

NBER WORKING PAPER SERIES

PERSISTENT PRIVATE INFORMATION

Noah Williams

Working Paper 13894

<http://www.nber.org/papers/w13894>

NATIONAL BUREAU OF ECONOMIC RESEARCH

1050 Massachusetts Avenue

Cambridge, MA 02138

March 2008

I thank Dilip Abreu, Narayana Kocherlakota, Chris Sleet, and Ted Temzelides for helpful comments. Financial support from the National Science Foundation is gratefully acknowledged. The views expressed herein are those of the author(s) and do not necessarily reflect the views of the National Bureau of Economic Research.

NBER working papers are circulated for discussion and comment purposes. They have not been peer-reviewed or been subject to the review by the NBER Board of Directors that accompanies official NBER publications.

© 2008 by Noah Williams. All rights reserved. Short sections of text, not to exceed two paragraphs, may be quoted without explicit permission provided that full credit, including © notice, is given to the source.

Persistent Private Information
Noah Williams
NBER Working Paper No. 13894
March 2008
JEL No. D82,D86,E21

ABSTRACT

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 income stream. The income stream is private information to the borrower and is persistent. I find that the optimal contract conditions on the agent's reported endowment as well as two additional state variables: the agent's utility and marginal utility under the contract. I show how persistence alters the nature of the contract, and consider an exponential utility example which can be solved in closed form. 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 endowment, and the distortions vanish as I approximate an i.i.d. endowment.

Noah Williams
Department of Economics
001 Fisher Hall
Princeton University
Princeton, NJ 08544-1921
and NBER
noahw@princeton.edu

Persistent Private Information

Noah Williams*

Department of Economics, Princeton University

E-mail: noahw@princeton.edu

Revised January 26, 2008

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 income stream. The income stream is private information to the borrower and is persistent. I find that the optimal contract conditions on the agent's reported endowment as well as two additional state variables: the agent's utility and marginal utility under the contract. I show how persistence alters the nature of the contract, and consider an exponential utility example which can be solved in closed form. 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 endowment, and the distortions vanish as I approximate an i.i.d. endowment.

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 model with persistent private information. In particular, I focus on a continuous time version of an insurance problem similar to 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.

By casting the model in continuous time, I am able to use powerful tools from stochastic control. In particular, I adapt the methods developed in Williams (2006) for hidden action problems to handle models with private information. I employ the

* I thank Dilip Abreu, Narayana Kocherlakota, Chris Sleet, and Ted Temzelides for helpful comments. Financial support from the National Science Foundation is gratefully acknowledged.

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

revelation principle to cast the problem in terms of providing incentives for truthful reporting. In particular, 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.² The promised marginal utility state reflects the persistence of the private information, and as this persistence vanishes this state is no longer necessary.

I first lay out a relatively general version of the model with a concave utility function where the endowment follows a (weakly) mean reverting diffusion process. At this level of generality, I obtain some partial results characterizing optimal contracts, which suggest some differences with the previous discrete time i.i.d. models in the literature. Then I turn to a class of examples with exponential preferences and a linear endowment process where the optimal contract can be determined explicitly. There I find that there are indeed some sharp differences between my results and those in the literature.

In particular, Rogerson (1985a) 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 "immiserization" 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 the nature of the shocks, but also on differences in the environments. In the discrete time version of my model, when deciding what to report in the current period given a contract, the agent trades off the current consumption he would get after making his report with the promised future utility following the report. In my continuous time formulation, the agent's endowment 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 endowment.³ Thus in effect current consumption is independent of the current report, and all that matters for the reporting choice is how future promised utility evolves under the contract. 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 which the inverse Euler equation embodies.

My examples also highlight the relationship between the persistence of the private information and the magnitude of the distortions this information causes. The largest distortions occur when the endowment is permanent, following a Brownian

²The use of promised utility as a state variable is now well known following 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. Similar marginal utility states have been used by Kydland and Prescott (1980) and in related contexts by Werning (2001), Abraham and Pavoni (2003), and Kapicka (2006), as well as my earlier paper Williams (2006).

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.

motion with drift. In this case the contract entails 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. However as I argue below, when the period length shrinks the distortions shrink as well, and thus the existing results are consistent with my findings.

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. My conversion of the reporting problem to a hidden action problem follows their analysis. 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. One particular difference is that given the risk neutrality they are able to define cash flows, their source of private information, as *increments* of a Brownian motion with constant drift.⁴ Thus private information in their case is i.i.d. However in my environment it is more natural, and more consistent with the continuous time asset pricing literature, to define the endowment 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 Tchisty (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 (2007) 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 (2007) 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 as discussed above 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

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

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. In section 6.3.3 I also discuss the implications of whether the endowment is the level or the increment of a diffusion process.

⁶See Williams (2006) for more discussion of the first-order approach to contracting in continuous time.

a discrete time model. He exploits some smoothness conditions to replace a promised utility function with its level and derivative. Assuming (but not proving) the validity of the first-order approach, he presents some numerical results with permanent private information which are broadly similar to one of my examples below.

The persistence of private information gives the contracting problem an adverse selection component as well. In a continuous time setting, Sannikov (2006) and Cvitanic and Zhang (2006) 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. Sannikov (2006) focuses on the risk neutral case and shows how the optimal contract can be implemented by an increasing credit line. Cvitanic and Zhang (2006) consider a somewhat more general model with risk averse agents, but focus on risk neutral examples. While I do not emphasize the screening component of contracts, one could apply methods similar to those I use here to extend their papers to allow the agents' types to 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 then analyzes the agent's problem of choosing his endowment report when facing a given contract, deriving necessary optimality conditions. Our two endogenous state variables arise naturally here. Section 4 characterizes truthful reporting contracts, deriving necessary conditions they must satisfy and presenting sufficient conditions for a contract to insure truthful revelation. The sufficient conditions are easy to check but are a bit stringent, so in many situations to be sure that the contract is indeed valid one must re-pose the agent's problem facing the contract. In section 5 we present a general approach to solve for optimal contracts, and discuss how consumption evolves under the contract. Unlike the discrete time case, utility need not decline over time in our environment. Then in section 6 we consider a class of exponential utility examples where the optimal contract can be determined in closed-form. This allows us to show how persistence matters for the contract, and we contrast our results with some of the others in the literature. Finally section 7 offers some brief concluding remarks.

2. THE MODEL

2.1. Overview

In this section I lay out the model, which is essentially a continuous time version of Thomas and Worrall (1990) with a continuous endowment process. The model consists of a risk averse agent who has a privately observable income stream, and he would like to borrow from a risk-neutral lender to stabilize his income. If the lender could observe the endowment stream, then he would absorb the risk and the fully stabilize the agent's consumption. However with the endowment being the private information of the agent, the principal must rely on the agent's reports of his income, and the agent would have an incentive to lie and report that his income is lower than it really is. Thus the key problem is to design an optimal lending contract which provides the agent the incentive to truthfully report his income. Relative to Thomas and Worrall

(1990) and the subsequent literature, the key innovation in my analysis is to allow for persistence in the agent's endowment process.

2.2. Basic Layout

I start by considering a finite horizon $[0, T]$, and later let $T \rightarrow \infty$. I follow the convention of using a plain letter to denote a whole path of a variable, thus for instance $b = (b_t)_{t=0}^T$. It will be convenient in what follows to work with specifications in which the agent's endowment is either stationary in levels or in logs. Thus denoting the agent's endowment as B , I will work with the transformation $b = f(B)$ which I assume is increasing, smooth, invertible, and has a smooth inverse. In my examples below, I focus on the case when $f(x) = x$, but another natural case is $f(x) = \log x$. As a simplification but slight abuse of terminology, I will refer to b as the endowment process. I suppose that it is given by a Markov diffusion process defined on a probability space with a Brownian motion W , which evolves as:

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

For simplicity, I assume that $\sigma > 0$ is a constant, and that $\mu : \mathbb{R} \rightarrow \mathbb{R}$ is twice continuously differentiable. If μ is constant, then the endowment process has i.i.d. increments, but clearly the level of the endowment is always persistent.⁷ I focus on the (weakly) stationary case, where μ is (weakly) decreasing and concave, thus $\mu' \leq 0$, $\mu'' \leq 0$. Thus with $\mu' = 0$ an innovation to the endowment has a permanent effect, while with $\mu' < 0$ the process is persistent but the effects of innovations die off.

