Optimal Dynamic Contracts with Hidden Actions in Continuous Time

Mark M. Westerfield*

University of Southern California

November, 2006

Abstract

I consider a dynamic hidden action problem in continuous time, and I present a general method for solving such problems. In a model with non-separable intermediate utility and the possibility that the agent earns observable outside income or perks, I show how the optimal contract can be written as a function only of the agent's financial wealth and implemented actions. I then show that the principal's choices of intermediate consumption and effort level have a real option property that has no analog in static or linear contracting models: the principal will employ the agent at a loss, hoping conditions improve, in order to keep the option of future employment open. Finally, I show how to relax the level of commitment to allow for state contingent employment termination and time-inconsistent incentives.

JEL Classifications: D0, D8, C7.

Keywords: Hidden Action, Continuous Time, Principal-Agent, Real Options, Moral Hazard.

---

*I would like to thank Daron Acemoglu, Tobias Adrian, Jakša Cvitanić, Harry Deangelo, Bengt Holmstrom, Ravi Jagannathan, Chris Jones, Oleksandr Lytvak, Paul Milgrom, Oguzhan Ozbas, Vincenzo Quadrini, Jianfeng Yu, and Fernando Zapatero for their helpful insights and comments.
Department of Finance and Business Economics, University of Southern California, Hoffman Hall-701, MC-1427, 701 Exposition Blvd., Ste. 701, Los Angeles, CA 90089-1427. Email: mwesterf@usc.edu. Website: http://www-rcf.usc.edu/~mwesterf/

1 Introduction

Principal-agent models are the most common way of analyzing economic situations that are not market-mediated, and they have been used extensively to address questions of how to provide incentives in settings where there are informational constraints and asymmetries. However, there has existed no simple way of understanding how the principal-agent relationship would work in general settings in which the agent's decisions and the underlying economic processes occur continuously. Moreover, as the problem is usually studied, there are a number of intuitive and technical problems, such as the potential infeasibility of the first-order approach and the calculational difficulties in many recursive formulations. I provide a general characterization of the dynamic principal-agent moral hazard/hidden action problem in continuous time, and I explicitly characterize its solution in an easy-to-understand way. I then analyze the solution, including its interesting real-option characteristics and comparative statics.

The basic hidden action model consists of a project owned by the principal but managed by the agent. It is assumed that the agent exerts some unobservable level of effort that determines how profitable the project is. My model allows for a fully general contract that pays both intermediate and terminal consumption. I allow the agent to have general, a non-separable utility function for intermediate consumption and effort ($u(c, \mu)$) and I allow the agent's choice of effort to affect the agent's wealth from observable outside activities or perks in addition to his salary. I also allow for stochastic volatility and assume only local (as opposed to global) concavity for the problem. With this, I can write the optimal contract as a single stochastic process taking the principal's choices as arguments – I can write the optimal contract for *any* effort level the principal chooses to implement in terms of the actions the contract induces and the agent's wealth. Importantly, this process involves only one state variable – the agent's financial wealth – and it does not involve solving any integrals or lagrange multiplier problems. I can then characterize the principal's problem as a maximization over two variables (effort and consumption or contract volatility and consumption). Any application can be solved from that point using simple and well understood optimal control (or numerical) techniques. However, I continue to consider the problem analytically to understand several properties of the contract, the most important of which has to do with the real option the principal has in employing the agent.

The primary contribution of this paper is a representation of any optimal contract for a very general class of principal-agent models that is simple to compute. In particular, the representation does not have any non-local wealth effects despite accepting any valid utility function: if the agent were terminated today with wealth W, then the wealth effects can be summarized as a function of the agent's coefficient of absolute risk aversion evaluated at W. This is surprising because it holds true even with non-CARA utility and with intermediate consumption, and it applies regardless of the effort choice that the contract induces. In other words, the contract is fully myopic, which differs from every other analysis of the dynamic principal-agent problem of which I am aware. The result follows from the continuous nature of the agent's decision problem: because each choice has an infinitesimal impact on the evolution of the economy, only the local curvature of the utility function matters. Because I can account for all the agent's costs and all the principal's promises as they occur, I can exactly identify the local curvature and characterize the agent's decision purely in terms of those costs and promises.

There is an interesting dynamic property of the optimal contract that has no analog in static models and is missing from linear contracting/exponential utility models. The principal has a real option with regards to employing the agent, and the principal's employment decision resembles options exercise decisions. While the principal will never terminate the agent if he can make a short term gain by employing him, the principal will frequently employ the agent even if he suffers a short term loss from doing so. The reason is that as the agent's past successes and failures accumulate, the principal's cost of reimbursing the agent for his effort changes, so the principal will employ the agent at a loss in the hope that the cost of employing the agent will move in an advantageous direction. Thus, the principal keeps his real option open by employing the agent and exercises the option by terminating him. The underlying variable that determines the option's value is the agent's financial wealth.

I can use the explicit form of the dynamic optimal contract to provide new insight into old properties. I am able to show, for example, that reward smoothing – the fact that if the agent does well today he will receive more consumption today and more consumption tomorrow – is a function of the principal's cost minimization decision and is not related to motivating the agent's effort choice. This result is because the agent's effort choice depends on the volatility of the agent's continuation value function rather than the distribution of intermediate consumption.

By analyzing the evolution of the agent's continuation value, I show exactly how the contract has a "memory". However, while comparative statics – for example, how past failures affect future opportunities – can be stated explicitly for a given date and state, they cannot be signed across the whole state space in most models, without strong assumptions on the form of the utility functions. Lastly, the analysis shows that standard comparative statics – that the consumption and effort choices of the principal are monotonic in the agent's continuation utility level – are not general. The consumption comparative static fails when there is an interaction between consumption and effort choice: when effort choice is constant in a particular region, the agent's consumption will be monotonic with his continuation utility, but this is not necessarily so when the principal is willing to change the implemented effort level. The effort comparative static fails because the principal faces two costs of employing the agent, the cost of reimbursing the agent for his effort and the cost of insuring the agent against risk. These costs may increase in opposite directions yielding a non-monotonic choice of effort.

The model can also be expanded to include a stochastic termination time – for example, the principal could choose to fire the agent or the agent could choose to quit if either party's continuation value fell low enough. If the principal can commit to letting the agent walk away with all of his earned wealth, then the agent can be kept on the edge of indifference between working and not working. My analysis shows what that level of wealth must be. The same condition allows me to state that the agent can be made willing to accept a re-negotiation in the contract if the principal wishes to implement time-inconsistent policies. However, this does require that the agent not have rational expectations about any potential re-negotiation. My analysis shows that the principal will wish to renegotiate the contract, and the agent will agree, any time it would improve the principal's welfare to transfer wealth to the agent.

