

On Dynamic Principal-Agent Problems in Continuous Time

Noah Williams*

Department of Economics, University of Wisconsin - Madison

E-mail: nmwilliams@wisc.edu

Revised September 17, 2008

I study the provision of incentives in dynamic moral hazard models with hidden actions and possibly hidden states. I characterize implementable contracts by establishing the applicability of the first-order approach to contracting. Implementable contracts are history dependent, but can be written recursively with a small number of state variables. When the agent's actions are hidden, but all states are observed, implementable contracts must take account of the agent's utility process. When the agent has access to states which the principal cannot observe, implementable contracts must also take account of the shadow value (in marginal utility terms) of the hidden states. As an application of my results, I explicitly solve a model with linear production and exponential utility, showing how allocations are distorted for incentive reasons, and how access to hidden savings further alters allocations.

1. INTRODUCTION

In many economic environments, there arise natural questions as to how to provide incentives in a dynamic setting with hidden information. Some important examples include the design of employment and insurance contracts and the choice of economic policy, specifically unemployment insurance and fiscal policy.¹ However the analysis of dynamic hidden information models rapidly becomes complex, even more so when some of the relevant state variables cannot be monitored. In this paper I develop a tractable method to analyze a class of models with hidden actions and hidden states.

I study the design of optimal contracts in a dynamic principal-agent setting. I consider situations in which the principal cannot observe the agent's actions, and the agent may have access to a privately observable state variable as well. One of the main difficulties in analyzing such models is in incorporating the agent's incentive compatibility constraints. Related issues arise in static principal-agent problems, where the main simplifying method is the first-order approach. This method replaces the incentive compatibility constraints by the first order necessary conditions from the agent's decision prob-

* I thank Dilip Abreu, Jakša Cvitanić, Narayana Kocherlakota, Nicola Pavoni, Chris Phelan, Yuliy Sannikov, Ivan Werning, and Xun Yu Zhou for helpful comments. I also especially thank Bernard Salanié (the editor) and two anonymous referees for comments that greatly improved the paper.

¹ The list of references here is long and expanding. A short list of papers related to my approach or setting include: Holmstrom and Milgrom (1987), Schattler and Sung (1993), Sung (1997), Fudenberg, Holmstrom, and Milgrom (1990), Rogerson (1985a), Hopenhayn and Nicolini (1997), and Golosov, Kocherlakota, and Tsyvinski (2003).

lem. Hence the problem becomes solvable with standard Lagrangian methods. However it has been known at least since Mirrlees (1999) (originally 1975), that this approach is not generally valid. Different conditions insuring the validity of the first-order approach in a static setting have been given by Mirrlees (1999), Rogerson (1985b), and Jewitt (1988). But there has been little previously work on the first-order approach in a dynamic setting, which has hindered the analysis of such models. Working in continuous time, I establish the validity of the first-order approach in a very general dynamic model.²

When only the agent's actions are hidden, under mild conditions the first-order approach gives a complete characterization of the class of implementable contracts, those in which the agent carries out the principal's recommended action. Even in the static setting, the conditions insuring the validity of the first order approach can be difficult to satisfy, and they become even more complex in more than one dimension. By contrast my conditions are quite natural, are easily stated in terms of the primitives of the model, and apply to multidimensional settings. I simply require that effort increase output but decrease utility, and that output and utility each be concave in effort. Thus my results here are quite general, and easier to apply than even in static settings. I discuss the reasons for this below, and show how my results relate to the static moral hazard literature.

When the agent has access to hidden states as well, I show that the first-order approach is valid under stronger concavity assumptions. As I discuss below, these assumptions are too strong for some typical formulations with hidden savings, although they may apply in other hidden state models.³ Nevertheless, even when the sufficient conditions fail, the first-order approach can be used to derive a candidate optimal contract. Then one can check whether the contract does indeed provide the proper incentives.⁴

To illustrate my methods, I study a more fully dynamic extension of Holmstrom and Milgrom (1987) in which the principal and agent have exponential preferences, and produce and save via linear technologies. In this setting I explicitly solve for the optimal contracts under full information, when the agent's effort action is hidden, and when the agent has access to hidden savings as well. I characterize how the informational frictions distort the labor/leisure margin (a "labor wedge") and the consumption/savings margin (an "intertemporal wedge").⁵ Moral hazard in production leads directly to a wedge between the agent's marginal rate of substitution and the marginal product of la-

² Independently in concurrent research, Sannikov (2008) has given a similar proof in a simpler model with only hidden actions. More recently (after earlier drafts of this paper), Westerfield (2006) analyzes a related model using different techniques, and Cvitanic, Wan, and Zhang (2006) analyze hidden action models with only terminal payments. But there are no comparable results of my generality, and none of these papers touch on the hidden action case. There are some related results on the first-order approach in other dynamic settings. For example, Phelan and Stacchetti (2001) use a first-order approach to characterize optimal Ramsey taxation.

³ See Allen (1985), Cole and Kocherlakota (2001), Werning (2002), Abraham and Pavoni (2008), Kocherlakota (2004a), and Doepke and Townsend (2006) for some alternative formulations with hidden savings. In Williams (2008), I show that the methods developed here can be used to characterize models with persistent private information.

⁴ Werning (2002) and Abraham and Pavoni (2008) follow similar procedures to numerically check whether the candidate optimal contract is implementable. In section 7 below, I do so analytically. Kocherlakota (2004a) explicitly solves a particular hidden savings model using different means and provides a discussion.

⁵ See Chari, Kehoe, and McGrattan (2007) on the role of the labor and intertemporal (or investment) wedges in fluctuations. Kocherlakota (2004b) shows how intertemporal wedges result from dynamic moral hazard problems.

bor. Moreover, the dynamic nature of the contracting problem leads to a wedge between the agent's intertemporal marginal rate of substitution and the marginal product of capital. If the principal can control the agent's consumption (or tax his savings), the contract will thus have an intertemporal wedge. However with hidden savings, the agent's consumption cannot be directly controlled by the principal. Thus there is no intertemporal wedge, but instead a larger labor wedge. Finally, I show how the contracts can be implemented by a history-dependent, performance-related payment and a tax on savings.

My analysis builds on and extends Holmstrom and Milgrom (1987) and Schattler and Sung (1993), who studied incentive provision in a continuous time model. In those papers, the principal and the agent contract over a period of time, with the contract specifying a one-time salary payment at the end of the contract period. In contrast to these *intertemporal* models, which consider dynamic processes but only a single transfer from the principal to the agent, in my fully *dynamic* model transfers occur continuously throughout the contract period.⁶ In such a dynamic setting, the agent's utility process becomes an additional state variable which the principal must respect. This form of history dependence is analogous to many related results in the literature, starting with Abreu, Pearce, and Stacchetti (1986)-(1990) and Spear and Srivastava (1987).

As noted above, the paper closest to mine is Sannikov (2008), who studies a model with hidden actions and no state variables other than the utility process. While there are similarities in some of our results, there is a difference in focus. Sannikov's simpler setup allows for more complete characterization of the optimal contract. In contrast my results cover much more general models with natural dynamics and with hidden states, but the generality implies that typically I can only provide partial characterizations. However, I do show how to use my methods to solve for optimal contracts in particular settings.

When the agent has access to hidden states, the principal must respect an additional state variable which summarizes the "shadow value" of the state in marginal utility terms. My results, which draw on some results in control theory due to Zhou (1996), are similar to the approach of Werning (2002) and Abraham and Pavoni (2008) in discrete time dynamic moral hazard problems. While I'm able to provide some conditions to prove the validity of the first-order approach *ex-ante* (which the discrete time literature has not), the conditions are rather stringent and thus in many cases I too must check implementability *ex-post* (like Werning and Abraham and Pavoni). However even in such cases, the continuous time methods here may have some computational advantages, as we can find an optimal contract by solving a partial differential equation. Section 7 provides one case where this can be done analytically.

An alternative approach is considered by Doepke and Townsend (2006), who develop a recursive approach to hidden states models by enlarging the state space. In their setup the state vector includes the agent's promised utility conditional on each value of the unobserved state. This requires discretization of the state space, and leads to a large

⁶ The recent paper by Cadenillas, Cvitanic, and Zapatero (2005) studies a general intertemporal model with full observation, and some of their results cover dynamic transfers. Detemple, Govindaraj, and Loewenstein (2005) extend the Holmstrom and Milgrom (1987) model to allow hidden states.

number of incentive constraints which need to be checked. While they provide some ways to deal with these problems, their results remain less efficient than mine, albeit requiring some less stringent assumptions.

2. OVERVIEW

In order to describe the main ideas and results of the paper, in this section I describe a simple version of a benchmark the dynamic moral hazard model. I then summarize the main results of the paper, and relate them to the literature. I defer formal detail to later sections, which also consider more general settings.

2.1. A moral hazard model: hidden actions

The model is set in continuous time, with a principal contracting with an agent over a finite horizon $[0, T]$. Output y_t is determined by the effort e_t put forth by the agent as well as Brownian motion shocks W_t . In particular:

$$dy_t = f(e_t)dt + \sigma dW_t, \quad (1)$$

where the f is twice differentiable with $f' > 0$, $f'' < 0$. Later sections allow for more complex specifications, as current output is allowed to depend on time, its own past values, current consumption or payments, and other state variables. As in the classic moral hazard literature, I focus here on the case where output is the sole state variable and it is observed by the principal, while he cannot distinguish between effort and shock realizations. Later I discuss cases when the agent observes some states that the principal cannot. The principal designs a contract that provides the agent consumption s_t and recommends an effort level \hat{e}_t . The agent seeks to maximize his expected utility:

$$E \left[\int_0^T u(s_t, e_t) dt + v(s_T) \right], \quad (2)$$

where u is twice differentiable with $u_e < 0$, $u_{ee} < 0$. Later I consider more general preferences, allowing time dependence (for discounting) as well state dependence.

In section 4, I apply a stochastic maximum principle due to Bismut (1973) to characterize the agent's optimality conditions facing a specified contract.⁷ Similar to the deterministic Pontryagin maximum principle, the stochastic maximum principle defines a Hamiltonian, expresses optimality conditions as differentials of it, and derives "co-state" or adjoint variables. Here the Hamiltonian is:

$$H(e) = \gamma f(e) + u(s, e), \quad (3)$$

⁷ The basic maximum principle is further expounded in Bismut (1978). More recent contributions are detailed in Yong and Zhou (1999). Bismut (1975) and Brock and Magill (1979) give early important applications of the stochastic maximum principle in economics.

See equation (16) below for the more general case. The agent's optimal effort choice e^* is given by the first order condition:

$$\gamma f'(e^*) = -u_e(s, e^*), \quad (4)$$

where equation (20) below provides the more general case.

From the literature starting with Abreu, Pearce, and Stacchetti (1986)-(1990) and Spear and Srivastava (1987), we know that in a dynamic moral hazard setting a contract should condition on the agent's promised utility. Here the promised utility process q_t plays the role of a co-state. In particular, the co-state follows:

$$dq_t = -u(s_t, e_t^*)dt + \gamma_t \sigma dW_t, \quad q_T = v(s_T). \quad (5)$$

This follows from the more general (17) below, and its solution can be expressed:

$$q_t = E \left[\int_t^T u(s_t, e_t^*)dt + v(s_T) \middle| \mathcal{F}_t \right].$$

Thus q_t is the agent's optimal utility process, the remaining expected utility at date t when following an optimal control e^* . Here γ_t gives the sensitivity of the agent's promised utility to the fundamental shocks, and is the key for providing incentives. Note from (4) that under our assumptions on u and f we see that $\gamma > 0$. Therefore the agent's promised utility increases with positive shocks (positive movements in W_t).

One of the main results of the paper, Proposition 5.1, establishes that the first-order approach to contracting is valid in this environment. In particular, I show that if a contract $\{s_t, \hat{e}_t\}$ is consistent with the evolution of (5) —a “promise-keeping” constraint—where $\gamma_t = -u_e(s_t, \hat{e}_t)/f'(\hat{e}_t)$ from (4)—a local incentive compatibility constraint—then the contract is implementable. That is, the agent will find it optimal to choose \hat{e}_t when facing the contract. Moreover all implementable contracts can be written in this way, and the results hold for much broader specifications than outlined in this section. The key restrictions are the natural monotonicity and concavity assumptions on preferences and production. As mentioned in the introduction, this result is far stronger than what has been shown in discrete time or even in static settings. I next try to provide some intuition as to why the results are so strong here.