As my leading example, I suppose that the drift of endowment process is affine:

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

With $\lambda = 0$ we thus have a Brownian motion with drift, and increments to the endowment are i.i.d. and innovations have permanent effects. On the other hand with $\lambda > 0$ we have an Ornstein-Uhlenbeck process (see Karatzas and Shreve (1991)). This is a continuous-time version of a stationary Gaussian autoregressive process, having the following properties:

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

Thus μ_0/λ gives the unconditional mean of the process, while λ governs the rate of mean reversion and hence the persistence of the process. As mentioned above, the endowment process 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 $\bar{\sigma}^2/2$.

⁷As noted above, if we were to follow DeMarzo and Sannikov (2006) b_t would be the *cumulative* endowment process, and hence i.i.d. increments would correspond to an i.i.d. endowment. However apart from the risk neutral case which they consider, it is not clear how to define utility over the increments of a diffusion process.

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 must 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 we must have $y_t \leq b_t$. Thus the deposit (or report) is a verifiable statement of at least part of the realized endowment.⁸ Note that 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. Note that 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 $\omega_t = \{W_s : t \leq t\}$.

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 $y|_t = \{y_s : s \leq t\}$. I assume that the endowment process is initialized at a publicly known value b_0 , and that the principal knows the process which the agent's endowment follows (i.e. he knows the function μ and the constant σ), but he does not observe the realizations of it. However since 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:

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

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 insure 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 punish him.

Formally, the agent's report y must be absolutely continuous with respect to his endowment b . Hence the 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 endowment 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 endowment, we must have $\Delta_t \leq 0$. Integrating this evolution, and using that $y_0 = b_0$, we see that the report of the endowment process is equal to the truth plus the cumulative lies:

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

⁸Restricting reports to be less than or equal to the actual realization facilitates analyzing the incentive constraints. Removing this restriction at the outset is bit difficult due to the linearity of agent's reporting problem. In the previous discrete time models, such a restriction would not bind.

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:

$$dy_t = [\mu(y_t - m_t) + \Delta_t]dt + \sigma dW_t, \tag{1}$$

$$dm_t = \Delta_t dt, \tag{2}$$

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 endowment 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 specification of payments from the principal to the agent conditional on the agent’s reports. We assume that 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 $S_T : C[0, T] \rightarrow \mathbb{R}$ adapted to \mathcal{Y}_T . Note that 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

Recall that our baseline process b is the transformed endowment process, and let g denote the inverse of that transformation: $B = g(b)$. As noted above, in the examples below we focus on the case $g(x) = x$, but a convenient alternative is $g(x) = \exp(x)$. Note that for an arbitrary contract, the agent’s consumption at date t is simply his endowment plus his payment from the principal:

$$c_t = g(b_t) + s_t(y) = g(y_t - m_t) + s_t(y). \tag{3}$$

As the payment is history dependent, the agent’s consumption in general depends on the whole past history of reports. This makes direct analysis of the agent’s reporting problem difficult, as standard dynamic programming methods are not applicable. However as in Williams (2006), I follow Bismut (1978) and make the problem tractable by taking the key state variable to be the density of the reporting process instead of the reporting process directly. In particular, the state variable is the relative density $\Gamma_t(y)$ for the change of measure associated with different reporting policies. As I show in Williams (2006), it evolves as:

$$d\Gamma_t = \frac{\Gamma_t}{\sigma} [\mu(y_t - m_t) + \Delta_t] dW_t^0, \tag{4}$$

with $\Gamma_0 = 1$. Here W_t^0 is a Wiener process under the measure P^0 on $C[0, T]$, and is interpretable as the agent reporting that his endowment is a martingale. Alternative reporting strategies change the distribution of outcomes. Thus for each report y we can define a new measure $P(y)$ via $\frac{dP(y)}{dP^0} = \Gamma_T(y)$, and the process

$$W_t^y = W_t^0 - \int_0^t \frac{\mu(y_s - m_s) + \Delta_s}{\sigma} ds \tag{5}$$

is a Brownian motion under $P(y)$.

The covariation between reports and the stock of lies is a key factor in the model. Thus it is also useful to take $z_t = \Gamma_t m_t$ as the relevant hidden state variable. Simple calculations show that the evolution of this scaled variable is given by:

$$dz_t = \Gamma_t \Delta_t dt + \frac{z_t}{\sigma} [\mu(y_t - m_t) + \Delta_t] dW_t^0 \quad (6)$$

with $z_0 = 0$. As I show below, this transformation of variables simplifies the analysis by encapsulating the history dependence in a convenient way.

3. THE AGENT'S REPORTING PROBLEM

In this section I derive optimality conditions for the agent facing a given contract. The agent's preferences take a standard time additive form, with a smooth concave flow utility u and terminal utility U defined over consumption and the terminal payment, respectively. 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:

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

Here the first line uses the definition of consumption from (3) and takes the expectation with respect to the measure P_y over reporting outcomes discussed above, and the second line uses the density process defined above. The agent takes the contract s as given and solves:

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

subject to (4)-(6).

Under the change of variables, the agent's reporting problem is simply a control problem with random coefficients. As in Williams (2006), I apply a stochastic maximum principle from Bismut (1973)-(1978) to derive the agent's necessary optimality conditions. 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, I associate co-state variables with the state variables, and these co-states have specified terminal conditions. However to respect the information flow and satisfy the terminal conditions, in the stochastic problem the co-states are *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:

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

We will derive the optimality conditions and the evolution of the co-states using the transformed variables (Γ, z) , but then substitute back to express everything in terms of the original variables (y, m) .

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. Instead of considering the optimality and evolution in general, 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:

$$\gamma + Qm + p \geq 0. \quad (7)$$

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:

$$\gamma + p \geq 0. \quad (8)$$

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

$$\begin{aligned} dq_t &= \left[\rho q_t - \frac{\partial \mathcal{H}(\Gamma, z)}{\partial \Gamma} \right] dt + \gamma_t \sigma dW_t^0 \\ &= [\rho q_t - u(c_t)] dt + \gamma_t \sigma dW_t^y, \end{aligned} \quad (9)$$

$$\begin{aligned} q_T &= U(y_T - S_T). \\ dp_t &= \left[\rho p_t - \frac{\partial \mathcal{H}(\Gamma, z)}{\partial z} \right] dt + Q_t \sigma dW_t^0 \\ &= [\rho p_t + \gamma_t \mu'(y_t) + u'(c_t)g'(y_t)] dt + Q_t \sigma dW_t^y, \\ p_T &= -U'(c_T)g'(y_T). \end{aligned} \quad (10)$$

Here (9) and (10) carry out the differentiation, evaluate the result under truthful revelation, and change the Brownian motion as in (5). 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 must condition on, in addition to the current report. In effect, the principal is also 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 evolution consists of (1) with $\Delta_t = m_t = 0$, (9), and (10).

To help interpret the co-state equations, consider first the co-state (9). As can be easily verified, we can write its solution as:

$$q_t = E_y \left[\int_t^T e^{-\rho(s-t)} u(c_s) ds + e^{-\rho(T-t)} U(c_T) \middle| \mathcal{F}_t \right], \quad (11)$$

so that $q_0 = V(y)$, the agent's optimal utility under the contract. Thus the agent's optimal utility process becomes a state variable for the contracting problem. The idea of conditioning on utility as a state variable is well-known 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).