The seminal paper in the continuous time contracting literature is Holmstrom and Milgrom (1987) which shows that when the problem does not involve wealth effects or a changing economic environment, the resulting contract is linear in output. This is an important economic result that is further systematized in Schättler and Sung (1993), Sung (1995), and Ou-Yang (2003). Calculating equilibria with wealth effects and changing economic conditions has been mostly addressed in discrete time formulations, first in two period models such as Rogerson (1985b). Spear and Srivastava (1987) simplify the problem by using the agent's continuation value as a state variable and Phelan and Townsend (1991) extend the dynamic analysis to many other types of problems. A useful discussion of perks and other non-financial benefits and their economic effects is contained in Marino and Zábojník (2006).

A number of recent important papers address the continuous time problem in order to simplify the calculations. Sannikov (2006) created a continuous time model for which the solution can be characterized by an ordinary differential equation. DeMarzo and Sannikov (2004) shows how this approach can work in an application to agency costs and capital structure. Cadenillas, Cvitanić, and Zapatero (2005) shows how the problem can be solved with full information, while Cvitanić and Zhang (2006) extends the Holmstrom and Milgrom (1987) model with adverse selection and non-exponential utility functions. Williams (2004) uses a general approach, including hidden savings, and is able to characterize the solution as a system of forward-backward stochastic differential equations. While the economy I analyze is nearly as general as that in Williams (2004), my methods and the simplicity of the contract form go significantly further than other papers in obtaining closed form solutions to general problems and in evaluating the contract itself.

The paper proceeds as follows: Section 2 lays out the model, including the information sets and objective functions. Section 3 analyzes the optimal contract using the principal's assessment of the agent's problem and the agent's actual maximization decision. Section 4 follows with a simplification of the principal's problem and some comparative statics. The real options results can be found in section 4.2. Section 5 concludes.

2 The Model

Information Structure

Uncertainty is described by two independent Brownian motions, $\{B(t,\omega), Z(t,\omega)\}$, for $0 \leq t \leq T$, defined on a complete probability space $(\Omega, \mathcal{F}, \mathbb{P})$. $\mathcal{B}_t$ is the augmented filtration generated by $\{B_t, Z_t\}$. The probability space fulfills the usual conditions. All processes I consider are appropriately adapted to $\mathcal{B}_t$, and, unless otherwise specified, all expectations are taken with respect to the probability measure $\mathbb{P}$.

I also introduce the processes Y and w, such that

$$\begin{aligned} dY_t &= \mu_t dt + \sigma_t dB_t \\ dw_t &= g(\mu_t)dt + \phi_t dZ_t \end{aligned}$$

where μ_t and $g(\mu_t)$ will be defined later. Y_0 is a constant and $w_0 = 0$. $\mathcal{Y}_t$ is the augmented filtration generated by $\{Y_t, w_t\}$. σ_t and ϕ_t are positive $\mathcal{Y}_t$-measurable elements of $\mathcal{L}_2$.[1]

It will turn out that $\mathcal{B}_t$ represents the agent's information set, while $\mathcal{Y}_t$ represents the principal's information set. This means that the agent can observe every element in the economy, including the path of $\{B_t, Z_t\}$, while the principal can only observe the path of $\{Y_t, w_t\}$.[2]

Opportunities

There is a risky project in the economy that pays a cumulative dividend Y_t over the interval $[0, t]$ and terminates at time T. The rights to the projects are owned by the principal, but he hires an agent to undertake the project. The agent exerts a control, μ_t, that determines the evolution of Y:

$$dY_t = \mu_t dt + \sigma_t dB_t \tag{1}$$

I will assume that the agent chooses $\mu_t \in M$, where M is a compact set in which the smallest value is 0. Moreover, μ must be an element in $\mathcal{L}_1$. Both discrete (e.g. $M = \{0, H\}$) and continuous (e.g. $M = [0, H]$) choice sets are allowed.

The agent also has access to outside investment opportunities or perks that pay a cumulative amount w_t over the interval $[0, t]$ with

$$dw_t = g(\mu_t)dt + \phi_t dZ_t \tag{2}$$

The function g represents the effect of the agent's effort on his outside activities. For example, g could represent a financial cost of effort: as he spends more time and effort on the principal's

[1] The spaces $\mathcal{L}_1$ and $\mathcal{L}_2$ are

$$\begin{aligned} \mathcal{L}_1 &= \left\{ X : \int_0^T |X_t| dt < \infty \text{ a.s.} \right\} \\ \mathcal{L}_2 &= \left\{ X : \int_0^T X_t^2 dt < \infty \text{ a.s.} \right\} \end{aligned}$$

[2] Throughout the paper, I use $\mathcal{Y}_t$ to denote the filtration generated by Y and w, rather than just Y. It is not important that the noise in Y and the noise in w are uncorrelated, only that their correlation is not 1. Keeping the correlation less than one will ensure that the principal cannot use the paths of Y and w to infer the drifts.

project, his outside activities and investments suffer.[3] Similarly, g could represent an outside gain from effort: as the agent exerts more effort, he may gain knowledge that he can use outside of his work for the principal. Alternately, w_t could represent the financial value of perks associated with the agent's position. Thus, I do not restrict g to be increasing or decreasing (or even continuous), but I do assume that $g(0) = 0$ and $g(\mu) \in \mathcal{L}_1$. In either case, however, the principal can observe the w process. The term $\phi_t dZ_t$ represents a random return on investments or perks, and it also acts to prevent the principal from using the agent's outside income to infer μ from the path of w.

In return for the agent's labor, the principal offers the agent a contract or sharing rule that pays the agent a cumulative amount S_t (salary) over the interval $[0, t]$ and an intermediate consumption process c. The principal cannot directly observe μ, but he can observe the history of Y and w. As a result, the principal offers payments at time t that depend on the entire path of $\{Y, w\}$ on $[0, t]$. This captures the principal's imperfect information about the agent's controls. More rigorously, the principal is restricted to offering a contract in which S_t and c_t are $\mathcal{Y}_t$-measurable and $\mathrm{E}[S_t]$ exists and is finite. I will assume that $c_t \in C$, where C is a compact set. Moreover, c must be an element in $\mathcal{L}_1$. Both discrete (e.g. $C = \{0, H\}$) and continuous (e.g. $C = [0, H]$) choice sets are allowed.

The two income processes for the agent, S and w, make up the agent's financial wealth process W:

$$dW_t = dS_t + rW_t dt + dw_t \tag{3}$$

where $W_0 = S_0$. $r \geq 0$ is a constant and rW_t represents the fact that any wealth the agent has is invested at a constant risk-free rate. It is important that the principal's observations of S and w allow him to observe W as well.