2.2. Relation with previous results

2.2.1. Sufficient conditions for the first-order approach in static models

As discussed in the introduction, it is well-known that the first-order approach only holds in a static environment under rather stringent conditions, as spelled out by Mirrlees (1999), Rogerson (1985b), and Jewitt (1988). However these conditions are much easier to satisfy when an action by the agent has a (vanishingly) small effect on distribution of output.⁸ For this argument, I embed a static model in a dynamic framework. In

⁸For contemporary discussions of the first-order approach in static models, see chapter 4 of Bolton and Dewatripont (2005). The argument in this section was suggested by Bernard Salanié, the editor.

particular, let $T = 1$ and suppose that effort is constant at e_0 over the interval $[0, 1]$. Then output at date 1 is given by:

$$y_1 = f(e_0) + W_1.$$

Apart from the timing, this is a static model where W_1 is a standard normal shock with density ϕ , thus output has the density $\phi(y - f(e))$ and the normal distribution function $\Phi(y - f(e))$. The sufficient conditions for the first order approach are the monotone-likelihood ration property (MLRP) and the convexity of the distribution function condition (CDFC). In this setting the MLRP holds under the natural assumption that the production function is strictly increasing:

$$\frac{d}{dy} \left[\frac{\phi_e(y - f(e))}{\phi(y - f(e))} \right] = f'(e) > 0.$$

However the CDFC can't be assured, as we have:

$$\Phi_{ee}(y - f(e)) = \phi'(y - f(e))f'(e)^2 - \phi(y - f(e))f''(e).$$

As Jewitt (1988) notes, if $f(e) = e$ then the CDFC only holds ($\Phi_{ee} \geq 0$) if the density ϕ is increasing everywhere ($\phi' \geq 0$), which clearly is not the case for the normal density. In cases like this it is difficult to guarantee in advance that the first-order approach is valid.

Now divide the time interval $[0, 1]$ into N different sub-intervals of length Δ and suppose that the agent may choose different effort levels for each sub-interval. In this case, output is given by:

$$y_1 = \sum_{i=0}^{N-1} f(e_{\Delta i})\Delta + W_1.$$

This is a discretized (and integrated) version of equation (1) with $\sigma = 1$. For each choice e_t , $t \in \{0, \Delta, \dots, (N-1)\Delta = T\}$ we now show that the sufficient conditions hold for small enough Δ . The MLRP continues to hold for all Δ , as we have:

$$\frac{d}{dy} \left[\frac{\phi_{e_t} \left(y - \sum_{i=0}^{N-1} f(e_{\Delta i})\Delta \right)}{\phi \left(y - \sum_{i=0}^{N-1} f(e_{\Delta i})\Delta \right)} \right] = f'(e_t)\Delta > 0.$$

But now for the CDFC we have:

$$\Phi_{e_t e_t} \left(y - \sum_{i=0}^{N-1} f(e_{\Delta i})\Delta \right) = \phi' \left(y - \sum_{i=0}^{N-1} f(e_{\Delta i})\Delta \right) f'(e_t)^2 \Delta^2 - \phi \left(y - \sum_{i=0}^{N-1} f(e_{\Delta i})\Delta \right) f''(e_t)\Delta.$$

While this can't be guaranteed to be positive in general, for small enough Δ the second term dominates, and so the CDFC will hold if f is concave.

Thus by passing to the limit in which the agent takes frequent actions, with each action having a diminishingly small effect on the distribution of output, we can insure that the first-order approach is valid. The first-order conditions use only local information, and

when an action is relevant for a very short horizon such information is sufficient. This paper shows that this logic extends to much more general dynamic environments.

2.2.2. Models with two performance outcomes

While the preceding argument is informative, the argument in this paper is fairly different, and has similarities with simpler static models where output can take on only two values.⁹ In particular, a key step in my analysis uses the fact that we can express the agent's utility under the contract with a target effort policy $\{\hat{e}_t\}$ as:

$$\int_0^T u(s_t, \hat{e}_t) dt + v(s_T) = E_{\hat{e}} \left[\int_0^T u(s_t, \hat{e}_t) dt + v(s_T) \right] + V_T$$

where $E_{\hat{e}}[V_T] = 0$, and where the expectations are conditional on employing $\{\hat{e}_t\}$. This simply states that the realized utility is equal to expected utility plus a mean zero shock. Due to the special structure of the Brownian information flow, we can further express the "utility shock" V_T as a stochastic integral:

$$V_T = \int_0^T \gamma_t dW_t = \int_0^T \gamma_t (dy_t - f(\hat{e}_t) dt), \quad (6)$$

for some process $\{\gamma_t\}$.¹⁰ This representation is what allows me to depict the utility process as in (5) above, and to tie the utility sensitivity γ_t to the optimality conditions.

Similar results hold in a static model, again with effort fixed at e_0 on the interval $[0, 1]$. For this argument, it is convenient to work with separable preferences, so assume $u(s_t, e_t) = u(e_t)$. Then we can write:

$$u(\hat{e}_0) + v(s_1) = E[u(\hat{e}_0) + v(s_1(y_1)) | \hat{e}_0] + V_1(y_1),$$

where we emphasize that V_1 and s_1 are dependent on the realization of y_1 , and $E[V_1 | \hat{e}_0] = 0$. For an arbitrary output distribution we cannot say much more. But now suppose that $y_1 = 1$ with probability $f(e_0)$ and $y_1 = 0$ with probability $1 - f(e_0)$, where we of course assume $0 \leq f(e) \leq 1$.¹¹ In this case it is easy to verify that we can express the "utility shock" in a way similar to (6):

$$V_1(y_1) = \gamma(y_1 - f(\hat{e}_0)),$$

where the sensitivity term γ is the difference in the utility of consumption across output realizations: $\gamma = v(s_1(1)) - v(s_1(0))$. With this representation, we can then write the agent's problem given a contract as:

$$\max_e E[u(e) + v(s_1(y_1)) | e] = \max_e u(e) + v(s_1) - \gamma(y_1 - f(e)).$$

The first order condition for the optimal effort choice is:

$$\gamma f'(e^*) = -u'(e^*),$$

⁹The discussion of two performance outcome models follows chapter 4.1 of Bolton and Dewatripont (2005).

¹⁰See Theorem 4.15 and Problem 4.17 in Karatzas and Shreve (1991), which also state the required conditions.

¹¹Moreover, to obtain an interior solution we assume that $f(0) = 0$, $\lim_{e \rightarrow \infty} f(e) = 1$, and $f'(0) > 1$.

which is directly parallel to (4), the first-order condition in the dynamic model. In addition, this first-order condition is sufficient when $f'' < 0$ and $u'' < 0$, the same conditions as in my dynamic setting.

While the ease of analysis of the static two outcome case has been known, it is striking to see how much of it carries over to more complex dynamic environments.¹² Even though the diffusion setting of my dynamic model has a continuum of outcomes, and there may be general temporal dependence in output, the key local properties are similar to this static binary model with i.i.d. shocks.

2.3. A more complex case: hidden states

While thus far I have discussed models in which all states are observed, I also consider a more complicated environment in which the agent has access to a hidden state variable. In such situations, a recursive representation of the contract needs to condition on something more than just the promised utility. I use a first-order approach to show that, at least in some cases, it is sufficient to condition on the “shadow value” of the hidden state in marginal utility terms.

In particular, suppose that the agent’s consumption c_t cannot be observed, and that it influences the evolution of a hidden state m_t , which evolves as:

$$dm_t = b(m_t, s_t, c_t)dt.$$

The agent’s preferences are as in (2), but now with c_t replacing s_t and terminal utility $v(c_T, m_T)$. My motivating example is a model of hidden savings which has:

$$dm_t = (rm_t + s_t - c_t)dt, \tag{7}$$

where m_t is the agent’s hidden wealth, which earns a constant rate of return r , gets inflows due to the payment s_t from the principal, and has outflows for consumption. The more general specification of b also captures the persistent private information specification in Williams (2008), where m is the cumulation of the agent’s past misreports of his privately observed income. More general cases are considered in (12) below.

With the hidden state, the agent’s Hamiltonian now becomes:

$$H(m, e, c) = (\gamma + Qm)f(e) + pb(m, s, c) + u(c, e). \tag{8}$$

Optimizing (8) by choice of c yields:

$$-b_c(m, s, c^*)p = u_c(e^*, c^*). \tag{9}$$

While γ is once again the sensitivity of the utility process (5), p is a second co-state variable and Q is its sensitivity:

$$dp_t = -b_m(m_t, s_t, c_t)p_tdt + Q_t\sigma dW_t, \quad p_T = v_m(c_T, m_T). \tag{10}$$

See equation (18) below for the more general case.

¹²Similar binomial approximations have been widely used in asset pricing (see Cox, Ross, and Rubinstein (1979) for an influential example), and Hellwig and Schmidt (2002) study such approximations in a setting like mine.

The interpretation of p_t is most apparent in the hidden savings model (7). In that case $b_c = -1$ and thus (9) shows that p is the marginal utility of consumption. Moreover since $b_m = r$, the solution to (10) can be written:

$$p_t = E \left[e^{r(T-t)} v_m(c_T, m_T) \mid \mathcal{F}_t \right].$$

By the law of iterated expectations, this implies that an Euler equation holds, as for $s > t$:

$$u_c(c_t, e_t) = E \left[e^{r(s-t)} u_c(c_s, e_s) \mid \mathcal{F}_t \right].$$

The agent's optimality conditions pin down the expected evolution of the marginal utility of consumption, while the contract influences the volatility of marginal utility Q_t .

The paper's main result with hidden states, Proposition 5.2, provides necessary and sufficient conditions for the first-order approach to be valid in this environment. In particular, I show that necessary conditions for a contract to be implementable are the local incentive compatibility constraints (4) and (9), the promise-keeping constraint (5), and the "extended promise keeping" constraint (10). These conditions are also sufficient when the Hamiltonian (8) is concave in (m, e, c) . As I discuss in more detail in section 5.3 below, this concavity restriction is stringent and in fact fails in the hidden savings case. Nevertheless a constructive solution strategy when the sufficient conditions fail is to use the first-order approach to derive a candidate optimal contract, and then check ex-post whether solution is indeed implementable. Werning (2002) and Abraham and Pavoni (2008) do this numerically in their models, and I pursue this strategy in an example in section 7 below, where it can be done analytically.

2.4. Overview of the rest of the paper

The next several sections fill in the formal detail behind the results in this section. In section 3, I formally lay out the general model, and then discuss how a change of variables is useful in analyzing it. Section 4 derives optimality conditions for the agent facing a given contract, and then section 5 characterizes the set of implementable contracts. I briefly discuss the principal's choice of a contract in section 6. Then I turn to a fully solved example in section 7, where I find the optimal contracts under full information, hidden actions, and hidden savings, and illustrate the implications of the information frictions in a dynamic environment. Finally, section 8 provides some concluding remarks. An appendix provides proofs and technical conditions for all of the results in the text.

3. THE MODEL

3.1. The general model

The environment is a continuous time stochastic setting, with an underlying probability space (Ω, \mathcal{F}, P) , on which is defined an n_y -dimensional standard Brownian motion. Information is represented by a filtration $\{\mathcal{F}_t\}$, which is generated by the Brownian motion W_t (suitably augmented). I consider a finite horizon $[0, T]$ which may be arbitrarily long, and in the application in section 7 below I let $T \rightarrow \infty$.

The actions of the agent and the principal affect the evolution of an observable state $y_t \in \mathbb{R}^{n_y}$ called *output* and a hidden state $m_t \in \mathbb{R}$. I assume that the hidden state is scalar for notational simplicity only. The agent has two sets of controls, $e_t \in A \subset \mathbb{R}^{n_e}$ which affects the observable state y_t and will be referred to as *effort*, and *consumption* $c_t \in B \subset$

\mathbb{R} which affects the unobservable state m_t . The principal takes actions $s_t \in S \subset \mathbb{R}^{n_s}$, referred to as the *payment*. The evolution of the states is given by:

$$dy_t = f(t, y_t, m_t, e_t, s_t)dt + \sigma(t, y_t, s_t)dW_t \quad (11)$$

$$dm_t = b(t, y_t, m_t, c_t, s_t)dt \quad (12)$$

with y_0, m_0 given. Note that the agent does not directly affect the diffusion coefficient σ . Due to the Brownian information structure, if the agent were to control the diffusion then his action would effectively be observed. Later results make regularity assumptions on f, b , and σ . The principal observes y_t but not e_t or W_t . Moreover he knows the initial level m_0 but does not observe c_t or $\{m_t : t > 0\}$. Since there is no shock in (12), even though the principal cannot observe m_t , if he were to know agents' decisions he could deduce what it would be.¹³

Let C^{n_y} be the space of continuous functions mapping $[0, T]$ into \mathbb{R}^{n_y} . I adopt the convention of letting a bar over a variable indicate an entire time path on $[0, T]$. Note then that the time path of output $\bar{y} = \{y_t : t \in [0, T]\}$ is a (random) element of C^{n_y} , which defines the principal's observation path. I define the filtration $\{\mathcal{Y}_t\}$ to be the completion of the σ -algebra generated by y_t at each date. A *contract* specifies a set of recommended actions (\hat{e}_t, \hat{c}_t) and a corresponding payment s_t for all t as a function of the relevant history. In particular, the set of admissible contracts \mathcal{S} is the set of \mathcal{Y}_t -predictable functions $(s, \hat{e}, \hat{c}) : [0, T] \times C^{n_y} \rightarrow S \times A \times B$.¹⁴ Thus the contract specifies a payment at date t that depends on the whole past history of the observations of the state up to that date (but not on the future). The recommended actions have no direct impact on the agent, as he is free to ignore the recommendations. Since the agent optimizes facing the given contract, the set of admissible controls \mathcal{A} for the agent are those \mathcal{F}_t -predictable functions $(\bar{e}, \bar{c}) : [0, T] \times C^{n_y} \rightarrow A \times B$. I assume that the sets A and B can be written as the countable union of compact sets. A contract will be called *implementable* if given the contract (s, \hat{e}, \hat{c}) the agent chooses the recommended actions: $(\hat{e}, \hat{c}) = (\bar{e}, \bar{c})$.

