$, this implies:

$$Q_t \leq -E_{\hat{y}} \left[\int_t^T e^{-\rho(\tau-t)} (u_{bb}(s_\tau, \hat{y}_\tau + m_t) + 2\lambda Q_\tau) d\tau \middle| \mathcal{F}_t \right]. \quad (\text{A.10})$$

But since $u_{bbb} \geq 0$ and $m_t < 0$, we know $-u_{bb}(s_\tau, \hat{y}_\tau + m_t) \geq -u_{bb}(s_\tau, \hat{y}_\tau)$. Therefore:

$$Q_t \leq -E_{\hat{y}} \left[\int_t^T e^{-\rho(\tau-t)} (u_{bb}(s_\tau, \hat{y}_\tau) + 2\lambda Q_\tau) d\tau \middle| \mathcal{F}_t \right]$$

is sufficient to guarantee that (A.10) holds. Since \hat{y} was arbitrary, this is the same as (17) in part 2 of the theorem. Clearly, setting $\lambda = 0$ gives (18) in part 3.

Thus we have shown that $\Delta_t = m_t = 0$ for all t is the solution of the control problem above, under the conditions (16) and (17), or (18) when $\lambda = 0$. Clearly with $\Delta_t = m_t = 0$ the maximized objective in (A.7) is zero, so $V(\hat{y}) - q_0 \leq 0$. Thus the utility gain from lying is bounded by zero, and the contract ensures truthful revelation.

A.3. CALCULATIONS FOR THE EXAMPLES

A.3.1. Calculations for the Hidden Endowment Examples

We first verify the guess (26) of the form of the value function. From the optimality conditions (24)-(25) and the form of the guess we get:

$$\begin{aligned} s &= \frac{\log \theta}{\theta} + \frac{\log(J_q + \theta J_p)}{\theta} - y \\ &= \frac{\log \theta}{\theta} + \frac{\log(j_2 + h'(k)(k - \theta)) - \log(-q)}{\theta} - y, \\ Q &= \frac{pJ_{qp}}{J_{pp}} = -qk \left(\frac{h'(k)}{h''(k)} + k \right) \equiv -q\hat{Q}(k). \end{aligned}$$

Then substituting these into the HJB equation (19) we get:

$$\begin{aligned} j_0 &= \frac{\log \theta}{\rho\theta} - \frac{\mu_0}{\rho(\rho + \lambda)}, \\ j_1 &= -\frac{1}{\rho + \lambda}, \\ j_2 &= \frac{1}{\rho\theta}, \end{aligned}$$

while $h(k)$ satisfies the second order ODE:

$$\rho h(k) = \frac{1}{\theta} \log \left(\frac{1}{\rho\theta} + h'(k)(k - \theta) \right) + \lambda h'(k)k + \frac{\sigma^2 k^2}{2} \left(\frac{1}{\rho\theta} - \frac{h'(k)^2}{h''(k)} \right). \quad (\text{A.11})$$




FIGURE 1. The function $h(k)$ from the principal’s cost function with $\lambda = 0.05$. The dotted lines show the optimal choice $k_0^* = \rho\theta/(\rho + \lambda)$ and its corresponding value $h(k_0^*)$.

Thus we’ve verified the guess. The agent’s consumption under the contract can then be written:

$$c(q, k) = \frac{\log(1/\rho + \theta h'(k)(k - \theta))}{\theta} - \frac{\log(-q)}{\theta} \equiv -\frac{\log(-q\hat{c}(k))}{\theta}.$$

While it does not seem feasible to solve for $h(k)$ analytically, it is relatively simple to solve the ODE (A.11) numerically.¹ To find the optimal initial condition, we numerically solve for h and then pick the minimum. We have done this for a range of different parameterizations, all yielding the same result $k_0^* = \frac{\rho\theta}{\rho + \lambda}$. One numerical example is illustrated in Figure 1 which plots $h(k)$ when $\lambda = 0.05$.² The minimum clearly occurs at k_0^* , with the cost function increasing rapidly for larger k and quite slowly for smaller k . Similar results were found for a range of different parameterizations. Note as well that since $h'(k_0^*) = 0$ we have $\hat{c}(k_0^*) = \rho$ and $\hat{Q}(k_0^*) = (k_0^*)^2$. For the principal’s cost, note that at k_0^* we have:

$$h(k_0^*) = \frac{1}{\rho\theta} \log\left(\frac{1}{\rho\theta}\right) + \frac{\sigma^2\theta}{2(\rho + \lambda)^2}.$$

A.3.2. Calculations Verifying Incentive Compatibility

We now verify that the sufficient conditions (16) and (17) from Theorem 4.1 hold. Under the optimal contract we have:

$$u_{bb}(s_t, y_t) = v''(y_t + s_t) = -\theta^2 \exp(-\theta\bar{c}(q_t)) = \theta^2 \rho q_t.$$

¹It is *relatively* simple, since the equation is a second order ODE. However the equation has singularities which complicate matters somewhat. In practice solving (A.11) as an implicit ODE as worked best.