In this paper, I present two different versions of the principal-agent problem. The first is the version outlined above when $dw_t \neq 0$. In this model, the agent has some outside opportunities that bring in some noisy income but that income is affected when the agent exerts effort on the principal's project. The second version sets $dw_t = 0$ (so $g(\mu_t) = \phi_t = 0$) and represents a model in which the principal simply awards the agent a consumption stream

[3] Since ϕ_t is $\mathcal{Y}_t$-measurable, the problem is substantively unchanged if I say the agent's outside income is $\phi_t W_t dZ_t$ plus some drift, where W_t is the agent's total wealth. As we will see, the fact that the principal can observe w implies that the principal can observe the agent's total wealth as well. Some readers might find $\phi_t W_t dZ_t$ more intuitively appealing as a form of risky investment return.

over time with a lump-sum payment at the end of the economy.

Objectives – Agent

The agent's objective function is

$$\mathrm{E}\left[\int_0^T e^{-rt}u\left(c_t,\mu_t\right)dt+e^{-rT}v\left(W_T\right)\right] \tag{4}$$

where u and v are utility functions. As is standard, I assume that u and v are three times continuously differentiable.

This specification allows for two potential costs or benefits associated with the agent's effort. The first cost is a dis-utility of effort, captured by the argument μ_t in the function $u\left(c_t,\mu_t\right)$. This can represent any time-specific effort cost, and it includes models in which effort is additively separable (ex: $u(c_t,\mu_t)=h(c_t)-j(\mu_t)$) as well as models in which the agent pays a time specific, financial opportunity cost for his effort (ex: $u(c_t,\mu_t)=h(c_t-j(\mu_t))$).

The second potential cost or benefit from effort is the financial effect of effort $g(\mu_t)$, already mentioned, with $g(0)=0$ and g twice-continuously differentiable. It will turn out that the two different costs of effort can have very different economic consequences.

The agent has an outside opportunity that he values with a certainty equivalent utility of $\hat{U}$. The agent will only accept the principal's contract if

$$\max_{\mu}\mathrm{E}\left[\int_0^T e^{-rt}u\left(c_t,\mu_t\right)dt+e^{-rT}v\left(W_T\right)\right]\geq\hat{U} \tag{5}$$

The participation constraint is enforced only at time 0 (when the contract is signed), although I will show in the discussion of commitment that an analogous constraint can be enforced on the entire $[0,T]$ interval.

Assuming the agent accepts the contract, his problem is to find μ^* so that

$$\begin{aligned}\mu^* \quad &\in \quad \arg\max_{\mu}\ \mathrm{E}\left[\int_0^T e^{-rt}u\left(c_t,\mu_t\right)dt+e^{-rT}v\left(W_T\right)\right] \\ \text{s.t.}\quad &\text{(i)}\quad dY_t=\mu_t dt+\sigma_t dB_t \\ &\text{(ii)}\quad dW_t=dS_t+rW_t dt+g(\mu_t)dt+\phi_t dZ_t\end{aligned} \tag{6}$$

where S_t and c_t are understood to be functions of Y and w.

If μ^* solves the agent's problem for $\{S, c\}$, then I say that $\{S, c\}$ *implements* μ^* (equivalently, μ^* is *incentive compatible* with $\{S, c\}$).

Objectives – Principal

The principal's objective function[4] is

$$\mathrm{E}\left[Y_0 - S_0 + \int_0^T e^{-rt}\left(dY_t - dS_t\right)\right] \tag{7}$$

The principal's problem is to maximize his objective function subject to the constraints that the agent accepts the contract and that the agent behaves optimally. Thus, the principal's problem is to find $\{S^*, c^*\}$ so that[5]

$$\begin{aligned} \{S^*, c^*\} \quad &\in \quad \arg\max_{S,c} \mathrm{E}\left[Y_0 - S_0 + \int_0^T e^{-rt}\left(dY_t - dS_t - c_t dt\right)\right] \\ \text{s.t.} \quad (\text{i}) \quad & \mu[S,c] \text{ solves the agent's problem} \\ (\text{ii}) \quad & dY_t = \mu_t[S,c]dt + \sigma_t dB_t \\ (\text{iii}) \quad & dW_t = dS_t + rW_t dt + g(\mu_t[S,c])dt + \phi_t dZ_t \\ (\text{iv}) \quad & \mathrm{E}\left[\int_0^T e^{-rt} u\left(c_t, \mu_t\right)dt + e^{-rT} v\left(W_T\right)\right]\Bigg|_{\mu=\mu[S,c]} \geq \hat{U} \end{aligned} \tag{8}$$

Additional Assumptions

In addition, I require the problem to exhibit "local concavity" and a type of invertibility:

Assumption 1 [Concavity]: *For any positive real number p_1, real number p_2, and feasible value of c_t, it is the case that*

$$p_1\mu_t + u(c_t, \mu_t) + p_2 g(\mu_t)$$

has a unique maximum across μ_t for $\mu_t \in M$.

[4] The results about the form of the contract in section 3 do not require that the principal is risk neutral. The same results would hold if the principal's objective function were any function that was increasing in the principal's consumption. However, some of the results from section 4 do require risk neutrality.

[5] I assume that the principal has the choice to undertake the project or not. It may be the case that the project evolution and the various constraints make the principal worse off with the project than without – his certainty equivalent wealth from undertaking the project may be less than zero or fail to exist. In this case, I say that the principal rejects the project.

I call this local concavity because it is assessed state-by-state rather than for the problem as a whole.

Assumption 2 **[Invertibility]:***For any feasible processes μ, c, and value W_T, it is the case that an X exists so that*

$$\mathrm{E}\left[\int_0^T e^{-rt}u\left(c_t,\mu_t\right)dt + e^{-rT}v\left(W_T\right)|\mathcal{B}_t\right] = \int_0^t e^{-rs}u\left(c_s,\mu_s\right)ds + e^{-rt}v(X)$$

This amounts to assuming that no matter what the contract is, the agent has a certainty equivalent wealth. The assumption is easily met for common utility functions such as CARA and CRRA.

Equilibrium

An equilibrium consists of a contract $\{S^*, c^*\}$ and an implemented level of effort μ^*. μ^* must be adapted to the agent's information set $\mathcal{B}_t$ and solve the agent's problem (6). $\{S^*, c^*\}$ must be adapted to the principal's information set $\mathcal{Y}_t$ and solve the principal's problem (8).

3 The Agent's Problem

In this section I will describe the optimal contract as a function of the equilibrium level of effort. This will allow me to write the principal's problem of choosing a contract as simply a choice of what effort level to implement. Since a major point of this paper is finding an easy-to-apply solution method, I will present the method in the text and leave the rigorous details to the proofs in the appendix.

3.1 Contract Value and Uniqueness

The first step in identifying an optimal contract is to notice that there are many different salary processes that result in the same choices for the principal and for the agent. In fact, any two salary processes that deliver the same present value of payments over $[0, T]$ will result in the same choices by the agent; the principal will also be indifferent between them. The agent is indifferent because he can borrow and save at rate r, while the principal is risk neutral.