3.2. A change of variables

For a given payment $s(t, \bar{y})$ the evolution of the state variables (11)-(12) can be written:

$$\begin{aligned} dy_t &= \bar{f}(t, \bar{y}, m_t, e_t)dt + \bar{\sigma}(t, \bar{y})dW_t, \\ dm_t &= \bar{b}(t, \bar{y}, m_t, c_t)dt, \end{aligned} \quad (13)$$

where the coefficients now include the contract: $\bar{f}(t, \bar{y}, m_t, e_t) = f(t, y_t, m_t, e_t, s(t, \bar{y}))$ and so on. The history dependence in the contract thus induces history dependence in the state evolution, making the coefficients of (13) depend on elements of C^{n_y} . This dependence is also inherited by the agent's preferences, as we see below. Such dependence would complicate a direct approach to the agent's problem, as a function would be a state variable.

As in Bismut (1978), I make the problem tractable by taking the key state variable to be the density of the output process rather than the output process itself. In particular,

¹³ With some additional notational complexity only, an observable shock to m_t could be included. However I exclude the more general settings with partial information where some random elements are not observed. Detemple, Govindaraj, and Loewenstein (2005) consider an intertemporal partial information model.

¹⁴ See Elliott (1982) for a definition of predictability. Any left-continuous, adapted process is predictable.

let W_t^0 be a Wiener process on C^{n_y} , which can be interpreted as the distribution of output resulting from an effort policy which makes output a martingale. Different effort choices by the agent change the distribution of output. Thus the agent's effort choice is a choice of a probability measure over output, and I take the relative density Γ_t for this change of measure as the key state variable. Details of the change of measure are given in Appendix A.1, where I show that the density evolves as:

$$d\Gamma_t = \Gamma_t \bar{\sigma}^{-1}(t, \bar{y}) \bar{f}(t, \bar{y}, m_t, e_t) dW_t^0, \quad (14)$$

with $\Gamma_0 = 1$.¹⁵ The covariation between the observable and unobservable states is also a key factor in the model. Thus is it also useful to take $x_t = \Gamma_t m_t$ as the relevant unobservable state variable. Simple calculations show that its evolution is:

$$dx_t = \Gamma_t \bar{b}(t, \bar{y}, x_t/\Gamma_t, c_t) dt + x_t \bar{\sigma}^{-1}(t, \bar{y}) \bar{f}(t, \bar{y}, x_t/\Gamma_t, e_t) dW_t^0, \quad (15)$$

with $x_0 = m_0$. By changing variables from (y, m) in (13) to (Γ, x) in (14)-(15) the state evolution is now a stochastic differential equation with random coefficients. Instead of the key states directly depending on their entire past history, the coefficients of the transformed state evolution depend on \bar{y} which is a fixed, but random, element of the probability space. This leads to substantial simplifications, as I show below.

4. THE AGENT'S PROBLEM

4.1. The agent's preferences

The agent has standard expected utility preferences defined over the states, his controls, and the principal's actions. Preferences take a standard time additive form, with flow utility u and a terminal utility v . In particular, for an arbitrary admissible control policy $(\bar{e}, \bar{c}) \in \mathcal{A}$ and a given payment $s(t, \bar{y})$, the agent's preferences are:

$$\begin{aligned} V(\bar{e}, \bar{c}) &= E_{\bar{e}} \left[\int_0^T u(t, y_t, m_t, c_t, e_t, s_t) dt + v(y_T, m_T) \right] \\ &= E_{\bar{e}} \left[\int_0^T \bar{u}(t, \bar{y}, m_t, c_t, e_t) dt + v(y_T, m_T) \right] \\ &= E \left[\int_0^T \Gamma_t \bar{u}(t, \bar{y}, m_t, c_t, e_t) dt + \Gamma_T v(y_T, m_T) \right]. \end{aligned}$$

Here the first line uses the expectation with respect to the measure $P_{\bar{e}}$ over output induced by the effort policy \bar{e} , as discussed in Appendix A.1. The second line substitutes in the contract and defines $\bar{u}(t, \bar{y}, m_t, c_t, e_t) = u(t, y_t, m_t, c_t, e_t, s(t, \bar{y}))$, and the third uses the density process defined above. The agent's problem is to solve:

$$\sup_{(\bar{e}, \bar{c}) \in \mathcal{A}} V(\bar{e}, \bar{c})$$

subject to (14)-(15), given s . Any admissible control policy (e^*, c^*) that achieves the maximum is an optimal control, and it implies an associated optimal state evolution (Γ^*, x^*) .

¹⁵ Similar ideas are employed by Elliott (1982), Schattler and Sung (1993), and Sannikov (2008) who use a similar change of measure in their martingale methods. Their approach does not apply in the hidden state case however.

4.2. The agent's optimality conditions

Under the change of variables, the agent's problem is one of control with random coefficients. 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 Hamiltonian function H as follows:

$$H(t, y, m, e, c, \gamma, p, Q) = (\gamma + Qm)\bar{f}(t, y, m, e) + p\bar{b}(t, y, m, c) + \bar{u}(t, y, m, c, e). \quad (16)$$

Here γ and Q are n_y dimensional vectors, and p is a scalar.

As in the deterministic theory, optimal controls maximize the Hamiltonian, and the evolution of the adjoint (or co-state) variables is governed by differentials of the Hamiltonian. The adjoint variables corresponding to the states (Γ_t, x_t) satisfy the following:

$$\begin{aligned} dq_t &= -[\gamma_t \bar{f}(t) - (\gamma_t + Q_t m_t) \bar{f}_m(t) + (\bar{b}(t) - \bar{b}_m(t) m_t) p_t + (\bar{u}(t) - \bar{u}_m(t) m_t)] dt + \gamma_t \sigma(t) dW_t^0 \\ &= [(\gamma_t + Q_t m_t) \bar{f}_m(t) - (\bar{b}(t) - \bar{b}_m(t) m_t) p_t - (\bar{u}(t) - \bar{u}_m(t) m_t)] dt + \gamma_t \sigma(t) dW_t^{\bar{e}}, \end{aligned} \quad (17)$$

$$q_T = v(y_T, m_T) - v_m(y_T, m_T) m_T.$$

$$\begin{aligned} dp_t &= -[\bar{b}_m(t) p_t + (\gamma_t + Q_t m_t) \bar{f}_m(t) + Q_t \bar{f}(t) + \bar{u}_m(t)] dt + Q_t \sigma(t) dW_t^0 \\ &= -[\bar{b}_m(t) p_t + (\gamma_t + Q_t m_t) \bar{f}_m(t) + \bar{u}_m(t)] dt + Q_t \sigma(t) dW_t^{\bar{e}}, \end{aligned} \quad (18)$$

$$p_T = v_m(y_T, m_T).$$

For a given (\bar{e}, \bar{c}) , I use the shorthand notation $\bar{b}(t) = \bar{b}(t, \bar{y}, m_t, c_t)$ for all the functions, and (17)-(18) use the change of measure from W_t^0 to $W_t^{\bar{e}}$. The adjoint variables follow backward stochastic differential equations (BSDEs), as they have specified terminal conditions but unknown initial values.¹⁶

In the following, I say that a process $X_t \in L^2$ if $E \int_0^T X_t^2 dt < \infty$. The first result gives the necessary conditions for optimality. As with all later propositions, required assumptions and proofs are given in Appendices A.2 and A.3, respectively.

PROPOSITION 4.1. *Suppose that Assumptions A.1, A.2, and A.3 hold, and that u also satisfies Assumption A.2. Let $(e^*, c^*, \Gamma^*, x^*)$ be an optimal control-state pair. Then there exist \mathcal{F}_t -adapted process (q_t, γ_t) and (p_t, Q_t) in L^2 (with $\gamma_t \sigma(t)$ and $Q_t \sigma(t)$ in L^2), that satisfy (17) and (18). Moreover the optimal control (e^*, c^*) satisfies for almost every $t \in [0, T]$ almost surely:*

$$H(t, \bar{y}, m_t^*, e_t^*, c_t^*, \gamma_t, p_t, Q_t) = \max_{(e, c) \in A \times B} H(t, \bar{y}, m_t^*, e, c, \gamma_t, p_t, Q_t). \quad (19)$$

Suppose in addition that A and B are convex and (u, f, b) are continuously differentiable in (e, c) . Then an optimal control (e^, c^*) must satisfy for all $(e, c) \in A \times B$, almost surely:*

$$\begin{aligned} H_e(t, \bar{y}, m_t^*, e_t^*, c_t^*, \gamma_t, p_t, Q_t)(e - e_t^*) &\leq 0 \\ H_c(t, \bar{y}, m_t^*, e_t^*, c_t^*, \gamma_t, p_t, Q_t)(c - c_t^*) &\leq 0 \end{aligned} \quad (20)$$

¹⁶ See El Karoui, Peng, and Quenez (1997) for an overview of BSDEs in finance. In particular, these BSDEs depend on forward SDEs, as described in Ma and Yong (1999).

I stress that these are only necessary conditions for problem. As discussed above, first order conditions such as (20) may not be sufficient to characterize an agent's incentive constraints, and so the set of implementable contracts may be smaller than that characterized by the first order conditions alone. However, I establish the validity of my first-order approach in the next section.

5. IMPLEMENTABILITY OF CONTRACTS

Now I characterize the class of controls that can be implemented by the principal by appropriately tailoring the contract. I focus first on settings when there are no hidden states, which is simpler and where my results are stronger. Then I add hidden states.

5.1. Hidden actions

With no hidden states, we can dispense with m and the separate control c for the agent. Thus a contract specifies a payment and a recommended action (s, \hat{e}) . I focus on interior target effort policies, and build in the incentive constraints via the first order condition (20) which thus reduces to an equality at \hat{e} . Here (20), the generalization of (4), leads to a representation of the target volatility process $\hat{\gamma}_t$:

$$\begin{aligned}\hat{\gamma}_t \equiv \hat{\gamma}(t, \bar{y}, \hat{e}_t, s_t) &= -\bar{u}_e(t, \bar{y}, \hat{e}_t) \bar{f}_e^{-1}(t, \bar{y}, \hat{e}_t) \\ &= -u_e(t, y_t, \hat{e}_t, s(t, \bar{y})) f_e^{-1}(t, y_t, \hat{e}_t, s(t, \bar{y}))\end{aligned}\quad (21)$$

Here the first line introduces notation for the target volatility function and then applies (20), assuming that \bar{f}_e is invertible, while the second uses the definitions of \bar{f} and \bar{u} .

If the agent were to carry out the recommended actions, he would obtain the expected utility $V(\hat{e}, \hat{c})$. I assume that the agent has an outside reservation utility level V_0 , and thus the contract must satisfy the *participation constraint* $V(\hat{e}, \hat{c}) \geq V_0$. Further, I say that a contract satisfies *promise-keeping* if (s, \hat{e}) imply a solution \hat{q} of the BSDE (5) with volatility process $\hat{\gamma}_t$ given by (21). This simply ensures that the policy is consistent with the utility evolution. Such contracts lead to a representation as in (6) above:

$$v(y_T) = \hat{q}_T = \hat{q}_0 - \int_0^T u(t, y_t, \hat{e}_t, s_t) dt + \int_0^T \hat{\gamma}_t \sigma(t, y_t, s_t) dW_t^{\hat{e}}, \quad (22)$$

for some $\hat{q}_0 \geq V_0$, which builds in the participation constraint.

Similarly, with the specified volatility $\hat{\gamma}_t$ the Hamiltonian from (16) can be represented explicitly in terms of the target effort \hat{e} :

$$H^*(t, e) = \hat{\gamma}(t, y_t, \hat{e}_t, s_t) f(t, y_t, e, s_t) + u(t, y_t, e, s_t). \quad (23)$$

I say a contract (s, \hat{e}) is *locally incentive compatible* if for almost every $(t, y_t) \in [0, T] \times \mathbb{R}^{n_y}$ the following holds almost surely:

$$H^*(t, \hat{e}_t) = \max_{e \in A} H^*(t, e). \quad (24)$$

This is a direct analogue of the maximum condition (19) with the given volatility process $\hat{\gamma}$. It states that, with the specified adjoint process, the target control satisfies the

agent's optimality conditions. With these definitions, I now have my main result, which establishes the validity of the first-order approach in this case.

PROPOSITION 5.1. *In addition to the assumptions of Proposition 4.1, suppose that $f_e(t, y, e, s)$ is invertible for every $(t, y, e, s) \in [0, T] \times \mathbb{R}^n \times A \times S$, so that $\hat{\gamma}_t$ in (21) is well-defined. Then a contract $(s, \hat{e}) \in \mathcal{S}$ is implementable in the hidden action case if and only if it: (i) satisfies the participation constraint, (ii) satisfies promise-keeping, and (iii) is locally incentive compatible.*

I now give some sufficient conditions which simplify the matter even further by guaranteeing that the local incentive compatibility condition holds. Natural assumptions are that u and f are all concave in e and that u_e and f_e have opposite signs, as increased effort lowers utility but increases output. These assumptions, which I state explicitly as Assumptions A.4 in Appendix A.2, imply that the target adjoint process $\hat{\gamma}$ from (21) is positive. This allows me to streamline Proposition 5.1.