²The preference parameters are $\theta = 1$ and $\rho = 0.1$. For the endowment, we scale the drift and diffusion so that we approximate an i.i.d. $N(1, 0.15^2/2)$ process as $\lambda \rightarrow \infty$. Thus we set $\sigma = 0.15\sqrt{\lambda}$, and $\mu_0 = \lambda$.

In addition, we have:

$$Q_t = -(k_0^*)^2 q_t = -\theta^2 \left(\frac{\rho}{\rho + \lambda} \right)^2 q_t$$

Thus since for all $\lambda > 0$

$$\left(\frac{\rho}{\rho + \lambda} \right)^2 < \frac{\rho}{2\lambda}$$

we have that (16) holds with a strict inequality when $\lambda > 0$. (The condition is irrelevant when $\lambda = 0$.) Moreover:

$$\begin{aligned} -E_y \left[\int_t^\infty e^{-\rho(\tau-t)} [u''(y_\tau + s_\tau) + 2\lambda Q_\tau] d\tau \middle| \mathcal{F}_t \right] &= -\theta^2 \left(\rho - \frac{2\lambda\rho^2}{(\rho + \lambda)^2} \right) \int_t^\infty e^{-\rho(\tau-t)} E_y [q_\tau | \mathcal{F}_t] d\tau \\ &= -\theta^2 \frac{\rho^2 + \lambda^2}{(\rho + \lambda)^2} q_t \\ &\geq -\theta^2 \frac{\rho^2}{(\rho + \lambda)^2} q_t = Q_t \end{aligned}$$

where we use the fact that q_t is a martingale and the expression for Q_t above. Thus (17) holds for all $\lambda \geq 0$ (with strict inequality when $\lambda > 0$) under the optimal contract.

To directly verify incentive compatibility when $\lambda > 0$ we now re-solve the agent's problem given a contract. In particular, the contract specifies a payment $s(y_t, q_t) = -\log(-\rho q_t)/\theta - y_t$ so that under an arbitrary reporting strategy the agent's consumption is $c_t = -\log(-\rho q_t)/\theta - m_t$. Clearly the agent's problem depends on the endogenous state variable q_t , which under his information set evolves as follows:

$$\begin{aligned} dq_t &= -\frac{\sigma\rho\theta}{\rho + \lambda} q_t dW_t^y \\ &= -\frac{\rho\theta}{\rho + \lambda} q_t (\sigma dW_t + [\mu(y_t - m_t) + \Delta_t - \mu(y_t)] dt) \\ &= -\frac{\rho\theta}{\rho + \lambda} q_t (\lambda m_t + \Delta_t) dt - \frac{\sigma\rho\theta}{\rho + \lambda} q_t dW_t, \end{aligned}$$

where we've changed measure from the principal's to the agent's information set.

The agent's problem is then to maximize his utility over reporting strategies subject to this law of motion for q_t and the evolution (3) for m_t . The problem no longer depends on the level of the endowment y_t directly. Letting $V(q, m)$ denote his value function, we see that it satisfies the HJB equation:

$$\rho V = \max_{\Delta \leq 0} \left\{ -\exp(-\theta[\bar{c}(q) - m]) - V_q \frac{\rho\theta}{\rho + \lambda} q (\lambda m + \Delta) + V_m \Delta + \frac{1}{2} V_{qq} \left(\frac{\sigma\rho\theta}{\rho + \lambda} q \right)^2 \right\}.$$

Truth-telling is optimal if the following analogue of the incentive constraint (9) holds:

$$-V_q \frac{\rho\theta}{\rho + \lambda} q + V_m \geq 0.$$

It is easy to verify that the following function satisfies the HJB equation with $\Delta = 0$:

$$V(q, m) = \frac{q \exp(\theta m)(\rho + \lambda)}{\rho + \lambda + \theta \lambda m}.$$

Moreover, with this function we have:

$$-V_q \frac{\rho \theta}{\rho + \lambda} q + V_m = \frac{q \exp(\theta m) \theta^2 \lambda^2 m}{(\rho + \lambda + \theta \lambda m)^2} \geq 0$$

so that truth-telling is indeed optimal no matter what m is. In particular, we assumed that $m_0 = 0$ and since the agent never lies then $m_t = 0$ and $V(q_t, 0) = q_t$ for all t . Thus the contract does indeed provide incentives for truthful revelation.