However, when agent's endowment process is persistent we require an additional state variable that encapsulates the effect of this persistence. Notice that the agent's marginal utility of consumption appears in the evolution of p_t , so we refer to it as the marginal utility state. It is the marginal value in utility terms of the hidden state variable m_t evaluated at $m_t = 0$, and thus gives in essence the marginal cost of not lying. This additional state is necessary because with a persistent endowment a lie in the current period affects both the principal's expectations of future endowments and the agent's expected future consumption.

In particular, suppose that the constraint (8) binds almost everywhere and consider the affine specification above. Then (10) becomes:

$$dp_t = [(\rho + \lambda)p_t + u'(c_t)g'(y_t)] dt + Q_t \sigma dW_t^y.$$

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

$$p_t = -E_y \left[\int_t^T e^{-(\rho+\lambda)(s-t)} u'(c_s) g'(y_s) ds + e^{-(\rho+\lambda)(T-t)} U'(C_T) g'(y_T) \middle| \mathcal{F}_t \right]. \quad (12)$$

Thus p_t is the negative of the agent's optimal marginal utility (of additional endowment y_t) under the contract. As noted in the introduction, similar state variables have been used in related contexts by Werning (2001), Abraham and Pavoni (2003), Kapicka (2006), and Williams (2006). Here the persistence of the endowment effectively acts as an additional discount, since larger λ means faster mean reversion and hence shortens the effective life of the private information. In the limit as $\lambda \rightarrow \infty$ we approximate an i.i.d. endowment process, and p_t converges to an i.i.d. random variable as well. Thus in the i.i.d. case, a contract only need condition on the utility process, while persistent endowments require 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. Under full information a risk-neutral principal 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's endowment is private information, the principal induces truthful revelation by making the agent's consumption and hence utility vary with his report. In particular note that (8) implies $\gamma_t \geq -p_t \geq 0$ so that promised utility q_t has a positive loading on the endowment

shocks – promised utility increases with higher endowment realizations. With persistent private information, the principal also chooses Q_t , the loading of the promised marginal utility state on the endowment shocks. As we show below, a sufficient condition for truthful revelation is $Q_t \leq 0$, and thus in this case the promised marginal utility ($-p_t$) also has a positive loading on endowment shocks. Thus when considering whether to lie and report $\Delta_t < 0$, the agent must balance the higher utility promise he would get (consistent with higher consumption overall) with the higher marginal utility promise (consistent with less current consumption).

4. TRUTHFUL REVELATION CONTRACTS

I now characterize the types of contracts which can induce the agent to truthfully report his endowment. 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 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 insure that the agent's utility under the contract is greater than under autarky. This simply puts a lower bound on the initial utility promise:

$$q_0 \geq V^a(y_0), \tag{13}$$

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 (9)-(10). The incentive constraint is the instantaneous truth-telling condition (8) which insures 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 (2003) and Williams (2006) for example), but I show that in our setting they can be ruled out via a simple restriction on the contract, namely that $Q_t \leq 0$. However this restriction is stronger than necessary, as I show in the examples later in the paper.

Moreover, in order to show that these conditions are sufficient for truthful revelation we require some additional curvature assumptions. In particular, we need that the utility be concave in the transformed endowment y_t . That is we define:

$$\tilde{u}(y_t) = u(g(y_t) + s_t), \quad \tilde{U}(y_T) = U(g(y_T) + s_T)$$

and we assume:

$$\tilde{u}''(y_t) = u'(g(y_t) + s_t)g'(y_t)^2 + u''(g(y_t) + s_t)g''(y_t) \leq 0$$

and similarly $\tilde{U}''(y_T) \leq 0$. When we work with the level of the endowment $B = g(y) = y$ this is clearly equivalent to requiring u and U be concave. However when we work

with the log endowment, so $B = \exp(y)$, in addition to the concavity of u and U this clearly requires some further restrictions on $\{s_t\}$ and/or (u, U) . In particular, consider the constant relative risk aversion case $u(c) = c^{1-\theta}/(1-\theta)$. Then it is easy to see that the concavity of \tilde{u} requires:

$$(\theta - 1)B_t \geq s_t.$$

Thus we require at a minimum $\theta > 1$, or else this restriction would imply $s_t \leq 0$ which, apart from the trivial contract of $s_t \equiv 0$, would not satisfy the participation constraint. Put differently, this restriction limits how much consumption smoothing the principal is able to provide for the agent and still insure truthful revelation.

I now provide the key theoretical result characterizing the class of truthful revelation contracts. The theorem adapts my previous results from Williams (2006), which in turn build on Schattler and Sung (1993) and Zhou (1996). However the setting here leads to simpler sufficient conditions and 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 and the transformation g are increasing and concave, and that the drift μ is decreasing and concave. Then we have:*

1. *Any truthful revelation contract $\{s_t, \gamma_t, Q_t\}$ satisfies (i) the participation constraint (13), (ii) the promise keeping constraints (9)-(10), and (iii) the incentive constraint (8).*

2. *Any contract satisfying the conditions (i)-(iii), having $Q_t \leq 0$ for all t , and insuring that $\tilde{u}(y_t)$ and $\tilde{U}(y_t)$ are concave for all y_t is a truthful revelation contract.*

3. *If $\mu'(y) \equiv 0$, then any contract satisfying the conditions (i)-(iii), and insuring that $\tilde{u}(y_t)$ and $\tilde{U}(y_t)$ are concave for all y_t is a truthful revelation contract.*

Note that the incentive constraint (8) and $Q_t \leq 0$ together are sufficient to insure that the agent's optimality condition (7) holds, since $m_t \leq 0$. However these conditions are stronger than necessary, as the agent's optimality condition only requires (7) and thus could be consistent with $Q_t \geq 0$. Note also that when the endowment is permanent, so $\mu'(y) = 0$, we do not need this restriction on Q_t . Further, the concavity of \tilde{u} and \tilde{U} puts some additional restrictions on the class of contracts, apart from the special (but important) case where we work with the level of the endowment, in which case the concavity of preferences is natural. Nonetheless, the theorem provides some simple sufficient conditions to check whether a given contract does in fact provide incentives for truthful revelation.

However the sufficient condition may be overly stringent in some environments, as in an example below where it fails. (We also consider a permanent endowment case where the sufficient conditions do hold.) In such cases, to be sure that the contract does indeed insure truthful revelation, we 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 (2003) conduct. While re-solving the agent's problem typically would require numerical methods, as indeed would finding the contract itself, in the examples below everything can be solved analytically. Thus in practice the most useful part of the theorem is part 1, which leads to the "relaxed" principal's problem in the terminology of Rogerson (1985b). In the cases covered by parts 2-3 of the theorem the solution of the relaxed problem agrees with the original problem, but more generally this must be verified.

5. OPTIMAL CONTRACTS ON AN INFINITE HORIZON

I now turn to the principal's problem of optimal contract design. For simplicity, I consider the infinite horizon contract. 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 (9) and (10) 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 (13), (ii) the promise keeping constraints (9)-(10), and (iii) the incentive constraint (8). We focus on the relaxed problem, and do not impose the sufficient conditions for truthful revelation from Theorem 4.1, but instead check ex-post whether they are satisfied (and if not, we verify that the contract is incentive compatible). Note also that while the participation constraint places a bound on q_0 , p_0 is 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 .

As noted above, the principal's contract choice is a relatively standard stochastic control problem. Thus we can analyze it either using standard dynamic programming methods or via a maximum principle. Here I use both for different purposes. Dynamic programming is useful as a constructive method for finding a solution, either explicitly in special cases or numerically. I mostly focus on this approach. However the maximum principle is useful for obtaining some partial characterizations of the dynamics of optimal contracts, which I consider briefly below.