COROLLARY 5.1. *In addition to the conditions of Proposition 5.1, suppose that Assumptions A.4 hold and the function H^* from (23) has a stationary point in A for almost every (t, y_t) . Then a contract $(s, \hat{e}) \in \mathcal{S}$ is implementable in the hidden action case if and only if it satisfies: (i) the participation constraint and (ii) promise-keeping.*

Under these conditions, implementable contracts are those which condition on the target utility process, which in turn builds in incentive compatibility through $\hat{\gamma}_t$ and participation through \hat{q}_0 . Moreover, the key conditions are the natural curvature restrictions in Assumptions A.4, and thus my results here are quite general.

5.2. Hidden states

When there are hidden states, the principal must forecast them. The initial m_0 is observed, but from the initial period onward the principal constructs the target as in (12):

$$d\hat{m}_t = b(t, \hat{m}_t, y_t, \hat{c}_t, s_t)dt. \quad (25)$$

Now the principal cannot directly distinguish whether m deviates from \hat{m} .

Proposition 4.1 above implies that the agent's optimality conditions in this case include two adjoint processes, and so a contract s must respect both of these. Thus a contract satisfies the *extended promise-keeping* constraints if (s, \hat{e}, \hat{c}) and the implied target hidden state \hat{m} imply solutions $(\hat{q}, \hat{\gamma})$ and (\hat{p}, \hat{Q}) of the adjoint equations (17)-(18). These constraints ensure that the contract is consistent with the agent's utility process, as well as his shadow value of the hidden states. Thus I extend my representation of the Hamiltonian (16) with these particular adjoint processes as in (23):

$$H^{**}(t, m, e, c) = (\hat{\gamma}_t + \hat{Q}_t m) f(t, y_t, m, e, s_t) + \hat{p}_t b(t, y_t, m, c, s_t) + u(t, y_t, m, c, e, s_t). \quad (26)$$

Then parallel to (22) above, a contract which satisfies the extended promise-keeping constraints has the representation:

$$\begin{aligned} v(y_T, \hat{m}_T) &= \hat{q}_T + \hat{p}_T m_T = \hat{q}_0 + \hat{p}_0 m_0 + \int_0^T dq_t + \int_0^T d(pm)_t \\ &= \hat{q}_0 + \hat{p}_0 m_0 - \int_0^T [u(t, y_t, \hat{m}_t, \hat{e}_t, \hat{c}_t, s_t)] dt + \int_0^T [\hat{\gamma}_t + \hat{Q}_t \hat{m}_t] \sigma(t, y_t, s_t) dW_t^{\hat{e}}. \end{aligned} \quad (27)$$

So now we require $V(\hat{e}, \hat{c}) = \hat{q}_0 + \hat{p}_0 m_0 \geq V_0$ for participation.

For almost every $(t, y_t) \in [0, T] \times \mathbb{R}^{n_y}$, the *local incentive compatibility* constraint now requires:

$$H^{**}(t, \hat{m}_t, \hat{e}_t, \hat{c}_t) = \max_{(e, c) \in A \times B} H^{**}(t, \hat{m}_t, e, c), \quad (28)$$

almost surely. This states that given the specified adjoint processes, if the agent has the target level of the hidden state then the target control satisfies his optimality conditions. However I also must rule out cases where the agent would accumulate a different amount of the hidden state and choose different actions. These cases are ruled out by the assumption that the Hamiltonian H^{**} is concave in (m, e, c) . This is analogous to the condition which insures the sufficiency of the maximum principle in Zhou (1996). Unfortunately, this concavity assumption is somewhat high-level, and is too strong for many applications as mentioned in section 2. I discuss this in more detail in section 5.3 below.

The main result of this section characterizes implementable contracts in the hidden state case. The necessary conditions are those from the first-order approach, which become sufficient under the concavity assumption.

PROPOSITION 5.2. *Under the assumptions of Proposition 5.1, an implementable contract $(s, \hat{e}, \hat{c}) \in \mathcal{S}$ in the hidden state case with target wealth \hat{m} satisfies (i) the participation constraint, (ii) extended promise-keeping, and (iii) is locally incentive compatible. If in addition, the Hamiltonian function H^{**} from (26) is concave in (m, e, c) and the terminal utility function v is concave in m , then any admissible contract satisfying the conditions (i)-(iii) is implementable.*

As above, simple sufficient conditions imply that the local incentive compatibility condition holds. In addition to the previous assumptions, it is natural to assume that u_c and b_c have opposite signs, as for example consumption increases utility but lowers wealth. Assumptions A.5 state this, along with requiring that $u(c, e)$ be separable.

COROLLARY 5.2. *In addition to the conditions of Proposition 5.2 (including the concavity restrictions), suppose that Assumptions A.4 and A.5 hold and the function H^{**} from (26) has a stationary point in $A \times B$ for almost every (t, y_t) . Then a contract $(s, \hat{e}, \hat{c}) \in \mathcal{S}$ which satisfies (i) the participation constraint and (ii) extended promise-keeping is implementable.*

5.3. The sufficient conditions in the hidden state case

Many applications lack the joint concavity in the hidden state and effort necessary for H^{**} to be concave, and thus for implementability to be guaranteed in the hidden state case. Assuming the necessary differentiability, it is easy to see that the Hessian of the H^{**}

function in (26) with respect to (m, e) is given by:

$$\begin{bmatrix} H_{mm}^{**} & H_{me}^{**} \\ H_{em}^{**} & H_{ee}^{**} \end{bmatrix} = \begin{bmatrix} \hat{p}_t b_{mm}(t) + u_{mm}(t) & \hat{Q}_t f_e(t) + u_{me}(t) \\ \hat{Q}_t f_e(t) + u_{me}(t) & (\hat{\gamma}_t + \hat{Q}_t m_t) f_{ee}(t) + u_{ee}(t) \end{bmatrix},$$

where I assume for simplicity $f_{mm} = f_{em} = 0$. In a standard hidden savings model like (7), b is linear and u is independent of m . Hence the upper left element of the matrix is zero. It is clear that in this case this matrix typically fails to be negative semidefinite.¹⁷

The only possibility would be $\hat{Q}_t = 0$, in which case the savings decision and effort decision would be completely uncoupled. But when the payment to the agent depends on the realized output, as it of course typically will in order to provide incentives, we have $\hat{Q}_t \neq 0$. At a minimum, the sufficient conditions require that at least one of the drift or diffusion of the hidden state or the preferences be strictly concave in m . Moreover the curvature must be enough to make the determinant of the Hessian matrix non-negative.

While the concavity of H^{**} is restrictive, it is stronger than necessary. In particular, I show in Appendix A.3 that the following condition is sufficient:

$$E_{\bar{e}} \left[\int_0^T \left(\hat{\gamma}_t [f(e_t, m_t) - f(\hat{e}_t, \hat{m}_t) - f_m(\hat{e}_t, \hat{m}_t) \Delta_t] + \hat{Q}_t m_t [f(e_t, m_t) - f(\hat{e}_t, \hat{m}_t)] \right. \right. \quad (29) \\ \left. \left. - \hat{Q}_t \hat{m}_t [f(e_t, \hat{m}_t) - f(\hat{e}_t, \hat{m}_t) + f_m(\hat{e}_t, \hat{m}_t) \Delta_t] + \hat{p}_t [b(m_t, c_t) - b(\hat{m}_t, c_t) - b_m(\hat{m}_t, \hat{c}_t) \Delta_t] \right) dt \right] \leq 0,$$

where I suppress the arguments of f and b which are inessential here. This expression can be used to insure the validity of the first order approach in certain special cases. For example, in the permanent shock model in Williams (2008), f is a constant, $b(m, c) = c$, and $\hat{m} = 0$. Thus the left side of (29) is identically zero, and so the sufficient conditions hold. In addition, in the hidden savings case of (7) the inequality (29) reduces to:

$$E_{\bar{e}} \left[\int_0^T \hat{Q}_t m_t (f(e_t) - f(\hat{e}_t)) dt \right] \leq 0. \quad (30)$$

This condition cannot be guaranteed in general, as it depends on the relationship between wealth and effort under the contract. However, I show in section 7 that with exponential utility the agent's optimal choice of e_t is independent of m_t under the optimal contract. Thus $f(e_t) = f(\hat{e}_t)$ so (30) holds, and the contract is implementable.

Although my results cover certain applications of interest, the sufficient conditions are rather stringent. But even when my sufficient conditions fail, the methods can be used to derive a candidate optimal contract. Then one can check ex-post whether the contract is in fact implementable, as I do for my example in section 7 below.

6. THE PRINCIPAL'S PROBLEM

I now turn briefly to the problem of the principal, who seeks to design an optimal contract. Since there is little I can establish at this level of generality, I simply set up and discuss the principal's (relaxed) problem here. In designing an optimal contract,

¹⁷ A similar problem in a related setting appears in Kocherlakota (2004a).

the principal chooses a payment and target controls for the agent to maximize his own expected utility, subject to the constraint that the contract be implementable. For a given contract $(\bar{s}, \bar{e}, \bar{c})$ in \mathcal{S} the principal has expected utility preferences given by:

$$J(\bar{s}, \bar{e}, \bar{c}) = E \left[\int_0^T U(t, y_t, m_t, c_t, e_t, s_t) dt + L(y_T, m_T) \right].$$

This specification is quite general, which matters little since we simply pose the problem. Assuming that Proposition 5.2 holds, the principal's problem is to solve:

$$\sup_{(\bar{s}, \bar{e}, \bar{c}) \in \mathcal{S}} J(\bar{s}, \bar{e}, \bar{c})$$

subject to the state evolution (11)-(12), the adjoint evolution (17)-(18), and the agent's optimality conditions (20). Here (y_0, m_0) are given, and (q_0, p_0) are to be determined subject to the constraints:

$$q_0 + p_0 m_0 = V(\bar{e}, \bar{c}, \bar{s}) \geq V_0, \quad q_T = v(y_T, m_T) - v_m(y_T, m_T) m_T, \quad p_T = v_m(y_T, m_T).$$

In general, this is a difficult optimization problem, which typically must be solved numerically. In the next section I analyze an example which can be solved explicitly.

7. A FULLY SOLVED EXAMPLE

In this section I study a model with exponential preferences and linear evolution that is explicitly solvable. This allows me to fully describe the optimal contract and its implementation, as well as to characterize the distortions caused by the informational frictions.

7.1. An example

I now consider a model in which a principal hires an agent to manage a risky project, with the agent's effort choice affecting the expected return on the project. The setup is a more fully dynamic version of Holmstrom and Milgrom (1987). In their environment consumption and payments occur only at the end of the period and output is i.i.d., while I include intermediate consumption (by both the principal and agent) and allow for persistence in the underlying processes. In addition, mine is a closed system, where effort adds to the stock of assets but consumption is drawn from it.¹⁸ I also consider an extension with hidden savings, where the agent can save in a risk-free asset with a constant rate of return. In this case, my results are related to the discrete time model of Fudenberg, Holmstrom, and Milgrom (1990), who show that the hidden savings problem is greatly simplified with exponential preferences. Much of the complications of hidden savings comes through the interaction of wealth effects and incentive constraints, which exponential utility does away with. However my results are not quite as simple as theirs, as

¹⁸The preferences here also differ from Holmstrom and Milgrom (1987), as they consider time multiplicatively separable preferences while I use time additively separable ones.

in my model the principal is risk averse and the production technology is persistent.¹⁹ Nonetheless, hidden savings affects the contract in a fairly simple way.

I assume that the principal and agent have identical exponential utility preferences over consumption, while the agent has quadratic financial costs of effort:

$$u(t, c, e) = -\exp\left(-\rho t - \lambda\left(c - \frac{e^2}{2}\right)\right), \quad U(t, d) = -\exp(-\rho t - \lambda d).$$

Thus both discount at the rate ρ , and d is the principal's consumption, interpreted as a dividend payment. For simplicity, I consider an infinite horizon version of the model, thus letting $T \rightarrow \infty$ in my results above. The evolution of the asset stock is linear with additive noise and the agent may have access to hidden savings as in (7):

$$\begin{aligned} dy_t &= (ry_t + Be_t - s_t - d_t) dt + \sigma dW_t \\ dm_t &= (rm_t + s_t - c_t) dt. \end{aligned} \tag{31}$$

Here s_t is the principal's payment to the agent, r is the expected return on assets in the absence of effort and also the return on wealth, and B represents the productivity of effort. In this case the agent's saving is redundant, as all that matters are total assets $y_t + m_t$. Without loss of generality I can have the principal do all the saving, and thus have the target $\hat{m}_t \equiv 0$ and so $\hat{c}_t = s_t$.²⁰

Since I consider an infinite horizon problem, it is easiest to work with the agent's *discounted* utility process (still denoted q_t), which follows:

$$dq_t = \left[\rho q_t + \exp\left(-\lambda\left(c_t - \frac{e_t^2}{2}\right)\right) \right] dt + \gamma_t \sigma dW_t.$$

I now successively solve for the optimal contract when the principal has full information, when the agent's effort choice is hidden, and when the agent's savings are hidden as well.