Definition 1 [**Equivalence**]: *Any two salary processes S and $\hat{S}$ are* equivalent *if*

$$S_0 + \int_0^T e^{-rt} dS_t = \hat{S}_0 + \int_0^T e^{-rt} d\hat{S}_t$$

almost surely.[6]

Even though the two contracts deliver different salary payments to the agent, the total value of the payments is the same. In that case, the two contracts result in the same choices and utilities for both the principal and the agent:

Proposition 1 [**Contract Equivalence**]: *Assume that the contracts $\{S, c\}$ and $\{\hat{S}, c\}$ have equivalent salary processes. Then the solution to the agent's problem is the same for both contracts. In addition, the principal and the agent achieve the same level of utility under the two contracts.*

If two salary processes give rise to the same value of W_T (almost surely across paths of $\{Y, w\}$), then the two salary processes are equivalent.

This proposition means that a valid strategy for finding an optimal contract is to find the value of W_T that results from the agent using his optimal controls, and then find an S_t that, combined with the agent's optimal controls and w_t, generates W_T. That S_t will be a salary process for which all optimal salary processes are equivalent.

Intuitively, knowing W_T reveals how much wealth the agent has at T, while knowing the agent's effort choices and w reveals how much the agent has acquired along the way. Putting those together reveals how much wealth the agent has received through the contract. Since the distribution of the discounted value of that wealth over time does not matter (equivalence of salary implies the same optimal control), I need only find one particular salary process that results in the given W_T. I go through this procedure in the next section.

3.2 The Contract Form

In solving the agent's problem, I am looking for all the salary processes (S_t) that induce a particular control choice by the agent. However, the principal *assumes* the agent behaves

[6]The equality is almost sure across paths of Y and w rather than B and Z because S_t is $\mathcal{Y}_t$-measurable rather than $\mathcal{B}_t$-measurable.

optimally (8i) when he sets S_t. I will use the fact that any optimal contract for the principal must in turn create an optimal control and a utility process for the agent. However, instead of using Girsanov's theorem and the weak method to solve for the agent's optimal control (as in Sannikov (2006) or Schättler and Sung (1993)), I will unravel the agent's utility using the principal's information set. I will then combine the principal's assessment of the agent's terminal wealth with his assumptions about the agent's control choice to find the payment process.

The first step in writing down the contract form is to define the agent's utility process using the principal's information set. The key observation is that knowing the agent's final wealth level W_T is the same as knowing the agent's terminal utility level $(e^{-rT}v(W_T))$. Since the agent's wealth level is $\mathcal{Y}_T$-measurable, we find the $\mathcal{Y}_t$-measurable evolution of the agent's utility function:

$$\mathcal{V}_t = \mathrm{E}\left[\int_0^T e^{-rt}u\left(c_t,\mu_t\right)dt + e^{-rT}v\left(W_T\right)|\mathcal{Y}_t,\mu_t=\mu_t^*\right] \tag{9a}$$

$$= \int_0^t e^{-rs}u\left(c_s,\mu_s^*\right)ds + e^{-rt}\mathcal{U}_t \tag{9b}$$

Here, $\mathcal{V}_t$ is the principal's estimate of the agent's total expected utility, while $\mathcal{U}_t$ is the total utility the principal expects the agent to receive in the future. Both expressions are evaluated using only the principal's information set – the information on which the actual payments are based. They are also evaluated under the constraint that the agent is actually choosing the level of effort and consumption that the principal wishes to implement, $\mu = \mu_t^* = \mu[S,c]$, because the principal assumes that the agent behaves optimally (8i) when determining payment. We will ignore, for the moment, the value the agent's optimal control actually takes.

Solving the principal's problem implies that the agent's participation constraint must bind exactly. Since the principal's and agent's information sets coincide at time 0, it must also be the case that $\mathcal{V}_0 = \mathcal{U}_0 = \hat{U}$.

The second step is to represent $\mathcal{V}_t$ in a tractable way. Because $\mathcal{V}_t$ is a martingale with respect to the information set $\mathcal{Y}_t$ (by the law of iterated expectations), we can use a Martingale Representation Theorem to show that there exist β_t and γ_t processes so that

$$d\mathcal{V}_t = e^{-rt}\beta_t\left(dY_t - \mu_t^*dt\right) + e^{-rt}\gamma_t\left(dw_t - g(\mu_t^*)dt\right) \tag{10}$$

where $dY_t - \mu_t^* dt$ and $dw_t - g(\mu_t^*)dt$ have zero drift (and dY_t and dw_t are what the principal can directly observe). Substituting (10) into (9b) and using Ito's lemma yields

$$d\mathcal{U}_t = r\mathcal{U}_t dt - u\left(c_t, \mu_t^*\right)dt + \beta_t\left(dY_t - \mu_t^* dt\right) + \gamma_t\left(dw_t - g(\mu_t^*)dt\right) \tag{11}$$

as an expression for the evolution of the agent's future expected utility. In equations (10) and (11), β_t and γ_t represent the principal's control over the volatility of the agent's utility.

The third step is to use the principal's estimate of the agent's utility ($\mathcal{U}_t$) to find the value of the principal's payment to the agent (S_t through W_T). Since $\mathcal{U}_T = v(W_T)$, I will examine the wealth process W_t with $\mathcal{U}_t = v(W_t)$ and $\hat{U} = v(W_0)$ and find the salary process S_t that supports it. This will result in one particular salary and wealth process pair that produces the required utility process and terminal wealth. However, proposition 1 shows that any other salary and wealth pair must be equivalent to the pair I find.[7]

To proceed, I substitute $\mathcal{U}_t = v(W_t)$ into (11), and using Ito's lemma, find that

$$\begin{aligned} & v'(W_t)dW_t + \frac{1}{2}v''(W_t)\left(\text{vol}(W_t)\right)^2 dt \\ = \; & rv(W_t)dt - u\left(c_t, \mu_t^*\right)dt + \beta_t\left(dY_t - \mu_t^* dt\right) + \gamma_t\left(dw_t - g(\mu_t^*)dt\right) \end{aligned} \tag{12}$$

where $\text{vol}(W_t)$ is the volatility of W_t. Matching diffusion terms shows that

$$\left(\text{vol}(W_t)\right)^2 = \left(\frac{\beta_t \sigma_t}{v'}\right)^2 + \left(\frac{\gamma_t \phi_t}{v'}\right)^2$$