Beginning with dynamic programming, and 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:

$$\rho J = \min_{\{s, \gamma \geq -p, Q\}} \left\{ s + J_y \mu(y) + J_q [\rho q - u(g(y) + s)] + J_p [\rho p + \gamma \mu'(y) + u'(g(y) + s)g'(y)] + \frac{\sigma^2}{2} [J_{yy} + J_{qq} \gamma^2 + J_{pp} Q^2 + 2(J_{yq} \gamma + J_{yp} Q + J_{pq} \gamma Q)] \right\} \quad (14)$$

where we suppress the arguments of J and its derivatives. Given the solution to (14), the initial value p_0 is chosen to maximize $J(y_0, q_0, p_0)$.

5.2. Full Information

To gain some insight into the problem, consider first the full information case where the principal directly observes the agent's endowment. We verify here the typical result that a risk neutral principal absorbs all the risk, completely stabilizing the agent's consumption. In terms of the previous discussion, the principal's problem in this case is to maximize J subject only to the participation constraint. It's possible to simply 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 , which determines the cash flow he obtains, 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 .

Denoting the full information value function $J^*(y, q)$, we see that it solves an HJB equation similar to (14):

$$\rho J^* = \min_{\{s, \gamma\}} \left\{ s + J_y^* \mu(y) + J_q^* [\rho q - u(g(y) + s)] + \frac{\sigma^2}{2} [J_{yy}^* + J_{qq}^* \gamma^2 + 2J_{yq}^* \gamma] \right\} \quad (15)$$

The optimality conditions are:

$$J_q^* u'(g(y) + s) = 1, \quad \gamma = -J_{yq}^* / J_{qq}^*.$$

Denoting $\bar{c}(q) = u^{-1}(\rho q)$ as the constant consumption consistent with promised utility q , it is straightforward to verify that the solution to (15) is given by:

$$J^*(y, q) = \frac{\bar{c}(q)}{\rho} + j(y) \quad (16)$$

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

$$\rho j(y) = -g(y) + j'(y)\mu(y) + \frac{1}{2}j''(y)\sigma^2. \quad (17)$$

Below we provide an explicit solution of this ODE for a particular specification of g and μ .

Thus the optimal full information contract indeed calls for complete stabilization, setting $s_t = \bar{c}(q_t) - g(y_t)$, 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 . In addition to being a useful benchmark on its own, the full information solution provides a lower bound on the principal's cost in the hidden information case.

5.3. Hidden Information

We now return to the hidden information problem, where it is difficult to characterize the solution in the general case in much detail. However note that it is likely to be the case that the incentive constraint (8) binds, so that $\gamma = -p$. We verify this conjecture below in particular examples, but 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 the principal wants to induce as little volatility in the agent's utility as possible in order to induce truthful revelation. In particular, this incentive constraint ties the current volatility of promised utility γ to the promised marginal utility $-p$. Thus a positive innovation in the reported endowment y_t leads to an increase in the promised utility q_t (as $\gamma_t > 0$). As well, if $Q_t \leq 0$ as is sufficient for truthful revelation, the innovation leads to higher future volatility of promised utility (as p_t decreases, and hence future γ_t increases).

The γ choice is constrained, while the other optimality conditions from (14) are:

$$J_q u'(g(y) + s) - J_p u''(g(y) + s)g'(y) = 1, \quad (18)$$

$$J_{pp}Q + J_{yp} - pJ_{qp} = 0. \quad (19)$$

Relative to the full information case, the first order condition for the payment (18) has an additional term coming from the effect of consumption on the marginal utility state p_t . The condition (19) for Q_t balances the effects of the variability of p_t with the covariations of p_t with y_t and q_t .

A general strategy for solving for the optimal contract, which could form the basis of a numerical algorithm, is to start with an initial guess $J^0(y, q, p)$ for J . Then (18) can be solved to obtain $s^0(y, q, p)$, and (19) can be solved for $Q^0(y, q, p)$. Plugging these (and $\gamma = -p$) back into (14) we can solve the PDE for an updated value $J^1(y, q, p)$. This process could then be iterated until convergence. (Usually it is advisable to include a dampening factor, so that the updated value is a convex combination of J^0 and J^1 .) After finding J we can then find the optimal starting point p_0 as discussed above, and check to see if the sufficient conditions from Theorem 4.1 hold. If not, we would solve the agent's reporting problem given the contract in order to verify incentive compatibility. While such a procedure could be followed numerically, it is not feasible to obtain explicit analytic solutions except in special cases. We consider one such case below.

5.4. The Dynamics of Optimal Contracts

We can gain some additional insight into the dynamics of optimal contracts by applying a stochastic maximum principle (see Yong and Zhou (1999)). Rather than giving an full treatment, we focus on the aspects of this approach which lead to partial

characterizations of contracts.⁹ Defining $\{(z_t^y, Z_t^y), (z_t^q, Z_t^q), (z_t^p, Z_t^p)\}$ as the principal's co-states associated with y_t , q_t , and p_t , the principal's Hamiltonian can be written:

$$H(y, q, p) = z^y \mu(y) + z^q [\rho q - u(g(y) + s)] + z^p [(\rho - \mu'(y))p + u'(g(y) + s)g'(y)] + \sigma(Z^y - Z^q p + Z^p Q) - s,$$

where we've assumed that the incentive constraint (8) binds. The principal's optimal choice of payment s minimizes the Hamiltonian, giving the optimality condition:

$$-z^q u'(g(y) + s) + z^p u''(g(y) + s)g'(y) = 1. \quad (20)$$

Comparing (20) with (18), we can identify $z^q = -J_q$, $z^p = -J_p$.¹⁰ By differentiating the Hamiltonian we obtain the evolution of the principal's co-states (and thus the partial derivatives of the value function):

$$\begin{aligned} dz_t^q &= \left[\rho z_t^q - \frac{\partial H}{\partial q} \right] dt + Z_t^q dW_t \\ &= Z_t^q dW_t, \end{aligned} \quad (21)$$

$$\begin{aligned} dz_t^p &= \left[\rho z_t^p - \frac{\partial H}{\partial p} \right] dt + Z_t^p dW_t \\ &= [\mu'(y_t)z_t^p + \sigma Z_t^q] dt + Z_t^p dW_t. \end{aligned} \quad (22)$$

These conditions have some implications for the dynamics of optimal contracts. In particular, suppose that $z_t^p = 0$, so that the agent's marginal utility process implies no cost for principal. We saw above that this was the case in the full information problem, and we also argued that when the agent's endowment process is i.i.d. the principal need not condition on p_t . In either case, when $z_t^p = 0$ from the optimality condition (20) we then see that $z_t^q = -1/u'(c_t)$ and from (21) we then have that the inverse of the agent's marginal utility follows a (local) martingale. (In the full information case, $Z_t^q = 0$ and thus the martingale is a constant process.) This is an analogue of the "inverse Euler equation" which was shown by Rogerson (1985a) and Golosov, Kocherlakota, and Tsyvinski (2003) to characterize the dynamics of optimal contracts with information frictions. The principal's marginal cost of the agent's utility z_t^q is also a martingale in Thomas and Worrall (1990). Since the promised utility is the only state variable in their environment, they used this to deduce that promised utility converges to minus infinity almost surely. Zhang (2007) applied the martingale convergence theorem to the inverse of the agent's marginal utility to obtain a similar conclusion.

However in our more general environment, neither of these arguments apply. In general, $z_t^p \neq 0$ so that while the principal's marginal cost of promised utility is a martingale, this has no direct implications for the level of promised utility. Furthermore,

⁹Cvitanić, Wan, and Zhang (2007) use a stochastic maximum principle approach to characterize contracts in a moral hazard setting.