A.3.3. Calculations for the Taste Shock Example

First we solve the HJB equation (21) in the full information case. Under our parameterization, this can be written:

$$\rho J^* = e^{-\frac{1}{\theta}y} (J_q^*)^{\frac{1}{\theta}} + J_y^* \mu_0 + J_q^* \rho q - \frac{1}{1-\theta} e^{-\frac{1}{\theta}y} (J_q^*)^{\frac{1}{\theta}} + \frac{\sigma^2}{2} \left[J_{yy}^* - \frac{(J_{yq}^*)^2}{J_{qq}^*} \right].$$

Then we guess that the solution is of the form

$$J^* = A \exp\left(\frac{1}{1-\theta}y\right) ((1-\theta)q)^{\frac{1}{1-\theta}}.$$

Carrying out the appropriate differentiation, and substituting the expression into the HJB equation, we see that there is a common $\exp\left(\frac{1}{1-\theta}y\right) ((1-\theta)q)^{\frac{1}{1-\theta}}$ factor on both sides of the equation. Collecting constants, we then have that A solves:

$$A \left(-\rho + \frac{\mu_0}{1-\theta} + \frac{\rho}{1-\theta} + \frac{\sigma^2}{2} \left[\frac{1}{(1-\theta)^2} - \frac{1}{\theta(1-\theta)^2} \right] \right) = \frac{\theta}{1-\theta} A^{\frac{1}{\theta}}.$$

Solving for A , simplifying, and using the value of μ_0 yields the expression in the text.

The solution of the hidden information case proceeds in much the same way, only now with a different γ . In particular, the HJB equation (19) can now be written:

$$\rho J = e^{-\frac{1}{\theta}y} J_q^{\frac{1}{\theta}} + J_y \mu_0 + J_q \rho q - \frac{1}{1-\theta} e^{-\frac{1}{\theta}y} J_q^{\frac{1}{\theta}} + \frac{\sigma^2}{2} [J_{qq} q^2 + J_{yy} - 2J_{yq} q].$$

Then we guess that the solution is again of the form

$$J = B \exp\left(\frac{1}{1-\theta}y\right) ((1-\theta)q)^{\frac{1}{1-\theta}}.$$

Once again there will be a common $\exp\left(\frac{1}{1-\theta}y\right) ((1-\theta)q)^{\frac{1}{1-\theta}}$ factor on both sides of the equation. Collecting constants, we then have that B solves:

$$B \left(-\rho + \frac{\mu_0}{1-\theta} + \frac{\rho}{1-\theta} + \frac{\sigma^2}{2} \left[\frac{\theta}{(1-\theta)^2} + \frac{1}{(1-\theta)^2} - \frac{2}{(1-\theta)^2} \right] \right) = \frac{\theta}{1-\theta} B^{\frac{1}{\theta}}.$$

Solving for B , simplifying, and using the value of μ_0 yields the expression in the text.