7.2. Full information

Although not considered above, the full information case can be analyzed using my methods as well. I let $J(y, q)$ denote the principal's value function, where the agent's promised utility only matters because of the participation constraint. The principal can freely choose the volatility term γ , as he need not provide incentives. The principal's Hamilton-Jacobi-Bellman (HJB) equation can thus be written:

$$\begin{aligned} \rho J(y, q) = \max_{c, d, e, \gamma} \left\{ -\exp(-\lambda d) + J_y(y, q) [ry + Be - c - d] + J_q(y, q) [\rho q + \exp(-\lambda(c - e^2/2))] \right. \\ \left. + \frac{1}{2} J_{yy}(y, q) \sigma^2 + J_{yq}(y, q) \gamma \sigma^2 + \frac{1}{2} J_{qq}(y, q) \gamma^2 \sigma^2 \right\} \end{aligned}$$

¹⁹In their model Fudenberg, Holmstrom, and Milgrom (1990) show that there are no gains to long term contracting, and that an optimal contract is completely independent of history. The first result relies on the risk neutrality of the principal, while the second relies on technology being history independent as well. Neither condition holds in my model, and I find that the optimal contract is history dependent and that hidden savings alter the contract.

²⁰This relies on risk being additive. Otherwise varying m may affect the risk of the total asset stock $y + m$, and the principal would face a portfolio problem.

The first order conditions for (d, c, e, γ) are then:

$$\begin{aligned}\lambda \exp(-\lambda d) &= J_y \\ \lambda \exp(-\lambda(c - e^2/2)) &= -J_y/J_q \\ \lambda e \exp(-\lambda(c - e^2/2)) &= -BJ_y/J_q \\ \gamma &= -J_{yq}/J_{qq}\end{aligned}$$

Taking ratios of the conditions for (c, e) gives $e = B$.

Due to the exponential preferences and linear of the evolution, it is easy to verify that the value function is the following:

$$J(y, q) = \frac{J_0}{q} \exp(-r\lambda y),$$

where the constant J_0 is given by:

$$J_0 = \frac{1}{r^2} \exp\left(\frac{2(r - \rho)}{r} - \frac{\lambda B^2}{2} + \frac{\sigma^2 \lambda^2 r}{4}\right).$$

The optimal policies are thus:

$$\begin{aligned}e^{fi} &= B, \quad \gamma^{fi}(q) = -\frac{r\lambda q}{2}, \\ c^{fi}(q) &= \frac{B^2}{2} - \frac{\log r}{\lambda} - \frac{\log(-q)}{\lambda}, \\ d^{fi}(y, q) &= -\frac{\log(J_0 r)}{\lambda} + \frac{\log(-q)}{\lambda} + ry.\end{aligned}$$

The agent's optimal effort is constant and his consumption does not depend directly on output, but instead is linear in the log of the utility process (which however is a function of output). The principal's consumption is linear in the log of the agent's utility process and also linear in current output. These policies imply that the state variables evolve as follows:

$$\begin{aligned}dy_t &= \left[\frac{2(r - \rho)}{r\lambda} + \frac{\sigma^2 \lambda r}{4} \right] dt + \sigma dW_t \\ dq_t &= (\rho - r)q_t dt - \frac{\sigma \lambda r}{2} q_t dW_t.\end{aligned}$$

Thus output follows an arithmetic Brownian motion with constant drift, while the utility process follows a geometric Brownian motion. The expected growth rate of utility is constant and equal to the difference between the subjective discount rate ρ and the rate of return parameter r .

7.3. The hidden action case

I now turn to the case where the principal cannot observe the agent's effort e_t . Note that the sufficient conditions from Corollary 5.1 above are satisfied. Therefore the agent's

effort level is determined by his first order condition (20), which here is:

$$\gamma B = \lambda e \exp(-\lambda(c - e^2/2)). \quad (32)$$

The principal must now choose contracts which are consistent with this local incentive compatibility condition. The principal's HJB equation now becomes:

$$\begin{aligned} \rho J(y, q) = \max_{c, d, e} \{ & -\exp(-\lambda d) + J_y(y, q) [ry + Be - c - d] + J_q(y, q) [\rho q + \exp(-\lambda(c - e^2/2))] \\ & + \frac{1}{2} J_{yy}(y, q) \sigma^2 + J_{yq}(y, q) \gamma(c, e) \sigma^2 + \frac{1}{2} J_{qq}(y, q) \gamma(c, e)^2 \sigma^2 \} \end{aligned}$$

where I substitute $\gamma = \gamma(c, e)$ using (32). The first order conditions for (d, c, e) are then:

$$\begin{aligned} \lambda \exp(-\lambda d) &= J_y, \\ -J_y - J_q \lambda \exp(-\lambda(c - e^2/2)) - J_{yq} \sigma^2 \lambda \gamma(c, e) - J_{qq} \sigma^2 \lambda \gamma(c, e)^2 &= 0, \\ J_y B + J_q \lambda e \exp(-\lambda(c - 1/2 e^2)) + J_{yq} \sigma^2 \frac{1 + \lambda e^2}{e} \gamma(c, e) + J_{qq} \sigma^2 \frac{1 + \lambda e^2}{e} \gamma(c, e)^2 &= 0. \end{aligned}$$

A special feature of this example is that the value function and the optimal policies take the same form as the full information case, albeit with different key constants. In particular, the value function is of the same form as above,

$$J(y, q) = \frac{J_1}{q} \exp(-r\lambda y)$$

for some constant J_1 . The optimal policies are thus:

$$\begin{aligned} e^{ha} &= e^*, \quad \gamma^{ha}(q) = -\frac{\lambda e^* k q}{B}, \\ c^{ha}(q) &= \frac{(e^*)^2}{2} - \frac{\log k}{\lambda} - \frac{\log(-q)}{\lambda}, \\ d^{ha}(y, q) &= -\frac{\log(J_1 r)}{\lambda} + \frac{\log(-q)}{\lambda} + ry, \end{aligned}$$

where (e^*, k, J_1) are constants. In appendix A.4.1, I provide the expressions that these constants satisfy. Under the optimal contract, the evolution of the states is now:

$$\begin{aligned} dy_t &= \left[\frac{(r + k - 2\rho)}{r\lambda} + \sigma^2 \lambda \left(\frac{r}{2} - \frac{e^* k}{B} + \frac{(e^*)^2 k^2}{r B^2} \right) \right] dt + \sigma dW_t \\ dq_t &= (\rho - k) q_t dt - \frac{\sigma \lambda e^* k}{B} q_t dW_t. \end{aligned}$$

While the form of the policy functions is the same as in the full information case, the constants defining them differ. Solving for the values of the constants is a simple numerical task, but explicit analytic expressions are not available. To gain some additional

insight into the optimal contract, I expand e^* and k in σ^2 around zero. From (A.6) and (A.7) I have the following approximations:

$$\begin{aligned} e^* &= B - \sigma^2 \frac{r\lambda}{B} + o(\sigma^4) \\ k &= r - \sigma^2 r^2 \lambda^2 + o(\sigma^4). \end{aligned}$$

In turn, substituting these approximations into $c^{ha}(q)$ gives:

$$c^{ha}(q) = c^{fi}(q) + o(\sigma^4).$$

Thus the first order effects (in the shock variance) of the information frictions are a reduction in effort put forth, but no change in consumption. I also show below that k is the agent's effective rate of return, and thus this return decreases with more volatility. Moreover, effort varies with the parameters in a simple way: a greater rate of return parameter r or risk aversion parameter λ or smaller productivity values B lead to larger reductions in effort. Below I plot the exact solutions for a parameterized version of the model and show that the results are in accord with these first order asymptotics. Thus the information friction leads to a reduction in effort, but has little effect on consumption.

7.4. The hidden saving case

When the agent has access to hidden savings, some of the analysis is altered as we must consider the dynamics of m and the agent's choice of c . The agent's (current-value) Hamiltonian becomes:

$$H = (\gamma + Qm)(ry + Be - s - d) + p(rm + s - c) - \exp(-\lambda(c - e^2/2)),$$

and thus his optimality conditions for (e, c) are:

$$\begin{aligned} (\gamma + Qm)B &= \lambda e \exp(-\lambda(c - e^2/2)) \\ p &= \lambda \exp(-\lambda(c - e^2/2)). \end{aligned}$$

For the same reasons as discussed in section 5.3 above, the agent's Hamiltonian is not concave in (m, e, c) . Thus the sufficient conditions of Proposition 5.2 fail, and I cannot be sure the first-order approach is valid here. However I use it to derive a contract, and verify below that the candidate optimal contract is indeed implementable. The agent's optimality conditions determine the following functions:

$$\gamma(e, p) = \frac{ep}{B}, \quad c(e, p) = \frac{e^2}{2} - \frac{\log(p/\lambda)}{\lambda}.$$

Thus we also have $u(c(e, p), e) = -p/\lambda$. The evolution of the discounted marginal utility state, the discounted version of (18), is simply:

$$dp_t = (\rho - r)p_t dt + Q_t \sigma dW_t.$$

TABLE 1.

Comparison of the policies for the optimal contracts under full information, hidden actions, and hidden savings.

	Full Information	Hidden Action	Hidden Saving
Effort, e	B	$e^*(k) = \frac{B^3 r + \sigma^2 \lambda B r k}{B^2 r + 2\sigma^2 \lambda k^2}$	$e^*(r) = \frac{B^3 r + \sigma^2 \lambda B r^2}{B^2 r + 2\sigma^2 \lambda r^2}$
Consumption, c	$\frac{B^2}{2} - \frac{\log r}{\lambda} - \frac{\log(-q)}{\lambda}$	$\frac{e^*(k)^2}{2} - \frac{\log k}{\lambda} - \frac{\log(-q)}{\lambda}$	$\frac{e^*(r)^2}{2} - \frac{\log r}{\lambda} - \frac{\log(-q)}{\lambda}$
Dividend, d	$-\frac{\log(J_0 r)}{\lambda} + \frac{\log(-q)}{\lambda} + r y$	$-\frac{\log(J_1 r)}{\lambda} + \frac{\log(-q)}{\lambda} + r y$	$-\frac{\log(F(-\lambda r))}{\lambda} + \frac{\log(-q)}{\lambda} + r y$

For a given p , the principal's value function $J(y, q, p)$ solves the HJB equation:

$$\rho J = \max_{d, e, Q} \left\{ -\exp(-\lambda d) + J_y [r y + B e - c(e, p) - d] + J_q [\rho q + p/\lambda] + J_p (\rho - r) p + \sigma^2 \left(\frac{1}{2} J_{yy} + J_{yq} \gamma(e, p) + J_{yp} Q + \frac{1}{2} J_{qq} \gamma(e, p)^2 + J_{qp} \gamma(e, p) Q + \frac{1}{2} J_{pp} Q^2 \right) \right\} \quad (33)$$

where I substitute $\gamma = \gamma(e, p)$ and $c = c(e, p)$ and suppress arguments of J . Then if the value function is concave in p , the optimal initial condition p_0^* is determined by:

$$J_p(y, q, p_0^*) = 0.$$

Due to the proportionality of utility and marginal utility with exponential preferences, the value function can be written in the form:

$$J(y, q, p) = \frac{F(p/q)}{q} \exp(-r \lambda y)$$

for some univariate function F . In appendix A.4.2, I show that F is the solution of a second-order ordinary differential equation, and derive the principal's optimality conditions in terms of this function. Thus in general the contract depends upon p_t/q_t , the ratio of promised utility to the current marginal utility of consumption. However I also show in appendix A.4.2 that the optimal choice of initial marginal utility is $p_0^* = -\lambda r q_0$. With this initialization, the key ratio remains constant over the contract: $p_t/q_t = p_0^*/q_0 = -\lambda r$ for all t . Thus the contract does not need to condition on the additional marginal utility state. This is a consequence of the absence of wealth effects, and thus special to exponential utility. The policy functions are once again very similar, with only the constants differing. As discussed in more detail below, hidden savings limits the ability of the principal to provide intertemporal incentives, by requiring that the agent's effective return on assets be equated to the risk-free rate r .

7.5. Comparing the different cases

Table 1 summarizes the optimal contracts under full information, hidden actions, and hidden savings. As we've seen, the policy functions bear a strong resemblance to each other. In each case, effort is constant, consumption is linear in the log of the promised utility q , and the principal's dividend is linear in $\log(-q)$ and the current assets y . The consumption policies also depend negatively on the agent's effective rate of return on assets, which is r with full information, k with hidden actions, and r with hidden savings. As we've already seen, the main effect of the hidden action case relative to full

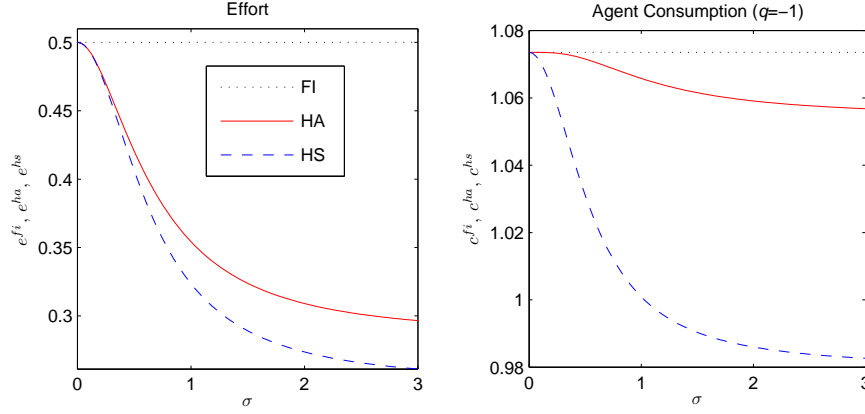


FIGURE 1. The agent's effort and consumption policies for different noise levels σ .

information is the reduction of effort. We've also seen that to first order there is no effect on consumption. This can be seen in the consumption policies, as effort falls but so does the effective return as $k < r$, leading to at least partially offsetting effects.