¹⁰See Yong and Zhou (1999) for general results linking dynamic programming and the maximum principle, of which this is a simple example.

when $z_t^p \neq 0$ then we cannot identify z_t^q with the inverse of the agent's marginal utility. So the inverse Euler equation and the related immiserization results need not hold in our environment. In fact, we provide examples below where they fail.

6. A CLASS OF EXPONENTIAL UTILITY EXAMPLES

In this section, we provide a class of examples which allows for explicit solutions. In particular, we suppose that the agent has exponential utility, and work directly with the level of the endowment, which we assume has affine drift. In terms of our notation above, we use the following functional forms:

$$u(c) = -\exp(-\theta c), \quad g(y) = y, \quad \mu(y) = \mu_0 - \lambda y.$$

As is well known, exponential utility with linear evolution often leads to explicit solutions, and this is once again the case here. In particular, we first suppose that the endowment is permanent ($\lambda = 0$), which leads to some simplifications in the analysis. We fully solve for the optimal contract, verify that it is indeed incentive compatible, and discuss the implications of the contract. We then turn to the more general case where the endowment is persistent ($\lambda > 0$), again provide explicit solutions, discuss how persistence affects the contract, and compare our results to those in the literature.

6.1. Full Information

As a benchmark, we first consider the full information case where full stabilization of consumption is optimal. It is easy to verify that the solution of the ODE (17) determining the principal's cost is:

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

Inverting the agent's utility function, consumption under the contract is:

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

Using the solution for j , 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. Permanent Endowment

We now turn to the case where $\lambda = 0$ so that innovations are permanent. In particular, this implies that the endowment follows a Brownian motion with drift: $y_t = y_0 + \mu_0 t + \sigma W_t^y$. As we'll see, this special case leads to some simplifications.

Moreover, by part 3 of Theorem 4.1 we know that our conditions fully characterize truthful revelation contracts. We assume throughout this section that $\mu_0 \geq \theta\sigma^2/2$, so that the drift of the endowment is greater than the agent's required risk premium. This insures that the agent's utility under autarky is bounded for any discount rate $\rho > 0$. In particular, the agent's expected utility under autarky can be found by solving:

$$\rho V^a(y) = -\exp(-\theta y) + V_y^a(y)\mu_0 + \frac{\sigma^2}{2}V_{yy}^a(y).$$

The solution of this can be easily be seen to be:

$$V^a(y) = -\frac{1}{\rho + \theta\mu_0 - \theta^2\sigma^2/2} \exp(-\theta y).$$

As we show below, the structure of the principal's cost function in the hidden information case is similar to the full information cost. Typically an additional state variable (p_t) affects the problem. However with exponential preferences utility and marginal utility are proportional, $u'(c) = -\theta u(c)$. Thus when the endowment is permanent from the representations of the co-state processes in (11) and (12) we get directly that $p_t = \theta q_t$. Thus we can dispense with p_t as a separate state variable.¹¹ Notice that this also implies:

$$Q_t = \theta\gamma_t = -\theta p_t = -\theta^2 q_t \geq 0.$$

Thus while the condition that $Q_t \leq 0$ from part 2 of Theorem 4.1 fails, it is not required here. Nonetheless, below we directly verify that the optimal contract that we find does indeed provide incentives for truthful revelation.

6.2.1. The Optimal Contract

We guess that the principal's cost function can be written as follows:

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

for some constants (j_0, j_1, j_2) . We now verify this guess. From the optimality condition (18) we get:

$$s = \frac{\log(\theta J_q)}{\theta} - y = \frac{\log(j_2\theta) - \log(-q)}{\theta} - y.$$

Substituting this into the HJB equation (14) and using our guess we get:

$$\rho(j_0 + j_1 y - j_2 \log(-q)) = \frac{\log(j_2\theta) - \log(-q)}{\theta} - y + j_1\mu_0 + \frac{j_2}{q} \left(\rho q - \frac{q}{j_2\theta} \right) + \frac{\sigma^2 j_2 \theta^2}{2}.$$

¹¹This does *not* imply that $z_t^p = 0$ as in our discussion above. Rather it implies that we can work with the concentrated value function $J(y, q) = J(y, q, \theta q)$, so that $J_q(y, q) = J_q(y, q, \theta q) + \theta J_p(y, q, \theta q)$. Thus the principal's marginal cost of promised utility is now $-(z_t^q + \theta z_t^p)$, which is typically not a martingale.

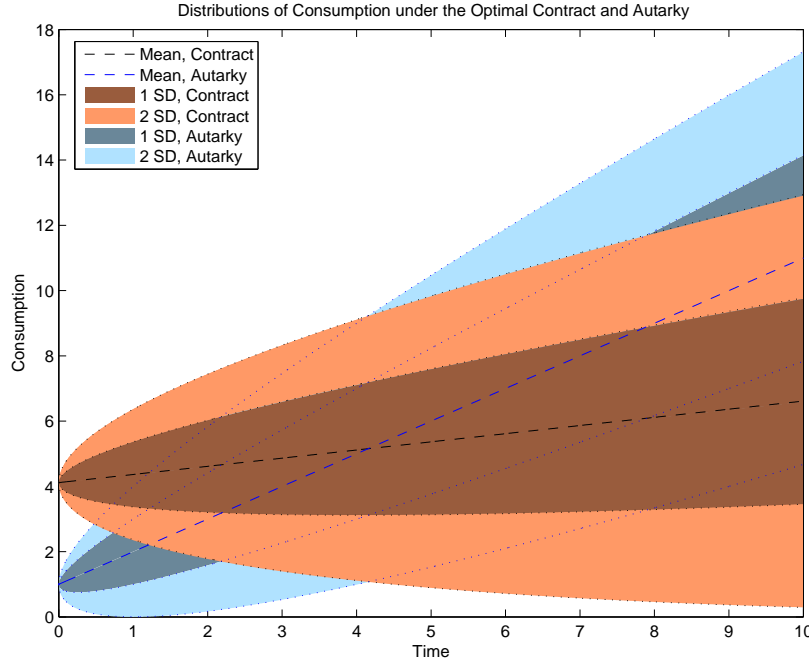


FIGURE 1. Distributions of consumption under the optimal contract and under autarky.

Solving for the constants, we then have:

$$\begin{aligned}
 j_0 &= -\frac{\log \rho}{\rho\theta} - \frac{\mu_0}{\rho^2} + \frac{\sigma^2\theta}{2\rho^2}, \\
 j_1 &= -\frac{1}{\rho}, \\
 j_2 &= \frac{1}{\rho\theta}.
 \end{aligned}$$

Thus we've verified the guess.

Note that this solution is very similar to the full information contract. In particular, we have the same optimal policy function $s = \bar{c}(q) - y$. Thus consumption as a function of promised utility is exactly the same whether the agent's endowment is private information or not. However in the hidden information case promised utility varies, while it is constant with full information. The principal's cost function is also nearly the same as the full information case, with the only change being the additional additive constant term $(\sigma^2\theta)/(2\rho^2)$. Thus 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.

The optimal contract has striking implications for the dynamics of consumption and promised utility. In particular, since $\gamma_t = -p_t = -\theta q_t$ and $c_t = \bar{c}(q_t)$ we have:

$$dq_t = -\sigma\theta q_t dW_t^y,$$

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

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

Therefore consumption tends to grow over time, having an upward trend which provides risk compensation for the increasing variability. As consumption tends to grow over time, promised utility tends toward its upper bound of zero. The distributions of consumption under the optimal contract and under autarky are shown in Figure 1.¹² There we see that the contract alters the time profile of consumption, providing more consumption up front in exchange for a smaller increase in consumption over time.