References

- Abraham, A. and N. Pavoni (2008). Efficient Allocations with Moral Hazard and Hidden Borrowing and Lending. *Review of Economic Dynamics* 11, 781–803.
- Abreu, D., D. Pearce, and E. Stacchetti (1986). Optimal Cartel Equilibria with Imperfect Monitoring. *Journal of Economic Theory* 39, 251–269.
- Abreu, D., D. Pearce, and E. Stacchetti (1990). Toward a Theory of Discounted Repeated Games with Imperfect Monitoring. *Econometrica* 58, 1041–1063.
- Atkeson, A. and R. E. Lucas (1992). On Efficient Distribution with Private Information. *Review of Economic Studies* 59, 427–453.
- Battaglini, M. (2005). Long-Term Contracting with Markovian Consumers. *American Economic Review* 95, 637–658.
- Biais, B., T. Mariotti, G. Plantin, and J.-C. Rochet (2007). Dynamic Security Design: Convergence to Continuous Time and Asset Pricing Implications. *Review of Economic Studies* 74, 345–90.
- Bismut, J. M. (1973). Conjugate Convex Functions in Optimal Stochastic Control. *Journal of Mathematical Analysis and Applications* 44, 384–404.
- Bismut, J. M. (1978). Duality Methods in the Control of Densities. *SIAM Journal on Control and Optimization* 16, 771–777.
- Breedon, D. (1979). An Intertemporal Asset Pricing Model with Stochastic Consumption and Investment Opportunities. *Journal of Financial Economics* 7, 265–296.
- Cvitanić, J., X. Wan, and J. Zhang (2009). Continuous-Time Principal-Agent Problems with Hidden Action and Lump-Sum Payment. *Applied Mathematics and Optimization* 59, 99–146.
- Cvitanić, J. and J. Zhang (2007). Optimal Compensation with Adverse Selection and Dynamic Actions. *Mathematics and Financial Economics* 1, 21–55.
- DeMarzo, P. M. and Y. Sannikov (2006). Optimal Security Design and Dynamic Capital Structure in a Continuous-Time Agency Model. *Journal of Finance* 61, 2681–2724.
- Fernandes, A. and C. Phelan (2000). A Recursive Formulation for Repeated Agency with History Dependence. *Journal of Economic Theory* 91, 223–247.
- Golosov, M., N. Kocherlakota, and A. Tsyvinski (2003). Optimal indirect and capital taxation. *Review of Economic Studies* 70, 569–587.
- Golosov, M. and A. Tsyvinski (2006). Designing Optimal Disability Insurance: A Case for Asset Testing. *Journal of Political Economy* 114, 257–279.
- Green, E. J. (1987). Lending and the Smoothing of Uninsurable Income. In E. C. Prescott and N. Wallace (Eds.), *Contractual Arrangements for Intertemporal Trade*, pp. 3–25. Minneapolis: University of Minnesota Press.
- Hausmann, U. G. (1986). *A Stochastic Maximum Principle for Optimal Control of Diffusions*. Essex, UK: Longman Scientific & Technical.
- Kapicka, M. (2006). Efficient Allocations in Dynamic Private Information Economies with Persistent Shocks: A First Order Approach. Working paper, University of California, Santa Barbara.
- Karatzas, I. and S. E. Shreve (1991). *Brownian Motion and Stochastic Calculus* (Second ed.). New York: Springer-Verlag.

- Kocherlakota, N. (2004). Figuring Out The Impact of Hidden Savings on Optimal Unemployment Insurance. *Review of Economic Dynamics* 7, 541–554.
- Kydland, F. E. and E. C. Prescott (1980). Dynamic Optimal Taxation, Rational Expectations and Optimal Control. *Journal of Economic Dynamics and Control* 2, 79–91.
- Liptser, R. S. and A. N. Shiryaev (2000). *Statistics of Random Processes* (Second ed.), Volume I. Berlin: Springer.
- Meghir, C. and L. Pistaferri (2004). Income Variance Dynamics and Heterogeneity. *Econometrica* 72, 1–32.
- Rogerson, W. P. (1985). Repeated Moral Hazard. *Econometrica* 53, 69–76.
- Sannikov, Y. (2007). Agency Problems, Screening and Increasing Credit Lines. Working paper, UC Berkeley.
- Sannikov, Y. (2008). A Continuous-Time Version of the Principal-Agent Problem. *Review of Economic Studies* 75, 957–984.
- Schattler, H. and J. Sung (1993). The First-Order Approach to the Continuous-Time Principal-Agent Problem with Exponential Utility. *Journal of Economic Theory* 61, 331–371.
- Spear, S. and S. Srivastava (1987). On Repeated Moral Hazard with Discounting. *Review of Economic Studies* 53, 599–617.
- Storesletten, K., C. I. Telmer, and A. Yaron (2004). Cyclical Dynamics in Idiosyncratic Labor Market Risk. *Journal of Political Economy* 112, 695–717.
- Tchisty, A. (2006). Security Design with Correlated Hidden Cash Flows: The Optimality of Performance Pricing. Working paper, NYU Stern School of Business.
- Thomas, J. and T. Worrall (1990). Income Fluctuation and Asymmetric Information: An Example of a Repeated Principal-Agent Problem. *Journal of Economic Theory* 51, 367–90.
- Werning, I. (2001). Optimal Unemployment Insurance with Unobservable Saving. Working paper, MIT.
- Williams, N. (2009). On Dynamic Principal-Agent Models in Continuous Time. Working paper, University of Wisconsin - Madison.
- Yong, J. and X. Y. Zhou (1999). *Stochastic Controls*. New York: Springer-Verlag.
- Zhang, Y. (2009). Dynamic Contracting with Persistent Shocks. *Journal of Economic Theory* 144, 635–675.
- Zhou, X. Y. (1996). Sufficient Conditions of Optimality for Stochastic Systems with Controllable Diffusions. *IEEE Transactions on Automatic Control* AC-41, 1176–1179.