I can now solve for the payment process from (12), substituting dW_t from the agent's budget constraint (3):

$$\begin{aligned} dS_t \; = \; & -\frac{u\left(c_t, \mu_t^*\right)}{v'(W_t)}dt + \frac{rv(W_t)}{v'(W_t)}dt - \frac{v''(W_t)}{v'(W_t)}\frac{1}{2}\left(\frac{\beta_t \sigma_t}{v'(W_t)}\right)^2 dt - \frac{v''(W_t)}{v'(W_t)}\frac{1}{2}\left(\frac{\gamma_t \phi_t}{v'(W_t)}\right)^2 dt \\ & -rW_t dt + \frac{\beta_t}{v'(W_t)}\left(dY_t - \mu_t^* dt\right) + \left(\frac{\gamma_t}{v'(W_t)} - 1\right)dw_t - \frac{\gamma_t}{v'(W_t)}g(\mu_t^*)dt \end{aligned}$$

[7] Critical to the procedure mentioned is that W_t is $\mathcal{Y}_t$-measurable (effectively observable to the principal through the S and w processes). An alternate procedure: Since the principal sets dS_t and observes dw_t, I could also say here that the principal sets dW_t through the budget constraint. Then, W_T is the time T value of a process W_t with $\mathcal{U}_t = v(W_t)$. This gives me a unique value for W_T in terms of dY_t, dw_t, β_t, γ_t, and μ_t^*, which define the evolution of $\mathcal{U}_t$. Using the budget constraint, I could find $S_0 + \int_0^T e^{-rt} dS_t$ uniquely and one value of dS_t that generates it. Then I would state that all other salary processes were equivalent to the one I found. This procedure involves slightly more steps than the discussion in the text but gives the same result.

with $v(S_0) = \hat{U}$.

Otherwise, if $dw_t = 0$, then $\phi_t = g(\mu_t) = 0$ and $\gamma_t = 0$ (since $\mathcal{V}_t$ can no longer depend on shocks to Z_t), and so

$$\begin{aligned} dS_t &= -\frac{u(c_t,\mu_t^*)}{v'(W_t)}dt + \frac{rv(W_t)}{v'(W_t)}dt - \frac{v''(W_t)}{v'(W_t)}\frac{1}{2}\left(\frac{\beta_t\sigma_t}{v'(W_t)}\right)^2 dt \\ &\quad -rW_t dt + c_t dt + \frac{\beta_t}{v'(W_t)}\left(dY_t - \mu_t^* dt\right) \end{aligned}$$

The payment processes (S_t) are expressed entirely as a function of things the principal observes (Y_t, W_t, and w_t), things the principal controls (β_t, γ_t, and c_t), and things the principal assumes ($\mu_t = \mu_t^*$). The principal controls the volatility of the agent's expected utility process (β_t and γ_t) by controlling the volatility of the agent's payment process (through the coefficients on dY_t and dw_t). After substituting $a(W_t) = -\frac{v''(W_t)}{v'(W_t)}$, I have a formal statement:

Theorem 2 [Contract Form]: *Assume the contract solves the principal's problem (8), implements μ^*, and $dw_t \neq 0$. Then the payment process is equivalent to the process S_t, with $\hat{U} = v(S_0)$ and*

$$\begin{aligned} dS_t &= -\frac{u(c_t,\mu_t^*)}{v'(W_t)}dt - dw_t - rW_t dt + r\frac{v(W_t)}{v'(W_t)}dt \\ &\quad + a(W_t)\frac{1}{2}\left(\frac{\beta(t,\mathcal{Y}_t)\sigma_t}{v'(W_t)}\right)^2 dt + a(W_t)\frac{1}{2}\left(\frac{\gamma(t,\mathcal{Y}_t)\phi_t}{v'(W_t)}\right)^2 dt \\ &\quad + \frac{\beta(t,\mathcal{Y}_t)}{v'(W_t)}\left(dY_t - \mu_t^* dt\right) + \frac{\gamma(t,\mathcal{Y}_t)}{v'(W_t)}\left(dw_t - g(\mu_t^*)dt\right) \end{aligned} \tag{13}$$

where $\beta(t,\mathcal{Y}_t)$ and $\gamma(t,\mathcal{Y}_t)$ are undetermined process (to be chosen by the principal) such that $e^{-rt}\beta_t\sigma_t$ and $e^{-rt}\gamma_t\phi_t$ are in $\mathcal{L}_2$ and $a(W_t) = -\frac{v''(W_t)}{v'(W_t)}$ is the agent's coefficient of absolute risk aversion.

If $dw_t = 0$, then the payment process is equivalent to the process S_t, with $\hat{U} = v(S_0)$ and

$$\begin{aligned} dS_t &= -\frac{u(c_t,\mu_t^*)}{v'(W_t)}dt - rW_t dt + r\frac{v(W_t)}{v'(W_t)}dt + a(W_t)\frac{1}{2}\left(\frac{\beta(t,\mathcal{Y}_t)\sigma_t}{v'(W_t)}\right)^2 dt \\ &\quad + \frac{\beta(t,\mathcal{Y}_t)}{v'(W_t)}\left(dY_t - \mu_t^* dt\right) \end{aligned} \tag{14}$$

Moreover, for both payment processes (13) and (14),

$$\mathrm{E}\left[\int_0^T e^{-rt}u\left(c_t,\mu_t\right)dt + e^{-rT}v\left(W_T\right)|\mathcal{B}_t,\ \mu=\mu^*\right] = \int_0^t e^{-rs}u\left(c_s,\mu_s^*\right)ds + e^{-rt}v(W_t) \quad (15)$$

The last equation (15) means that the agent's total financial wealth, W_t, is also the agent's certainty equivalent wealth under the particular contract given in equations 13 or 14. This fact is the reason I do not require any forward looking variables to form my contract representation and it is at the heart of many of the useful extensions that this method allows.

Discussion

Theorem 2 greatly constrains the set of potential optimal contracts. It shows that any optimal contract must be the outcome of a specific $\mathcal{Y}_t$-measurable process, S_t, characterized by the yet unknown volatility processes β_t and γ_t. Furthermore, the contract can be stated as a function of one state variable (W_t), which reflects the agent's financial wealth *and* his continuation value (15).[8] This is *not* a statement about the realized value of the payment the principal makes to the agent given that he uses the optimal control. Instead, it is a statement about the value of the payment the principal makes to the agent for *any* control the agent uses. One can see this result by noticing that the salary (13) is a function of both the optimal control (μ^* – known to the principal) and the control the agent actually uses (μ through dY_t). Theorem 2 shows that instead of studying all possible contracts, we can restrict our attention to contracts in which (13) holds. As a result, the agent's optimization decision will be easier to handle.