Our results stand in sharp contrast to the immiserization results of Thomas and Worrall (1990), where the agent's promised utility under the optimal contract tends toward minus infinity. This seems to be at least partly a consequence of the endowment process being permanent here as opposed to independently distributed in their setting. (In section 6.3.3 below we discuss the relationship with previous results in more detail.) 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 distribution of the endowment itself fans out, so tying promised utility (and hence consumption) to the endowment accomplishes this. But this comes at the cost of a collapse of risk sharing, as the contract insures that the agent's consumption moves one-for-one with innovations to his endowment. This is seen in Figure 1, where the consumption distribution under the contract fans out at the same rate under the contract as under autarky.

Moreover, with an i.i.d. endowment Thomas and Worrall (1990) show that it would become increasingly costly over time for the principal to allow the agent's utility to increase. In our case we do have $J_q = 1/q$ which increases to infinity as q increases to zero, but this only implies that *initial* utility promises are increasingly costly. How the principal's cost evolves over time depends on y_t and q_t , and in particular note that we have:

$$J(y_t, q_t, p_t) = J(y_0, q_0, p_0) + \frac{1}{\rho} \left(\frac{\sigma^2\theta}{2} - \mu_0 \right) t.$$

As we've assumed that the drift of the endowment μ_0 exceeds the required risk premium $\sigma^2\theta/2$, the principal's cost decreases over time, even as the agent's consumption tends to increase. Notice as well that the principal's cost evolves deterministically, as under the contract the agent now bears the risk.

6.2.2. Verifying Truthful Revelation and Implementation

¹²The figure uses the following parameters: $\rho = 0.1$, $\theta = 0.5$, $\sigma = 1$, $\mu_0 = 1$.

Although our sufficient condition from part 3 of Theorem 4.1 holds, it is useful to also directly verify that the contract is indeed consistent with truthful revelation. This helps to understand how the contract provides incentives for revelation. As discussed above, the contract ties the agent's promised utility and hence his consumption directly to his reports of the endowment. In particular, notice that there are two parts to the principal's payment s_t to the agent: a transfer from the agent due to y_t and a transfer to him due to q_t . 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.

More formally, note that under the contract but with an arbitrary reporting strategy we have:

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

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 insures 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.

This final depiction of consumption under the optimal contract makes clear that it can be implemented in a simple way. In particular, rather than depicting the transfers as $s_t = s(q_t, y_t)$, the principal can implement the optimal contract by simply providing a deterministic, time-dependent transfer:

$$s_t = \bar{c}(q_0) - y_0 + \left(\frac{\sigma^2\theta}{2} - \mu_0 \right) t.$$

In this form, the contract insures that the participation constraint is satisfied and is (weakly) consistent with truthful revelation, although the report is actually irrelevant.

6.3. Persistent Endowment

We now consider the more general case when the endowment is persistent but not necessarily permanent. The solution in this case has much of the same structure as in the permanent case, but now the contract and thus the cost depend on the ratio of the promised utility to promised marginal utility. When the endowment is permanent, this

ratio is necessarily constant. When the endowment is persistent but not permanent, this ratio may vary over time. But we show that under the optimal contract it is once again constant, albeit at a different level. We also show that the degree of risk sharing increases as the endowment becomes less persistent, and in the i.i.d. limit the contract converges to the full-information solution with complete stabilization of consumption. 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.

6.3.1. The Optimal Contract with an Arbitrary Initial Condition

We first consider the contract for an arbitrary initial condition for the marginal utility state p_0 . Later, we consider the choice of initial value, which we find makes p_t and q_t remain proportional over time. Thus first consider the principal's problem given the states (y, q, p) . Similar to above, we guess that the principal's cost function can be written as follows:

$$J(y, q, p) = j'_0 + j'_1 y - j_2 \log(-q) + h(p/q),$$

for some constants (j'_0, j'_1, j_2) and some function h . We now verify this guess. First, to simplify notation define $k = p/q$. In the permanent endowment case we have $k_t = \theta$ for all t , but when $\lambda > 0$ this ratio varies over time and becomes a state variable. From the optimality conditions (18)-(19) and the form of the guess we get:

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

Then substituting these into the HJB equation (14) 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:

$$\rho h(k) = \frac{1}{\theta} \log \left(\frac{1}{\rho \theta} + h'(k)(k - \theta) \right) + \lambda h'(k)k + \frac{\sigma^2 k^2}{2} \left(\frac{1}{\rho \theta} - \frac{h'(k)^2}{h''(k)} \right). \quad (23)$$

Thus we've verified the guess. Notice that the constant j_2 is the same as in the permanent endowment case, while j'_1 is the same as in the general full information case and reduces to the constant j_1 in the permanent endowment case when $\lambda = 0$.

The constant j'_0 is slightly different, but it insures that the solution here reduces to our previous findings when the endowment is permanent.

Unfortunately it does not seem feasible to solve for $h(k)$ analytically in general. Nonetheless, it is relatively simple to solve the ODE (23) numerically.¹³ Once we have a solution, it is easy to check whether our sufficient conditions from Theorem 4.1 hold. In particular, we can be sure that the contract is implementable as long as $Q_t \leq 0$ which here implies $\bar{Q}(k) \leq 0$. However this sufficient condition need not be met, and thus in general we need to verify whether the candidate optimal contract we find is indeed consistent with truthful revelation. We do so below after finding the optimal initial condition.

Note that the agent's consumption under the contract can then be written:

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

As in the previously analyzed cases above, the agent's consumption is independent of the current endowment y directly, and the direct dependence on the promised utility is logarithmic.¹⁴ However consumption now also depends on the ratio of the utility and marginal utility promises k , which in general will vary over time when the endowment is persistent but not permanent.

We also see that the dynamics of the co-states can be written:

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

Thus by applying Ito's lemma we have the dynamics of the ratio:

$$dk_t = [\hat{c}(k_t)(k_t - \theta) + \lambda k_t + \sigma^2 k_t^2 (k_t - \bar{Q}(k_t))] dt + \sigma k_t (k_t - \bar{Q}(k_t)) dW_t. \quad (24)$$

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.

6.3.2. 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. Of course given q_0 , which is pinned down by the participation constraint, choosing p_0 is equivalent to choosing k_0 . If the optimal choice is interior, we have $h'(k_0) = 0$. In particular, when $\lambda = 0$ and we know that $k_t = \theta$ for all t , so that it must be optimal to also start from $k_0 = \theta$ so $h'(\theta) = 0$. While the optimal choice of k_0 for $\lambda > 0$ is difficult to establish analytically, we verified numerically for a range of different λ values that the optimal initial condition is $k_0^* = \frac{\rho\theta}{\rho+\lambda}$. One numerical example is illustrated in Figure 2 which plots the function $h(k)$ when $\lambda = 0.05$.¹⁵ The

¹³It is *relatively* simple, since the equation is a second order ODE. However the equation has singularities which complicate matters somewhat. In practice solving (23) as an implicit ODE as it is written worked best.

¹⁴Obviously consumption does depend on the endowment history through q_t and k_t . But conditional on those endogenous variables, there is no further dependence on y_t .

¹⁵The 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 $\lambda \rightarrow \infty$ as discussed above. Thus we set $\sigma = 0.15\sqrt{\lambda}$, and $\mu_0 = \lambda$.




FIGURE 2. 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 $h(k_0^*)$.

minimum clearly occurs at k_0^* , with the cost function increasing rapidly for larger k and quite slowly for smaller k . The consumption rate function $\hat{c}(k)$ and the component $\bar{Q}(k)$ of the marginal utility state are shown in Figure 3. For small values of k , the consumption rate decreases only slightly while the volatility term \bar{Q} decreases quite rapidly. The converse happens for larger values of k , with the consumption rate increasing rapidly while the volatility increases only slightly. The optimal choice k_0^* is thus right on the cusp of large changes in the underlying functions.