The main difference between the hidden saving and the hidden action cases is the effective rate of return on the agent's savings. If the return on the agent's saving were k then the optimal policies in the hidden action and hidden savings cases would coincide. But the access to hidden saving limits the ability of the principal to provide intertemporal incentives. If $k < r < 4k$ then effort is decreasing in k , and hence consumption is as well.²¹ Since for small shocks we've seen that k is smaller than r by a term proportional to the shock variance, this will certainly hold for small shocks. Moreover $c^{ha}(q; k)$ is decreasing in k if $e^*(k)$ is, as effort falls but the effective return rises. Thus at least for small enough σ , with hidden savings the agent puts forth less effort and consumes less than with hidden actions alone.

The analytic results are also borne out numerically. In Figure 1 I plot the consumption functions for effort and consumption in the full information, hidden action, and hidden savings cases for a particular parameterization of the model. I set $r = 0.15$, $B = 0.5$, $\rho = 0.1$, $\lambda = 2$, and show the results for varying σ . In particular, the left panel plots effort versus σ , while the right panel plots consumption (evaluated at $q = -1$) versus σ . Clearly as $\sigma \rightarrow 0$ the cases all agree, as the information friction vanishes. Compared to full information, both effort and consumption fall under hidden actions. Moreover, for small σ effort falls but consumption is relatively unaffected, as the approximations suggest. When the agent has access to hidden savings, consumption and effort fall further. In addition, these effects are all monotone in σ .

The policy functions also provide explicit expressions for the inefficiency "wedges," discussed by Kocherlakota (2004b) and others. These measure how the information frictions lead the consumption and labor allocations to be distorted for incentive reasons. In

²¹Simple calculations give

$$\frac{de^*(k)}{dk} = \frac{\sigma^2 B^3 r \lambda (r - 4k) - 2\sigma^4 B r \lambda^2 k}{(B^2 r + 2\sigma^2 \lambda k^2)^2},$$

which is clearly negative if $r < 4k$.

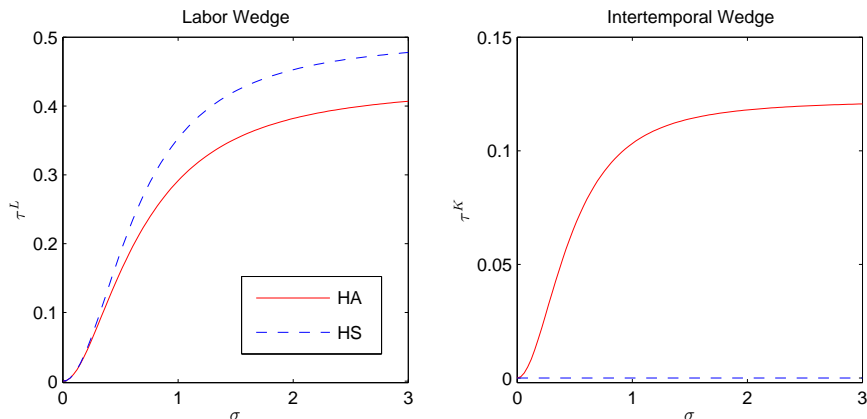


FIGURE 2. The labor and intertemporal wedges for different noise levels σ .

particular, suppose that the agent were able borrow and lend at the same risk free rate r as the principal. When the agent's saving is observable, the principal is able to tax it and drive its return down to k . Thus in the hidden action case, the contract introduces an intertemporal wedge τ^K , a gap between the intertemporal marginal rate of substitution and the marginal return on assets. This is simply given by the tax rate which drives the after-tax rate of return down to k :

$$\tau^K(k) = 1 - \frac{k}{r}.$$

However if the principal cannot observe the agent's savings, then of course he cannot tax it, so $\tau^K(r) = 0$.

By varying the payment to the agent for incentive reasons, the optimal contract also induces a labor wedge τ^L —a gap between the marginal productivity of effort and the marginal rate of substitution between consumption and effort. As the marginal product of effort is B and the marginal rate of substitution is $e^*(k)$, the labor wedge is simply:

$$\tau^L(k) = 1 - \frac{e^*(k)}{B}.$$

Thus in the hidden saving case, the labor wedge increases relative to the hidden action case, as $\tau^L(r) > \tau^L(k)$ for $k < r < 4k$.

Both the intertemporal and labor wedges are constant, and are shown in figure 2 for varying σ . The labor wedge is especially significant for this parameterization, as for $\sigma > 1.5$ it is comparable to labor income tax rates of more than 30% under hidden actions and roughly 40% under hidden savings. The intertemporal wedge is smaller, of course being identically zero under hidden savings and flattening out for large σ near an effective tax rate of 13% under hidden actions.

Finally, reductions in the principal's consumption provide a measure of the cost of the information frictions. As Table 1 shows, the different informational assumptions affect d only through the additive constants. Thus for each level of promised utility q and assets y ,




FIGURE 3. Reduction in the principal's dividend relative to the full information case for different noise levels σ .

the principal's consumption differs by a constant amount depending on the information structure. In Figure 3 I plot the reduction relative to the full information case in the level of the principal's dividend under hidden actions and hidden savings. The costs are relatively low, flattening out near 0.025 units of consumption under hidden actions and 0.03 units under hidden savings. As an example, with $y = 0$ and $q = -1$ the dividend is exactly equal to the constant term, and these level reductions imply a 1.5-2% fall in the principal's dividend. Of course with greater output or lower levels of promised utility, the proportional decline is much lower.

7.6. Implementing the optimal allocations

Thus far I have focused on the direct implementation of contracts, with the principal assigning consumption to the agent. In this section I show how the same outcomes can be achieved by giving the agent a performance-related payment, but then allowing him to invest in a risk-free asset and thus to choose his own consumption. This also allows me to verify that the contract is indeed implementable in the hidden saving case.

Suppose that the principal makes a payment to the agent of the history-dependent form $s_t = s(q_t; \hat{r}) = c^{ha}(q_t; \hat{r})$ where the after-tax return on saving is $\hat{r} = k$ in the hidden action case and $\hat{r} = r$ in the hidden savings case. This embodies an associated target effort $e^*(\hat{r})$ and evolution for q_t . The contract also specifies a tax rate $\tau^K(\hat{r}) = 1 - \hat{r}/r$ on saving. For the agent, the state q_t is simply part of the specification of the payment under the contract. From his vantage point, it evolves as:

$$\begin{aligned}
 dq_t &= (\rho - \hat{r})q_t dt - \frac{\sigma \lambda e^*(\hat{r}) \hat{r}}{B} q_t dW_t^{e^*} \\
 &= (\rho - \hat{r})q_t dt - \frac{\sigma \lambda e^*(\hat{r}) \hat{r}}{B} q_t \left(\frac{dy_t - (ry_t + B e^*(\hat{r}) - s(q_t; \hat{r}) - d_t) dt}{\sigma} \right) \\
 &= [\rho - \hat{r} - \lambda e^*(\hat{r}) \hat{r} (e_t - e^*(\hat{r}))] q_t dt - \frac{\sigma \lambda e^*(\hat{r}) \hat{r}}{B} q_t dW_t.
 \end{aligned}$$

Here I use the fact that $W_t^{e^*}$ is the driving Brownian motion under the optimal contract for the principal's information set. Moreover the agent begins with zero wealth $m_0 = 0$, which then evolves as:

$$\begin{aligned} dm_t &= ([1 - \tau^K(\hat{r})]rm_t + s(q_t; \hat{r}) - c_t) dt \\ &= \left(\hat{r}m_t + \frac{e^*(\hat{r})^2}{2} - \frac{\log(\hat{r})}{\lambda} - \frac{\log(-q_t)}{\lambda} - c_t \right) dt \end{aligned}$$

Thus the agent's value function $V(q, m)$ solves the HJB equation:

$$\begin{aligned} \rho V(q, m) &= \max_{c, e} \left\{ -\exp(-\lambda(c - e^2/2)) + V_m(q, m) [\hat{r}m + s(q; \hat{r}) - c] \right. \\ &\left. + V_q(q, m)q[\rho - \hat{r} - \lambda e^*(\hat{r})\hat{r}(e_t - e^*(\hat{r}))] + \frac{1}{2}V_{qq}(q, m)q^2 \frac{\sigma^2 \lambda^2 e^*(\hat{r})^2 \hat{r}^2}{B^2} \right\} \end{aligned}$$

It is easy to verify that the agent's value function is given by $V(q, m) = q \exp(-\lambda \hat{r} m)$. Substituting this into the HJB equation, and taking first order conditions for (c, e) gives:

$$\begin{aligned} \lambda \exp(-\lambda(c - e^2/2)) &= V_m = -\lambda \hat{r} q \exp(-\lambda \hat{r} m), \\ \lambda e \exp(-\lambda(c - e^2/2)) &= -\lambda e^*(\hat{r}) \hat{r} q V_q = -\lambda e^*(\hat{r}) \hat{r} q \exp(-\lambda \hat{r} m). \end{aligned}$$

Taking ratios of the two equations gives $e = e^*$, and thus the target effort level is implementable. Since effort is independent of m , the sufficient condition (29) holds. Thus the contract is indeed implementable.

It is also easy to directly verify implementability. The optimality condition for c gives:

$$c = \frac{e^*(\hat{r})^2}{2} - \frac{\log \hat{r}}{\lambda} - \frac{\log(-q)}{\lambda} + \hat{r}m.$$

Substituting this into the wealth equation gives $dm_t = 0$. Thus, if the agent begins with $m_t = 0$, then he will remain at zero wealth, will consume the optimal amount under the contract $c = c^{ha}(q; \hat{r})$, and will attain the value $V(q, 0) = q$. Therefore the policies associated with the optimal contracts can be implemented with this payment and savings tax scheme. In particular, the contract that we derived in the hidden saving case is indeed implementable, even though the sufficient concavity conditions do not hold.

8. CONCLUSION

In this paper I have established several key results for dynamic principal-agent problems. By working in a continuous time setting, I was able to take advantage of powerful results in stochastic control, which led to some sharp conclusions. I characterized the class of implementable contracts in dynamic settings with hidden actions and hidden states via a first-order approach, providing a general proof of its validity in a dynamic environment. I showed that implementable contracts must respect some additional state variables: the agent's utility process and the agent's "shadow value" (in marginal utility terms) of the hidden states. I then developed a constructive method for solving for

an optimal contract. The optimal contract is in general history dependent, but can be written recursively in terms of a small number of state variables.

As an application, I showed that as in Holmstrom and Milgrom (1987) the optimal contract is linear in a fully dynamic setting with exponential utility. However now the payment is linear in an endogenous object, the logarithm of the agent's promised utility under the contract. Moreover, I showed that the main effect of hidden actions is to reduce effort, with a smaller effect on the agent's implicit rate of return under the contract, which in turn affects consumption. Introducing hidden savings eliminates this second distortion, and increases the effort distortion.

Overall, the methods developed in the paper are tractable and reasonably general, making a class of models amenable to analysis. In Williams (2008) I apply the results in this paper to study borrowing and lending contracts when agents have persistent private information. Given the active literature on dynamic models with information frictions which has developed in the past few years, there is broad potential scope for further applications of my results.

APPENDIX

A.1. DETAILS OF THE CHANGE OF MEASURE

Here I provide technical detail associated with the change of measure in Section 3.2. I start by working with the induced distributions on the space of continuous functions, which I take to be the underlying probability space. Thus I let the sample space Ω be the space C^{n_y} , and let $W_t^0 = \omega(t)$ be the family of coordinate functions, and $\mathcal{F}_t^0 = \sigma\{W_s^0 \mid s \leq t\}$ the filtration generated by W_t^0 . I let P be the Wiener measure on $(\Omega, \mathcal{F}_T^0)$, and let \mathcal{F}_t be the completion of \mathcal{F}_t^0 with the null sets of \mathcal{F}_T^0 . This defines the basic (canonical) filtered probability space, on which is defined the Brownian motion W_t^0 . I now introduce some regularity conditions on the diffusion coefficient in (11) following Elliott (1982).

ASSUMPTIONS A.1. Denote by $\sigma_{ij}(t, y, s)$ the (i, j) element of the matrix $\sigma(t, y, s)$. Then I that for all i, j , for some fixed K independent of t and i, j :

1. σ_{ij} is continuous,
2. $|\sigma_{ij}(t, y, s) - \sigma_{ij}(t, y', s')| \leq K(|y - y'| + |s - s'|)$,
3. $\sigma(t, y, s)$ is non-singular for each (t, y, s) and $|(\sigma^{-1}(t, y, s))_{ij}| \leq K$.