The payment process is made up of four separate but interpretable parts. The first part is re-payment of direct costs:

$$-\frac{u\left(c_t,\mu_t^*\right)}{v'(W_t)}dt - dw_t - rW_t dt$$

These terms reflect the fact that the principal must pay the agent in order to compensate him for his direct costs. However, when the agent is paid more intermediate utility or gains a benefit from his effort, then the principal can pay the agent less by those amounts. Similarly, when the agent is able to capture the interest benefit from his wealth, the principal

[8]In making this statement, I am using the fact that c, β and γ are set by the principal, and we will see in section 5 that the principal's problem also has only one state variable: $v(W_t)$.

can reduce the salary by a corresponding amount. The second part of the salary process is promise keeping

$$r\frac{v(W_t)}{v'(W_t)}dt$$

which comes from the fact that the agent's wealth has to increase to offset the fact that later consumption and wealth are discounted.

The final two parts of the salary are the insurance terms

$$a(W_t)\frac{1}{2}\left(\frac{\beta_t\sigma_t}{v'(W_t)}\right)^2 dt + a(W_t)\frac{1}{2}\left(\frac{\gamma_t\phi_t}{v'(W_t)}\right)^2 dt$$

and the incentive terms

$$\frac{\beta_t}{v'(W_t)}\left(dY_t - \mu_t^* dt\right) + \frac{\gamma_t}{v'(W_t)}\left(dw_t - g(\mu_t^*)dt\right)$$

The incentive terms represent the agent's gain from exerting effort from the project or from his outside opportunities. The β and γ processes measure the "slope" of the contract and the agent's exposure to his own choice of effort and to the two sources of idiosyncratic risk, B and Z. Because the agent is risk averse, his exposure to risk must be compensated in proportion to the amount of risk. Thus, the insurance terms represent the standard hidden action cost of implementing a steep contract and a high level of effort.

The key feature of the representation of the terminal consumption process in (13) is that it exactly compensates the agent for his costs – both direct costs from $u(c_t, \mu_t)$ and $g(\mu_t)$ and indirect costs from risk – *at the moment they are incurred*. This means that agent is effectively myopic. To see this, observe that equation (15) implies that, given S_t, the agent's future expected utility gain on $(t, T]$ is always zero. So, while the agent can always affect his current situation (W_t), he does not change his future opportunities at the optimal effort level. The agent's myopia means that one can calculate all the quantities in the contract directly – none is based on a forward looking variable, like certainty equivalent wealth. This represents a significant advance over previous work that requires calculations involving forward looking expectations, or that assumes local risk aversion ($a(\cdot)$) is the same as global risk aversion (CARA).

3.3 The Agent's Optimal Control

Section 3.2 provides the form of any optimal contract in terms of the effort level that contract implements. We have the optimal contract as a function of the "slope" of the contract ($\{\beta, \gamma\}$), the agent's optimal control (μ^*), the agent's actual control (μ), and the agent's intermediate consumption (c). However, we do not yet know how to make the contract *incentive compatible*; we need to know how the principal sets $\{\beta, \gamma, c\}$ in order to control the agent's choice of μ. So, to proceed, we must find the correspondence between contract slope and effort level. I will use dynamic programming techniques to solve this problem.

Given the contract in theorem 2, the agent's problem is to maximize his objective function (6). It will be clear that by constructing the payment process according to the procedure in section 3.2, the optimal control problem becomes easy to solve (despite its apparent length – many terms will cancel).

I will use the continuous time version of the Hamilton-Jacobi-Bellman equation in the text. I will then present a verification theorem that confirms the HJB result and is proved in the appendix. To continue, if $V(t, W_t)$ represents the agent's value function with $V(T, W_T) = e^{-rT}v(W_T)$ (the terminal condition), then μ^* is the optimal control and V is the value function when μ_t^* solves

$$0 = \max_{\mu_t} \left[e^{-rt}u(c_t, \mu_t)dt + \mathrm{E}\left[dV(t, W_t)\right]\right]$$

First, we must establish the evolution of the variable of interest, W_t. I can combine the budget constraint (3) with the form of the payment process (13) and perform some algebra to obtain

$$\begin{aligned} dW_t \;\; = \;\; & a(W_t)\frac{1}{2}\left(\frac{\beta_t\sigma_t}{v'(W_t)}\right)^2 dt + a(W_t)\frac{1}{2}\left(\frac{\gamma_t\phi_t}{v'(W_t)}\right)^2 dt - \frac{u(c_t, \mu_t^*)}{v'(W_t)}dt + r\frac{v(W_t)}{v'(W_t)}dt \\ & + \frac{\beta_t}{v'(W_t)}\left(\mu_t - \mu_t^*\right)dt + \frac{\gamma_t}{v'(W_t)}\left(g(\mu_t) - g(\mu_t^*)\right)dt \qquad (16) \\ & + \frac{\beta_t}{v'(W_t)}\sigma_t dB_t + \frac{\gamma_t}{v'(W_t)}\phi_t dZ_t \end{aligned}$$

Second, using Ito's lemma, the HJB equation is

$$
\begin{aligned}
0 \quad = \quad & \max_{\mu_t} \Bigg[e^{-rt} u(c_t, \mu_t) + \frac{\partial}{\partial t} V + V_{WW} \frac{1}{2} \left(\left(\frac{\beta_t \sigma_t}{v'(W_t)} \right)^2 + \left(\frac{\gamma_t \phi_t}{v'(W_t)} \right)^2 \right) \\
& V_W \Bigg(a(W_t) \frac{1}{2} \left(\frac{\beta_t \sigma_t}{v'(W_t)} \right)^2 + a(W_t) \frac{1}{2} \left(\frac{\gamma_t \phi_t}{v'(W_t)} \right)^2 - \frac{u(c_t, \mu_t^*)}{v'(W_t)} + r \frac{v(W_t)}{v'(W_t)} \\
& + \frac{\beta_t}{v'(W_t)} (\mu_t - \mu_t^*) + \frac{\gamma_t}{v'(W_t)} (g(\mu_t) - g(\mu_t^*)) \Bigg) \Bigg]
\end{aligned}
\tag{17}
$$

However, theorem 2 and the derivation in section 3.2 have already shown that $V(t, W_t) = e^{-rt}\mathcal{U}_t = e^{-rt} v(W_t)$. So, we will use $V = e^{-rt} v(W_t)$ as the candidate value function. Using that candidate, we see that the second order conditions in (17) are met by assumption, that the maximum is obtained at

$$\mu_t^* = \arg\max_{\mu_t} u(c_t, \mu_t) + \beta_t \mu_t + \gamma_t g(\mu_t)$$

and that the right hand side of (17) equals zero when $\mu_t = \mu_t^*$.