The choice of the initial condition is significant, as it insures that k_t remains constant at k_0^* under the contract, and thus once again the promised utility and marginal utility states are proportional. In particular, by substituting $k = k_0^*$ and $h'(k_0^*) = 0$ in our expressions above, we have that the optimal consumption function is:

$$c(q, k_0^*) = -\frac{\log(-\rho q)}{\theta}$$

and thus $\hat{c}(k_0^*) = \rho$, just as we had previously. Making the same substitutions in the volatility of the marginal utility state gives:

$$\bar{Q}(k_0^*) = k_0^*.$$

Using this and the previous results in (24) we see that the ratio of the utility and marginal utility states is constant throughout the contract: $k_t = k_0^*$ for all t .

These results imply that, as in the permanent case, the agent's promised utility is a martingale under the contract. But now the agent's exposure to risk depends on the




FIGURE 3. The agent's consumption rate function $\hat{c}(k)$ and the component $\bar{Q}(k)$ of the volatility of the marginal utility state with $\lambda = 0.05$. The dotted lines show the optimal choice $k_0^* = \rho\theta/(\rho + \lambda)$ and the corresponding values $\hat{c}(k_0^*) = \rho$ and $\bar{Q}(k_0^*) = k_0^*$.

persistence of the information:

$$dq_t = -\frac{\sigma\rho\theta}{\rho + \lambda}q_t dW_t^y. \quad (25)$$

For $\lambda = 0$ we obtain our previous result in the permanent case, where there is no risk sharing. However as information becomes less persistent the amount of risk sharing increases. In particular, in the limit as we approach 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.

As in the permanent case, we can explicitly solve for consumption and the principal's cost. For the principal's cost, note that at k_0^* we have:

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

In turn, this implies:

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

so that the cost is equal to the full information cost plus a constant, as in the permanent case. The additional cost over the full information case from not observing the endowment goes to zero as the process becomes less persistent. For consumption, using exactly the same methods as above we have:

$$c_t = \bar{c}(q_0) + \frac{\sigma^2\theta\rho^2}{2(\rho + \lambda)^2}t + \frac{\sigma\rho}{\rho + \lambda}W_t^y.$$

So again consumption has a positive drift and increasing variability, but now both of these are tied to the degree of persistence.

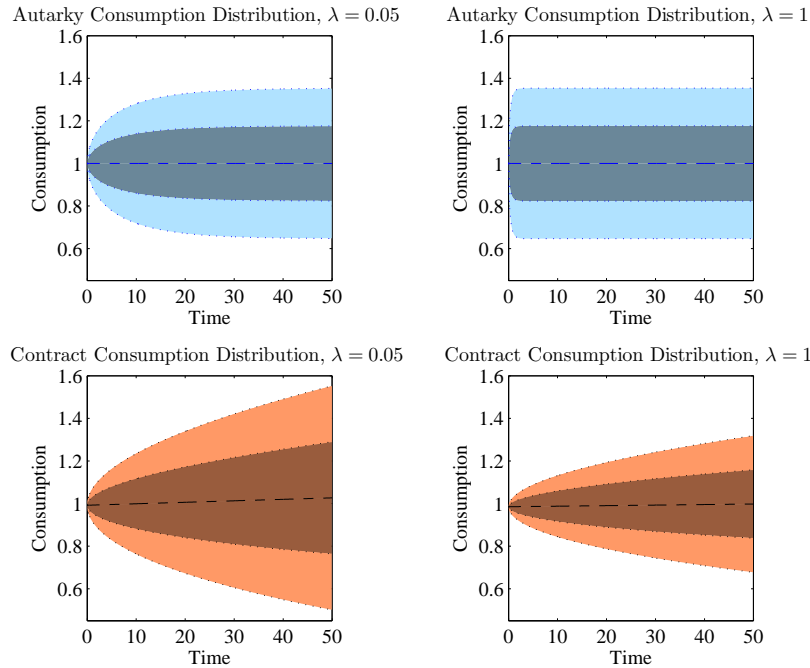


FIGURE 4. Distributions of consumption under the optimal contract and under autarky for two different values of λ .

The results are illustrated in Figure 4 which again plots the distributions of consumption under autarky and under the contract, with two different values of λ . Again we scale the drift and diffusion parameters with λ to maintain a constant unconditional mean and variance.¹⁶ The top two panels plot the distributions of the endowment, and hence the consumption distribution under autarky. The endowment is now stationary, so after being initialized at a common point (which we take to be 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 distribution under the optimal contract. Here we see that early in the contract 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 now, like the permanent endowment contract above but unlike the autarky case here, consumption is nonstationary under the contract (following a Brownian motion with drift). Thus consumption has an upward trend (which is difficult to see in the figure) and its distribution fans out over time.

6.3.3. Relationship with Previous Literature

Once again, our results are in sharp contrast to the discrete time models with i.i.d. private information. Again we see that by linking the agent's promised utility to the reported endowment shocks, the principal is able to achieve the fanning out of the

¹⁶In particular, we choose $\theta = 1$, $\rho = 0.1$, $\bar{\sigma} = 0.25$, and $\bar{\mu} = 1$.

utility distribution which is required to provide incentives, yet still have consumption increase over time. However here we see that how fast 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 insure truthful reporting.

The fact that we obtain efficiency in the limit is also due to the difference in the information structures between our continuous time formulation and the discrete time models in the literature. Note that we require the principal's payment to be a predictable function of the reported endowment process, which is a usual assumption in stochastic control and is crucial here for consumption to also be predictable. While this assumption is natural, it also means that the transfer from the principal to the agent does not respond instantaneously to the innovation in the current report. This was clear in our formulation above, as the agent's lies Δ_t affect the future transfers but not current consumption. When the endowment reverts very quickly to its mean, a current lie has little effect on the future value of the reported endowment.

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 be absolutely continuous we are able to define utility $U \circ Y = E \left[\int_0^T e^{-\rho t} u(y_t) dt \right]$. 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, except in the case of risk-neutrality as studied by DeMarzo and Sannikov (2006). While this technical difference may seem relatively minor, it insures 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 (25) 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.

These results are also implicit in Thomas and Worrall (1990). For example, 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, we can view our results as letting the period length shrink so we approach continuous time. But the consequence is the same – the efficiency gap 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.

As a final comparison, note that although the inverse Euler equation does not hold in our environment, it does indeed hold in the continuous time moral hazard model we analyzed in Williams (2006). 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 the 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.

6.3.4. Verifying Truthful Revelation

We noted above that our sufficient condition for truthful revelation fails, as we have $\bar{Q}(k_0^*) = k_0^* \geq 0$. So we now verify that the contract does indeed provide incentives for truthful revelation. While this was relatively trivial in the permanent case, we now must solve the agent's problem facing the 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:

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

where we've used the change of measure from the principal's information set to the agent's as in (5) and (A.3) in the appendix.

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 (2) 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:

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

Truth-telling is optimal if the following analogue of the incentive constraint (7) 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.

7. CONCLUSION

In this paper I have developed some relatively general methods to study optimal contracting problems with persistent private information. I have also shown how the methods lead to explicit solutions in a class of exponential utility examples. By casting the model in continuous time, I was able to use powerful tools from stochastic control. These allowed by 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, our second state variable is more novel, although it does have some precedents in the literature. These results are derived from analyzing the agent’s necessary conditions for optimality, and thus are not guaranteed to insure that the contract does provide incentives for truthful revelation. While we provide some simple sufficient conditions which guarantee truthful revelation, they are stronger than necessary as we illustrate in some examples. Thus to be sure that the contract is indeed valid, typically one must re-pose the agent’s reporting problem facing the contract. In the examples we were able to explicitly solve for both the contract and the agent’s reposed problem, but apart from special cases like these, solutions will 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. I have shown that the “inverse Euler equation” of Rogerson (1985a) and Golosov, Kocherlakota, and Tsyvinski (2003) need not hold, and the corresponding immiserization results of Thomas and Worrall (1990) may fail. In my examples the agent’s promised utility under the contract is a martingale and his consumption increases over time. The examples also highlight the close link between the persistence of the private information and the size of efficiency losses this information causes. In particular in the i.i.d. limit we obtain efficiency. This differs from the discrete time i.i.d. models in the literature, but the efficiency limit is inherent in those models as period length shrinks.