Of these, the most restrictive is perhaps the nonsingularity condition, which requires that noise enter every part of the state y . Under these conditions, there exists a unique strong solution to the stochastic differential equation:

$$dy_t = \bar{\sigma}(t, \bar{y}, \bar{Z})dW_t^0, \quad (\text{A.1})$$

with y_0 given. This is the evolution of output under an effort policy e^0 which makes $\bar{f}(t, \bar{y}, m_t, e_t^0) = 0$ at each date. Different effort choices alter the evolution of output by changing the distribution over outcomes in C^{n_y} .

Now I state some required regularity conditions on the drift function f .

ASSUMPTIONS A.2. I assume that for some fixed K independent of t :

1. f is continuous,
2. $|f(t, y, m, e, s)| \leq K(1 + |m| + |y| + |s|)$.

These imply the predictability, continuity, and linear growth conditions on the concentrated function \bar{f} which are assumed by Elliott (1982). Then for $\bar{e} \in \mathcal{A}$ I define the family of \mathcal{F}_t -predictable processes:

$$\Gamma_t(\bar{e}) = \exp\left(\int_0^t \bar{\sigma}^{-1}(v, \bar{y})\bar{f}(v, \bar{y}, m_v, e_v)dW_v^0 - \frac{1}{2}\int_0^t |\bar{\sigma}^{-1}(v, \bar{y})\bar{f}(v, \bar{y}, m_v, e_v)|^2 dv\right).$$

Under the conditions above, Γ_t is an \mathcal{F}_t -martingale (as the assumptions insure that Novikov's condition is satisfied) with $E[\Gamma_T(\bar{e})] = 1$ for all $\bar{e} \in \mathcal{A}$. Thus by the Girsanov theorem, I can define a new measure $P_{\bar{e}}$ via:

$$\frac{dP_{\bar{e}}}{dP} = \Gamma_T(\bar{e}),$$

and the process $W_t^{\bar{e}}$ defined by:

$$W_t^{\bar{e}} = W_t^0 - \int_0^t \bar{\sigma}^{-1}(v, \bar{y}) \bar{f}(v, \bar{y}, m_v, e_v) dv$$

is a Brownian motion under $P_{\bar{e}}$. Thus from (A.1), it's clear that the state follows the SDE:

$$dy_t = \bar{f}(t, \bar{y}, m_t, e_t) dt + \bar{\sigma}(t, \bar{y}) dW_t^{\bar{e}}.$$

Hence each effort choice \bar{e} results in a different Brownian motion. Γ_t defined above (suppressing \bar{e}) satisfies $\Gamma_t = E[\Gamma_T | \mathcal{F}_t]$, and thus is the relative density process for the change of measure.

A.2. ASSUMPTIONS FOR RESULTS

Here I list some additional regularity conditions. The first is a differentiability assumption for the agent's problem.

ASSUMPTION A.3. *The functions (u, v, f, b) are continuously differentiable in m .*

The next sets of assumptions guarantee that the local incentive constraint is satisfied in the hidden action and hidden state cases.

ASSUMPTIONS A.4. *The function u is concave in e , f is concave in e , and $u_e f_e \leq 0$.*

ASSUMPTIONS A.5. *The function u is separable in e and c and concave in (e, c) , b is concave in c , and $u_c b_c \leq 0$.*

A.3. PROOFS OF RESULTS

Proof (Proposition 4.1). This follows by applying the results in Bismut (1973)-(1978) to my problem. Under the assumptions, the maximum principle applies to the system with (Γ, x) as the state variables. I now illustrate the calculations leading to the Hamiltonian (16) and the adjoint equations (17)-(18). I first define the stacked system:

$$X_t = \begin{bmatrix} \Gamma_t \\ x_t \end{bmatrix}, \quad \Theta_t = \begin{bmatrix} q_t \\ p_t \end{bmatrix}, \quad \Lambda_t = \begin{bmatrix} \gamma_t \\ Q_t \end{bmatrix}.$$

Then note from (14)-(15) that X_t satisfies (suppressing arguments):

$$\begin{aligned} dX_t &= \Gamma_t \begin{bmatrix} 0 \\ \bar{b} \end{bmatrix} dt + \Gamma_t \begin{bmatrix} \sigma^{-1} \bar{f} \\ m_t \sigma^{-1} \bar{f} \end{bmatrix} dW_t^0 \\ &= M(t) dt + \Sigma(t) dW_t^0 \end{aligned}$$

Thus, as in Bismut (1978), the Hamiltonian for the problem is:

$$\tilde{H} = \Theta M + \text{tr}(\Lambda' \Sigma) + \Gamma \bar{u} = \Gamma H,$$

where H is from (16). As $\Gamma \geq 0$, the optimality condition (19) is the same with H or \tilde{H} . The adjoint variables follow:

$$\begin{aligned} d\Theta_t &= -\frac{\partial \tilde{H}}{\partial X}(t) dt + \Lambda_t dW_t^0 \\ \Theta_T &= \frac{\partial(\Gamma_T v(y_T, m_T))}{\partial X_T}. \end{aligned}$$

By carrying out the differentiation and simplifying I arrive at (17)-(18). ■

Proof (Proposition 5.1). The result is an extension of Theorem 4.2 in Schattler and Sung (1993). The necessity of the conditions follow directly from my results above: if the contract is implementable then it clearly must satisfy the participation constraint, and by Proposition 4.1 it must satisfy promise-keeping and be locally incentive compatible.

To show the converse, I verify that \hat{e} is an optimal control when the agent faces the contract (s, \hat{e}) . Recall that the expected utility from following \hat{e} is given by $V(\hat{e}) = \hat{q}_0$. From the results above, for any $\bar{e} \in \mathcal{A}$ the following holds:

$$\begin{aligned} V(\bar{e}) - V(\hat{e}) &= E_{\bar{e}} \left[\int_0^T [u(t, y_t, e_t, s_t) - u(t, y_t, \hat{e}_t, s_t)] dt + \int_0^T \hat{\gamma}_t \sigma(t, y_t, s_t) dW_t^{\bar{e}} \right] \\ &= E_{\bar{e}} \left[\int_0^T [u(t, y_t, e_t, s_t) - u(t, y_t, \hat{e}_t, s_t)] dt + \int_0^T \hat{\gamma}_t \sigma(t, y_t, s_t) dW_t^{\bar{e}} \right] \\ &\quad + E_{\bar{e}} \left[\int_0^T \hat{\gamma}_t [f(t, y_t, e_t, s_t) - f(t, y_t, \hat{e}_t, s_t)] dt \right] \\ &= E_{\bar{e}} \left[\int_0^T [H^*(t, e_t) - H^*(t, \hat{e}_t)] dt + \int_0^T \hat{\gamma}_t \sigma(t, y_t, s_t) dW_t^{\bar{e}} \right] \\ &\leq E_{\bar{e}} \left[\int_0^T \hat{\gamma}_t \sigma(t, y_t, s_t) dW_t^{\bar{e}} \right] = 0. \end{aligned}$$

Here the first equality uses (22), the second equality uses the definitions of the change of measure between $W^{\hat{e}}$ and $W^{\bar{e}}$, the third equality uses the definition of the H^* function in (23), the inequality follows from the local incentive constraint (24), and the final result uses the fact that the stochastic integral is a martingale due to the square integrability of the $\hat{\gamma}_t \sigma(t)$ process. Thus since \bar{e} was arbitrary, the agent can achieve at most the utility $V(\hat{e})$ which is greater than his reservation level by assumption. Thus \hat{e} is an optimal control, and so the contract is implementable. \blacksquare

Proof (Proposition 5.2). This proof combines ideas from Schattler and Sung (1993) with those in Zhou (1996). The necessity of the conditions follows directly as in Proposition 5.1 above, as it is a consequence of Proposition 4.1.

To show the converse, I consider an arbitrary (\bar{e}, \bar{c}) which imply a process $\bar{m} = \{m_t\}$. Under my assumptions the local incentive constraint (28) ensures that:

$$H_e^{**}(t, \hat{m}_t, \hat{e}_t, \hat{c}_t) = H_c^{**}(t, \hat{m}_t, \hat{e}_t, \hat{c}_t) = 0.$$

Thus by the concavity of H^{**} , for any (m_t, e_t, c_t) we get:

$$H^{**}(t, m_t, e_t, c_t) - H^{**}(t, \hat{m}_t, \hat{e}_t, \hat{c}_t) \leq \langle H_m^{**}(t, \hat{m}_t, \hat{e}_t, \hat{c}_t), m_t - \hat{m}_t \rangle. \quad (\text{A.2})$$

I extend the shorthand notation above to now write $b(\hat{t}) = b(t, \hat{m}_t, \hat{c}_t, s_t)$ and so on. Then I define $\Delta_t = m_t - \hat{m}_t$, and I can write its evolution as:

$$d\Delta_t = [b_m(\hat{t})\Delta_t + \alpha_t] dt, \quad (\text{A.3})$$

with $\Delta_0 = 0$ where I define:

$$\alpha_t = -b_m(\hat{t})\Delta_t + b(t) - b(\hat{t})$$

Then we have the following relationship between \hat{p}_t in (18) for the target policy and Δ_t in (A.3):

$$\begin{aligned} v_m(\hat{T})\Delta_t &= \hat{p}_T \Delta_T = \int_0^T [\Delta_t d\hat{p}_t + \hat{p}_t d\Delta_t] \\ &= \int_0^T \left(\alpha_t \hat{p}_t - \Delta_t \left[(\hat{\gamma}_t + \hat{Q}_t \hat{m}_t) f_m(\hat{t}) + \hat{Q}_t f(\hat{t}) + u_m(\hat{t}) \right] \right) dt + \int_0^T \Delta_t \hat{Q}_t \sigma(t) dW_t^0 \end{aligned}$$

Hence by the concavity of v I have:

$$E_{\bar{e}}[v(T) - v(\hat{T})] \leq E_{\bar{e}}[v_m(\hat{T})\Delta_T] = E_{\bar{e}} \left[\int_0^T \left(\hat{p}_t \alpha_t - \left[(\hat{\gamma}_t + \hat{Q}_t \hat{m}_t) f_m(\hat{t}) + u_m(\hat{t}) + \hat{Q}_t (f(t) - f(\hat{t})) \right] \Delta_t \right) dt \right], \quad (\text{A.4})$$

where the final term in \hat{Q}_t is due to the change from the base measure P_0 to $P_{\bar{e}}$ (associated with \bar{e}).

Then I proceed as in Proposition 5.1 and note that for any $(\bar{e}, \bar{c}) \in \mathcal{A}$ I have:

$$\begin{aligned}
V(\bar{e}, \bar{c}) - V(\hat{e}, \hat{c}) &= E_{\bar{e}} \left[\int_0^T [u(t) - u(\hat{t})] dt + \int_0^T [\hat{\gamma}_t + \hat{Q}_t \hat{m}_t] \sigma(t) dW_t^{\bar{e}} + v(T) - v(\hat{T}) \right] \\
&= E_{\bar{e}} \left[\int_0^T [u(t) - u(\hat{t}) + (\hat{\gamma}_t + \hat{Q}_t \hat{m}_t)(f(t) - f(\hat{t}))] dt + \int_0^T [\hat{\gamma}_t + \hat{Q}_t \hat{m}_t] \sigma(t) dW_t^{\bar{e}} + v(T) - v(\hat{T}) \right] \\
&\leq E_{\bar{e}} \left[\int_0^T \left(u(t) - u(\hat{t}) + [\hat{\gamma}_t + \hat{Q}_t \hat{m}_t] f(\hat{t}) + \hat{p}_t \alpha_t - \left[(\hat{\gamma}_t + \hat{Q}_t \hat{m}_t) f_m(\hat{t}) + u_m(\hat{t}) + \hat{Q}_t (f(t) - f(\hat{t})) \right] \Delta_t \right) dt \right] \\
&= E_{\bar{e}} \left[\int_0^T \left(u(t) - u(\hat{t}) - u_m(\hat{t}) \Delta_t + (\hat{\gamma}_t + \hat{Q}_t \hat{m}_t)(f(t) - f(\hat{t})) - (\hat{\gamma}_t + \hat{Q}_t \hat{m}_t) \Delta_t f_m(\hat{t}) \right) dt \right] \\
&\quad + E_{\bar{e}} \left[\int_0^T (\hat{p}_t [b(t) - b(\hat{t}) - b_m(\hat{t}) \Delta_t]) dt \right] \\
&= E_{\bar{e}} \left[\int_0^T [H^{**}(t) - H^{**}(\hat{t}) - H_m^{**}(\hat{t}) \Delta_t] dt \right] \leq 0.
\end{aligned} \tag{A.5}$$

Here the first equality uses (27), while the second equality uses the definitions of the change of measure between $P_{\bar{e}}$ and $P_{\hat{e}}$. The following inequality uses (A.4), while the next equality uses the definition of α_t . The following equality uses the definition of H^{**} in (26), and the final result follows from its concavity as in (A.2). Thus since (\bar{e}, \bar{c}) was arbitrary, we see that (\hat{e}, \hat{c}) is an optimal control, and the contract is thus implementable.