If $dw_t = 0$, then dW_t becomes

$$
\begin{aligned}
dW_t \quad = \quad & -\frac{u(c_t, \mu_t^*)}{v'(W_t)} dt + \frac{r v(W_t)}{v'(W_t)} dt + a(W_t) \frac{1}{2} \left(\frac{\beta_t \sigma_t}{v'(W_t)} \right)^2 dt \\
& + \frac{\beta_t}{v'(W_t)} (\mu_t - \mu_t^*) dt + \frac{\beta_t}{v'(W_t)} \sigma_t dB_t
\end{aligned}
\tag{18}
$$

and the HJB equation becomes

$$
\begin{aligned}
0 \quad = \quad & \max_{\mu_t} \Bigg[e^{-rt} u(c_t, \mu_t) + \frac{\partial}{\partial t} V + V_{WW} \frac{1}{2} \left(\frac{\beta_t \sigma_t}{v'(W_t)} \right)^2 \\
& V_W \left(a(W_t) \frac{1}{2} \left(\frac{\beta_t \sigma_t}{v'(W_t)} \right)^2 - \frac{u(c_t, \mu_t^*)}{v'(W_t)} + r \frac{v(W_t)}{v'(W_t)} + \frac{\beta_t}{v'(W_t)} (\mu_t - \mu_t^*) \right) \Bigg]
\end{aligned}
\tag{19}
$$

Again, we use the candidate value function $e^{-rt} v(W_t)$ to see that second order conditions in (19) are met by assumption, that the maximum is obtained at

$$\mu_t^* = \arg\max_{\mu_t} u(c_t, \mu_t) + \beta_t \mu_t$$

and that the right hand side of (19) equals zero when $\mu_t = \mu_t^*$.

A verification theorem proves that these are indeed the incentive compatible controls:

Theorem 3 *If* $dw_t \neq 0$ *and the principal offers the agent a contract equivalent to the one in (13), then the contract implements* μ^* *if and only if*

$$\mu_t^* = \arg\max_{\mu_t} u(c_t, \mu_t) + \beta(t, \mathcal{Y}_t)\mu_t + \gamma(t, \mathcal{Y}_t)g(\mu_t) \tag{20}$$

If $dw_t = 0$ *and the principal offers the agent a contract equivalent to the one in (14), then the contract implements* μ^* *if and only if*

$$\mu_t^* = \arg\max_{\mu_t} u(c_t, \mu_t) + \beta(t, \mathcal{Y}_t)\mu_t \tag{21}$$

While theorem 3 is phrased as an optimal control problem, it is really a condition on the unknown variables from section 3.2 (β and γ) that make μ^* the solution to the agent's problem. Equations 20 and 21 are the *incentive compatibility constraints* on the contract form given in theorem 2.

Theorem 3 also shows that under the process S_t in (13), the agent is myopic in his choices: In choosing his optimal effort choice, the agent maximizes only the sum of direct costs and benefits (20). In that equation, the first term ($u(c_t, \mu_t)$) represents the agent's loss of utility due to effort. The second and third terms represent the agent's gain and loss to the agent from the incentive parts of the contract. β_t is the agent's "slope" with respect to effort spent on the dividend, and γ_t is the agent's "slope" with respect to the effect of effort on the agent's outside income. This myopic-ness makes sense given that the agent is compensated for his costs as he goes along – since additional effort does not increase his future opportunities (see section 3.2), the agent is motivated at time t by the direct time t costs and benefits.

In addition, theorem 3 shows that "local concavity" for the agent (assumption 1) implies global concavity for the agent. Thus we do not require any of the more complicated assumptions that are necessary in discrete time (see, for example, Jewitt (1988)). The reason is that when decisions are continuous and the outcome of any one decision is infinitesimal, then local concavity everywhere assures us of global concavity.

3.4 Implementation

The results in the previous two sections allow me to make if and only if statements about form of the contract the principal chooses:

Theorem 4 [Implementation]: *Assume a given contract* $\{S, c\}$ *solves the principal's problem.*

If $dw_t \neq 0$*, the contract implements* μ^* *if and only if* S_t *is equivalent to (13) for which* $\beta(t, \mathcal{Y}_t)$*,* $\gamma(t, \mathcal{Y}_t)$*, and* μ_t^* *are related as given in (20).*

If $dw_t = 0$*, the contract implements* μ^* *if and only if* S_t *is equivalent to (14) for which* $\beta(t, \mathcal{Y}_t)$*, and* μ_t^* *are related as given in (21).*

The combination of the salary process (13 or 14) and the incentive compatibility constraints (20 or 21) gives the optimal payment process (S_t) as a function only of wealth (W_t) and the contract's "slope" ($\{\beta_t, \gamma_t\}$).

Discussion

The myopic representation – the conversion of a dynamic problem into a repeated static problem – is very powerful. Given *any* contract, theorem 4 determines the controls the contract will implement in terms of a representation of that contract. So the theorem shows how to construct a contract based on the desired controls and says that all contracts that implement those controls must have the same representation.

The myopic representation also explains why so many things that one might think are important for the agent's opportunity set do not, in fact, matter at all. Consider, for example, stochastic volatility, potentially non-Markovian controls, stochastic or deterministic changes in the cost function, etc. Each of these impacts the opportunity set of the agent, but the agent ignores these potential future effects. More importantly, *the agent appears to ignore his own impact on his own future opportunities.* For example, it may be the case that σ_t decreases as the dividend increases. One might think that a risk-averse agent might then exert extra effort to enter the region with less risk. In the myopic representation, the agent is dynamically compensated for changes in σ_t (or any other variable) so as to keep his control choice unaffected. In this example, the agent knows that if σ_t decreases, his compensation will decrease to reflect the lower levels of risk (the last term of the contract). Because the agent is exactly compensated for the risk he incurs, he has no incentive to strategically change his control to lower his exposure to risk. The resulting contract is balanced so as to keep the agent myopic.

Quitting, Firing, and Commitment

The results contained in theorem 4 actually require much less commitment than is assumed in the model. In fact, the model can be generalized to allow the principal to fire the agent or for the agent to quit at any time before T. In a model with early termination at τ, one only requires that the agent is able to walk away with his wealth W_τ at the time of termination.

If termination occurs at time τ, assume the agent's objective function is

$$\mathrm{E}\left[\int_0^\tau e^{-rt}u(c_t,\mu_t)dt + e^{-r\tau}v(W_\tau)\right]$$

Equation 15 shows that under the total commitment optimal contract given in (13), the agent receives zero future expected utility gain on $(\tau, T]$ regardless of whether he quits at τ, is fired at τ, or works until T. Even if the agent cannot commit to keep working, he is still indifferent in equilibrium between continuing or stopping work. Thus, the agent will never choose to quit regardless of his level of commitment.

If it is the principal that cannot commit, rather than the agent, the solution is just as simple. Whenever the principal desires to terminate employment, he can simply stop paying the agent. This does not make the agent any worse off and ensures that the principal will always receive a positive expected utility from a contract. In other words, as long as the agent can use his wealth when terminated early (the agent's objective function contains $e^{-r\tau}v(W_\tau)$), the solution for the optimal contract given in theorem 4 will be valid for any type of employment termination.