In closing, note that a wide variety of applications and extensions of the results here are worth pursuing. Although the exponential utility examples presented here are useful in understanding the role of information persistence, more general preference specifications will be crucial in more applied work seeking to quantify the effects of information frictions. With more general preferences, promised utility and promised marginal utility need not remain constant under the optimal contract. This additional source of variability may have interesting implications for consumption dynamics and asset prices.¹⁷ Another realm of application is to extend models of dynamic optimal taxation, such as Albanesi and Sleet (2006), to allow for persistent private information. As discussed above, Kapicka (2006) and Zhang (2007) have taken some steps in this direction with their related approaches. However there is much yet to be done in this area, and the tractability and generality of my approach may pay dividends.

¹⁷To analyze asset prices, one would like to embed the one-on-one contracting problem in a market setting, which is task in itself. Kocherlakota and Pistaferri (2007) take a more direct approach and analyze the shadow asset prices.

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, all of which are satisfied under our assumptions except for a linear growth condition on μ :

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

Clearly this holds under the affine specification that we mostly focus on.

Given the definition of the Hamiltonian in the text, for an arbitrary optimal, but not necessarily truthful, reporting strategy we have the evolution of the co-state variables:

$$\begin{aligned} dq_t &= \left[\rho q_t - \frac{\partial \mathcal{H}(\Gamma, z)}{\partial \Gamma} \right] dt + \gamma_t \sigma dW_t^0 \\ q_T &= U(y_T - S_T). \\ dp_t &= \left[\rho p_t - \frac{\partial \mathcal{H}(\Gamma, z)}{\partial z} \right] dt + Q_t \sigma dW_t^0 \\ p_T &= -U'(c_T)g'(y_T). \end{aligned}$$

Carrying out the differentiation, we have (suppressing arguments of functions):

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

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

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

Hence substituting these into the above evolution gives and using the change of measure (5) gives:

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

and thus we have (9)-(10) 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 u , U , and μ , we bound the utility difference by a linear approximation. Finally then we use the incentive constraint and the inequality on Q_t to show that the utility gain from deviation 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 (9)-(10), 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 (9) and integrating gives:

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

Further, using (10) and (2) along with the substitution for \tilde{p}_t gives:

$$\tilde{p}_T m_T = \int_0^T e^{-\rho t} [p_t \Delta_t + m_t \gamma_t \mu'(y_t) + m_t u'(g(y_t) + s_t) g'(y_t)] dt + \int_0^T e^{-\rho t} Q_t m_t \sigma dW_t^y. \quad (\text{A.2})$$

Also note that from (5) we have:

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

Now for an arbitrary reporting policy $\hat{y} = y - m$ we wish to compute the gain of deviating from the truthful reporting strategy y (which gives promised utility q_0):

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

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

$$\begin{aligned} E_{\hat{y}} [e^{-\rho T} [U(g(y_T - m_T) + S_T) - U(g(y_T) + S_T)]] &\leq -E_{\hat{y}} [e^{-\rho T} U'(g(y_T) + S_T) g'(y_T) m_T] \\ &= E_{\hat{y}} [\tilde{p}_T m_T] \end{aligned} \quad (\text{A.5})$$

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

$$\begin{aligned} V(\hat{y}) - q_0 &\leq E_{\hat{y}} \left[\int_0^T e^{-\rho t} [u(g(y_t - m_t) + s_t) - u(g(y_t) + s_t) + m_t u'(g(y_t) + s_t) g'(y_t)] dt \right] \\ &\quad + E_{\hat{y}} \left[\int_0^T e^{-\rho t} [p_t \Delta_t + m_t \gamma_t \mu'(y_t)] dt + \int_0^T e^{-\rho t} \sigma [\gamma_t + Q_t m_t] dW_t^y \right] \quad (\text{A.6}) \\ &\leq E_{\hat{y}} \left[\int_0^T e^{-\rho t} [p_t \Delta_t + m_t \gamma_t \mu'(y_t) + (\gamma_t + Q_t m_t) (\mu(y_t - m_t) + \Delta_t - \mu(y_t))] dt \right] \end{aligned}$$

where the second inequality uses the concavity of \tilde{u} , the change of variables from (A.3), and the fact that the stochastic integral with respect to $W^{\hat{y}}$ has expectation zero with respect to $P_{\hat{y}}$. Next the incentive constraint (8), the assumption that $Q_t \leq 0$, and the restriction that $m_t \leq 0$ insure that $(p_t + \gamma_t + Q_t m_t)\Delta_t \leq 0$. Using this in (A.6) and regrouping terms we have:

$$\begin{aligned} V(\hat{y}) - q_0 &\leq E_{\hat{y}} \left[\int_0^T e^{-\rho t} [\gamma_t (\mu(y_t - m_t) - \mu(y_t) + m_t \mu'(y_t)) + Q_t m_t (\mu(y_t - m_t) - \mu(y_t))] dt \right] \\ &\leq -E_{\hat{y}} \left[\int_0^T e^{-\rho t} Q_t m_t^2 \mu'(y_t) dt \right] \\ &\leq 0. \end{aligned}$$

Here the second inequality follows from the concavity of μ , and the final inequality follows from the assumptions that μ is decreasing and $Q_t \leq 0$. Hence the gain from deviating from truthful revelation is non-positive, so it is optimal for the agent to report his true endowment. Note that when $\mu'(y_t) \equiv 0$ then the final inequality is actually an equality for any $\{Q_t\}$ and thus we have part 3 of the theorem.

References

- Abraham, A. and N. Pavoni (2003). Efficient Allocations with Moral Hazard and Hidden Borrowing and Lending. Working paper, University College London.
- 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.
- Albanesi, S. and C. Sleet (2006). Dynamic Optimal Taxation with Private Information. *Review of Economic Studies* 73, 130.
- 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.
- Cvitanić, J., X. Wan, and J. Zhang (2007). Continuous-Time Principal-Agent Problems with Hidden Action and Lump-Sum Payment. Working paper, Caltech.

- Cvitanić, J. and J. Zhang (2006). Optimal Compensation with Adverse Selection and Dynamic Actions. Working paper, Caltech.
- 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.
- 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. (2003). Figuring Out The Impact of Hidden Savings on Optimal Unemployment Insurance. Working paper, Stanford University.
- Kocherlakota, N. and L. Pistaferri (2007). Asset Pricing Implications of Pareto Optimality with Private Information. Working paper, University of Minnesota.
- 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. (1985a). Repeated Moral Hazard. *Econometrica* 53, 69–76.
- Rogerson, W. P. (1985b). The First-Order Approach to Principal-Agent Problems. *Econometrica* 53, 1357–1367.
- Sannikov, Y. (2006). Agency Problems, Screening and Increasing Credit Lines. Working paper, UC Berkeley.
- 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. (2006). On Dynamic Principal-Agent Models in Continuous Time. Working paper, Princeton University.
- Yong, J. and X. Y. Zhou (1999). *Stochastic Controls*. New York: Springer-Verlag.
- Zhang, Y. (2007). Dynamic Contracting with Persistent Shocks. Working paper, University of Iowa.
- Zhou, X. Y. (1996). Sufficient Conditions of Optimality for Stochastic Systems with Controllable Diffusions. *IEEE Transactions on Automatic Control* AC-41, 1176–1179.