The weaker sufficient condition (29) stated in section 5.3 uses (A.5) along with the fact that the local incentive constraint implies that $H^{**}(t, \hat{m}_t, \hat{e}_t, \hat{c}_t) \geq H^{**}(t, \hat{m}_t, \bar{e}_t, \bar{c}_t)$. ■

Proof (Corollary 5.1 and 5.2). In either case, the stated assumptions insure that H^* is concave in e and H^{**} is concave in (e, c) . This is clear from (21) in the hidden action case. The parallel first order conditions give the result in the hidden state case, where the separability insures that I can simply look at the separate first order conditions. But by assumption, each function has a stationary point on the feasible set, which is then a maximum. ■

A.4. DETAILED CALCULATIONS FOR THE EXAMPLE

A.4.1. The hidden action case

Here we provide explicit expressions for the key constants which determine the policy functions in the hidden action case. Note that $u(t, c(q), e^*) = e^{-\rho t} k q$. Using this along with the form of the value function in the first order conditions for (e, c) determines the equations that e^* and k must satisfy:

$$\begin{aligned}
r - k + \sigma^2 r \lambda^2 e^* k / B - 2\sigma^2 \lambda^2 (e^*)^2 k^2 / B^2 &= 0, \\
-Br + \lambda e^* k - \sigma^2 r \lambda (1 + \lambda (e^*)^2) k / B - 2\sigma^2 \lambda e^* (1 + \lambda (e^*)^2) k^2 / B^2 &= 0.
\end{aligned} \tag{A.6}$$

These can be solved to get $e^*(k)$:

$$e^* = \frac{B^3 r + \sigma^2 \lambda B r k}{B^2 r + 2\sigma^2 \lambda k^2}, \tag{A.7}$$

then substituting this back into the first equation in (A.6) gives an expression for k . Notice that for $\sigma = 0$ these collapse to the full information solution $e^* = B$ and $k = r$. The constant J_1 is determined from the HJB equation after substituting in the optimal policies, which implies:

$$J_1 = \frac{1}{kr} \exp \left(\frac{r + k - 2\rho}{rA} - \lambda e^* \left(B - \frac{e^*}{2} \right) + \sigma^2 \lambda^2 \left(\frac{r}{2} - \frac{e^* k}{B} + \frac{(e^*)^2 k^2}{rB^2} \right) \right).$$

A.4.2. The hidden saving case

I now derive the optimality conditions in the hidden savings case. To simplify notation, I define $z = p/q$ as the argument of the F function, noting that $z \in (-\infty, 0)$. Thus the optimality condition for p_0^* translates to $F'(z_0^*) = 0$, or $p_0^* = (F')^{-1}(0)q_0$. Using the guess of the form of the solution in the HJB equation (33), I take the first order condition

for the dividend, and obtain a policy function similar to the hidden action case:

$$d^{hs}(y, q, z) = -\frac{\log(F(z)r)}{\lambda} + \frac{\log(-q)}{\lambda} + ry.$$

For the volatility of marginal utility, I find it convenient to normalize $Q_t = \hat{Q}_t p_t$. Then the optimal choices of e and \hat{Q} are determined by the first order conditions:

$$\begin{aligned} -\lambda r F(z)(B - e) + (F''(z)z^2 + 4F'(z)z + 2F(z)) \frac{\sigma^2 z^2 e}{B^2} \\ + \lambda r (F'(z)z + F(z)) \frac{\sigma^2 z}{B} - (2F'(z) + F''(z)z) \hat{Q} \frac{\sigma^2 z^2}{B} = 0, \end{aligned} \quad (\text{A.8})$$

$$F''(z)z \hat{Q} - \lambda r F'(z) - (2F'(z) + F''(z)z) \frac{ez}{B} = 0. \quad (\text{A.9})$$

Thus given F , these equations determine the optimal policies as a function of z alone: $e^{hs}(z)$ and $\hat{Q}(z)$. Substituting these into (33), I find that F is the solution of the second order ODE:

$$\begin{aligned} \rho F(z) = rF(z) - (F'(z)z + F(z))(\rho + z/\lambda) - \lambda r F(z) \left[B e^{hs}(z) - \frac{e^{hs}(z)^2}{2} + \frac{\log(-rzF(z)/\lambda)}{\lambda} \right] \\ + (\rho - r)zF'(z) + (F''(z)z^2 + 4F'(z)z + 2F(z)) \frac{\sigma^2 e^{hs}(z)^2 z^2}{2B^2} + (F'(z)z + F(z)) \frac{\sigma^2 \lambda r e^{hs}(z)z}{B} \\ + \sigma^2 \left(\frac{\lambda^2 r^2}{2} F(z) + F''(z)z^2 \frac{\hat{Q}(z)^2}{2} - \lambda r F'(z)z \hat{Q}(z) - (2F'(z) + F''(z)z^2) \frac{e^{hs}(z)\hat{Q}(z)z^2}{B} \right), \end{aligned}$$

with the boundary conditions $\lim_{z \rightarrow 0} F(z) = \infty$ and $\lim_{z \rightarrow -\infty} F(z) = \infty$. These conditions imply that the principal's utility falls without bound as the agent's current consumption or future promised utility increase without bound.

Thus I have reduced the solution of the three dimensional PDE for the value function to an ODE. The ODE is sufficiently complex that an explicit solution for arbitrary z does not seem feasible. However, I can obtain more explicit results for the value and the policies at the optimal starting value z_0^* . At z_0^* (A.8) and (A.9) imply:

$$\hat{Q}(z_0^*) = \frac{e^{hs}(z_0^*)z_0^*}{B}, \quad e_0^* = e^{hs}(z_0^*) = \frac{B^3 r - \sigma^2 B r z_0^*}{B^2 r + 2\sigma^2 (z_0^*)^2 / \lambda}.$$

Then using these in the ODE we get:

$$F(z_0^*) = \frac{-\lambda}{r z_0^*} \exp \left(\frac{\lambda r - z_0^* - 2\lambda \rho}{\lambda r} - \lambda e_0^* \left(B - \frac{e_0^*}{2} \right) + \sigma^2 \lambda^2 \left(\frac{r}{2} + \frac{e_0^* z_0^*}{\lambda B} + \frac{(e_0^* z_0^*)^2}{\lambda r B^2} \right) \right).$$

Note that if $z_0^* = -\lambda k$ then these results agree with the hidden action case, as $e^{hs}(-\lambda k) = e^*$ from (A.7) and $F(-\lambda k) = J_1$. But this is indeed the optimal choice of z_0 when the return on the agent's saving differs from the principal's and is equal to k . This can be verified analytically by implicitly differentiating the ODE for F and using the expressions for $e^{hs}(z_0^*)$ and $\hat{Q}(z_0^*)$ to solve for z_0^* . However in general I cannot analytically solve for the optimal starting value, and so I numerically solve the ODE for $F(z)$, then choose z_0^* as the minimizing value. (These calculations are straightforward but lengthy and thus are not included here. Details can be furnished upon request.) For the parametrization below, I find that $z_0^* = -\lambda r$. I have also verified the same result in alternative parameterizations.

These results imply that z_t is in fact constant over time, which follows from the evolution of z_t :

$$dz_t = - \left[(z_t/\lambda + r)z_t + \frac{\sigma^2 e^{hs}(z_t)^2 z_t^2}{B} \left(\hat{Q}(z_t) - \frac{e^{hs}(z_t)z_t}{B} \right) \right] dt + \sigma z_t \left(\hat{Q}(z_t) - \frac{e^{hs}(z_t)z_t}{B} \right) dW_t.$$

Using the expression for $\hat{Q}(z_0^*)$ above, note that $dz_t = 0$ if $z_0^* = -\lambda r$. Therefore utility and marginal utility remain proportional forever under the contract.

References

Abraham, A. and N. Pavoni (2008). Efficient Allocations with Moral Hazard and Hidden Borrowing and Lending. *Review of Economic Dynamics*, forthcoming.

- 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.
- Allen, F. (1985). Repeated Principal-Agent Relationships with Borrowing and Lending. *Economics Letters* 17, 27–31.
- Bismut, J. M. (1973). Conjugate Convex Functions in Optimal Stochastic Control. *Journal of Mathematical Analysis and Applications* 44, 384–404.
- Bismut, J. M. (1975). Growth and the Optimal Intertemporal Allocation of Risks. *Journal of Economic Theory* 10, 239–257.
- Bismut, J. M. (1978). Duality Methods in the Control of Densities. *SIAM Journal on Control and Optimization* 16, 771–777.
- Bolton, P. and M. Dewatripont (2005). *Contract Theory*. Cambridge: MIT Press.
- Brock, W. A. and M. J. Magill (1979). Dynamics Under Uncertainty. *Econometrica* 47, 843–868.
- Cadenillas, A., J. Cvitanić, and F. Zapatero (2005). Dynamic Principal-Agent Problems with Perfect Information. Working paper, USC.
- Chari, V. V., P. J. Kehoe, and E. McGrattan (2007). Business Cycle Accounting. *Econometrica* 75, 781–836.
- Cole, H. L. and N. Kocherlakota (2001). Efficient Allocations with Hidden Income and Hidden Storage. *Review of Economic Studies* 68, 523–542.
- Cox, J., S. Ross, and M. Rubinstein (1979). Option Pricing: A Simplified Approach. *Journal of Financial Economics* 7, 229–263.
- Cvitanić, J., X. Wan, and J. Zhang (2006). Continuous-Time Principal-Agent Problems with Hidden Action and Lump-Sum Payment. Working paper, USC.
- Detemple, J., S. Govindaraj, and M. Loewenstein (2005). Wealth-Robust Intertemporal Incentive Contracts. Working paper, Boston University.
- Doepke, M. and R. M. Townsend (2006). Dynamic Mechanism Design with Hidden Income and Hidden Storage. *Journal of Economic Theory* 126, 235–285.
- El Karoui, N., S. Peng, and M. C. Quenez (1997). Backward Stochastic Differential Equations in Finance. *Mathematical Finance* 7, 1–71.
- Elliott, R. J. (1982). *Stochastic Calculus and Applications*. New York: Springer-Verlag.
- Fudenberg, D., B. Holmstrom, and P. Milgrom (1990). Short-Term Contracts and Long-Term Agency Relationships. *Journal of Economic Theory* 51, 1–31.
- Golosov, M., N. Kocherlakota, and A. Tsyvinski (2003). Optimal Indirect and Capital Taxation. *Review of Economic Studies* 70, 569–587.
- Hellwig, M. F. and K. M. Schmidt (2002). Discrete-Time Approximations of the Holmstrom-Milgrom Brownian-Motion Model of Intertemporal Incentive Provision. *Econometrica* 70, 2225–2264.

- Holmstrom, B. and P. Milgrom (1987). Aggregation and Linearity in the Provision of Intertemporal Incentives. *Econometrica* 55, 303–328.
- Hopenhayn, H. A. and J. P. Nicolini (1997). Optimal Unemployment Insurance. *Journal of Political Economy* 105, 412–38.
- Jewitt, I. (1988). Justifying the First-Order Approach to Principal-Agent Problems. *Econometrica* 56, 1177–1190.
- Karatzas, I. and S. E. Shreve (1991). *Brownian Motion and Stochastic Calculus* (2nd ed.). New York: Springer-Verlag.
- Kocherlakota, N. R. (2004a). Figuring Out The Impact of Hidden Savings on Optimal Unemployment Insurance. *Review of Economic Dynamics* 7, 541–554.
- Kocherlakota, N. R. (2004b). Wedges and taxes. *American Economic Review (Papers and Proceedings)* 94, 109–113.
- Ma, J. and J. Yong (1999). *Forward-Backward Stochastic Differential Equations and Their Applications*, Volume 1702 of *Lecture Notes in Mathematics*. New York: Springer-Verlag.
- Mirrlees, J. A. (1999). The Theory of Moral Hazard and Unobservable Behaviour: Part I. *Review of Economic Studies* 66, 3–21.
- Phelan, C. and E. Stacchetti (2001). Sequential Equilibria in a Ramsey Tax Model. *Econometrica* 69, 1491–1518.
- 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. (2008). A Continuous-Time Version of the Principal-Agent Problem. *Review of Economic Studies* 75, 957–984.
- Schattler, H. and J. Sung (1993). The First-Order Approach to the Continuous-Time Principal-Agent Problem with Exponential Utility. *Journal of Economic Theory* 61, 331–371.
- Spear, S. and S. Srivastava (1987). On Repeated Moral Hazard with Discounting. *Review of Economic Studies* 53, 599–617.
- Sung, J. (1997). Corporate Insurance and Managerial Incentives. *Journal of Economic Theory* 74, 297–332.
- Werning, I. (2002). Optimal Unemployment Insurance with Unobservable Saving. Working paper, MIT.
- Westerfield, M. (2006). Optimal Dynamic Contracts with Hidden Actions in Continuous Time. Working paper, USC.
- Williams, N. (2008). Persistent Private Information. Working paper, University of Wisconsin.
- Yong, J. and X. Y. Zhou (1999). *Stochastic Controls*. New York: Springer-Verlag.
- Zhou, X. Y. (1996). Sufficient Conditions of Optimality for Stochastic Systems with Controllable Diffusions. *IEEE Transactions on Automatic Control* AC-41, 1176–1179.