Time Consistency

In parts of the macroeconomics literature (e.g. the literature started by Kydland and Prescott (1977)) policy makers will often try to implement time-inconsistent policies. One can imagine that a principal might wish to choose a time inconsistent contract – a contract, for example, that uses some punishment as an incentive to the agent, but a punishment that the principal does not wish to carry through in the event that the contract calls for it. In my setting, these types of contracts require that the agent not have rational expectations:

the agent must be fooled by the principal into thinking that the punishment will actually be acted on. If the agent has rational expectations instead, then any time-inconsistent contract is equivalent to the time consistent version that would arise after the re-negotiation, and I can say that I am dealing with the time consistent version.

If, however, the agent does not have rational expectations with regards to the principal's actions, then the model outlined in section 2 can be relaxed to allow for time-inconsistent contracts. To see how, observe that the condition on the agent's expected utility (15) implies that at any time t, the agent is indifferent to continuing or terminating the contract given in (13). Moreover, it is easy for the principal to offer the agent a contract for which the agent receives zero expected utility: use the contract form given by theorem 4, but set $\hat{U} = 0$. If the principal wishes to re-negotiate the contract at any time, the agent will agree to any new contract with the form in (13) that sets $S_0 = \hat{U} \geq 0$ because such a re-negotiation or new contract results in a payment to the agent of S_0, making the agent strictly better off. Thus, one only needs to add a condition in the principal's problem that says he will re-negotiate the contract whenever

$$\frac{\partial}{\partial W_t} \mathrm{E}\left[\int_t^T e^{-r(s-t)}\left(dY_s - dS_s\right)|\mathcal{Y}_t,\ \mu = \mu^*\right] > 1$$

This condition says that the value the principal obtains from transferring a unit of wealth to the agent exceeds the value from keeping it.[9]

Memory

It is commonly stated that in dynamic contracts the rewards to the agent from a positive realization of dY_t are spread over time and that the contract itself exhibits memory. From the representation in (13) and (20) we can see exactly how those results are generated. First, I consider the spreading of rewards over time. Equation 20 shows that the agent makes his effort choice at time t based only on the time t values of β, γ, and c. Effort does not directly depend at all on anything that happens before or after t. So consumption smoothing is entirely driven by the principal's maximization decision and his risk-sharing with the agent, not by what is required to motivate the agent. The specific value of c_t the principal chooses is demonstrated in section 4.3 for separable utility.

[9]Since the agent's future expected utility is always zero (15), this condition is equivalent to saying that the total surplus created by transferring a unit of wealth to the agent is greater than one.

The second issue is memory – to what extent the agent's opportunities change as a result of past successes and failures. The contract clearly exhibits memory: dS_t depends on W_t, which is the agent's financial wealth and summarizes the agent's continuation value. However, the direction is not clear – the derivative of $\mathrm{E}\left[dS_t\right]$ with respect to W_t cannot be signed without additional assumptions. Moreover, the value of the principal's and agent's shared discount rate r can change the apparent sign of the derivative.[10]

4 The Principal's Problem

Because I have shown how to implement a desired control choice, the principal's problem becomes one of finding the desired control to implement. In this sense, it is exactly the same as the agent's problem: the principal takes the contract form to be exogenous and simply makes an optimization decision. This is often called the principal's relaxed problem (following Rogerson (1985a)), and being able to frame problems in this way despite not using the first-order approach is a major advantage of the continuous time formulation.

The principal's problem, as originally stated, is to find

$$\begin{aligned}
\{S^*, c^*\} \quad &\in \quad \arg\max_{S,c} \mathrm{E}\left[Y_0 - S_0 + \int_0^T e^{-rt}\left(dY_t - dS_t - c_t dt\right)\right] \\
\text{s.t.} \quad &\text{(i)} \quad \mu[S,c] \text{ solves the agent's problem} \\
&\text{(ii)} \quad dY_t = \mu_t[S,c]dt + \sigma_t dB_t \\
&\text{(iii)} \quad dW_t = dS_t + rW_t dt + g(\mu_t[S,c])dt + \phi_t dZ_t \\
&\text{(iv)} \quad \mathrm{E}\left[\int_0^T e^{-rt} u\left(c_t, \mu_t\right) dt + e^{-rT} v\left(W_T\right)\right]\Bigg|_{\mu=\mu[S,c]} \geq \hat{U}
\end{aligned}$$

We have shown in section 3 that if the principal's problem has a solution, then the payment process must be equivalent to the one given in theorem 2, where β, γ, and μ are related by theorem 3. The principal's constraints (8i) mean that the principal conducts his optimization under the assumption that the agent behaves optimally, and so we can set

[10] Since v can be positive or negative, when r is large, then the impact of rv on both consumption and utility can be large and positive or large and negative.

$\mu = \mu^*$ along the path of S_t:

$$\begin{aligned} dS_t &= -\frac{u\left(c_t, \mu_t^*\right)}{v'(W_t)}dt + \frac{rv(W_t)}{v'(W_t)}dt + a(W_t)\frac{1}{2}\left(\frac{\beta_t \sigma_t}{v'(W_t)}\right)^2 dt + a(W_t)\frac{1}{2}\left(\frac{\gamma_t \phi_t}{v'(W_t)}\right)^2 dt \\ &\quad -rW_t dt - g(\mu_t^*)dt + \frac{\beta_t}{v'(W_t)}\sigma_t dB_t + \left(\frac{\gamma_t}{v'(W_t)} - 1\right)\phi_t dZ_t \end{aligned} \tag{22}$$

when $dw_t \neq 0$, or

$$dS_t = -\frac{u\left(c_t, \mu_t^*\right)}{v'(W_t)}dt + \frac{rv(W_t)}{v'(W_t)}dt + a(W_t)\frac{1}{2}\left(\frac{\beta_t \sigma_t}{v'(W_t)}\right)^2 dt - rW_t dt + \frac{\beta_t}{v'(W_t)}\sigma_t dB_t \tag{23}$$

when $dw_t = 0$.

However, it is often easier to consider the agent's total remaining utility as the state variable. Re-arranging and integrating the agent's budget constraint (3) (as in the proof of proposition 1) shows

$$S_0 + \int_0^T e^{-rt} dS_t = e^{-rT} W_T - \int_0^T e^{-rt}\left(g(\mu_t)dt + \phi_t dZ_t\right) \tag{24}$$

If we also define $v_t = v(W_t)$, then the original statement of the evolution of the agent's utility (11) and the $\mu = \mu^*$ substitution yields

$$dv_t = rv_t dt - u(c_t, \mu_t^*)dt + \beta_t \sigma_t dB_t + \gamma_t \phi_t dZ_t \qquad \text{if } dw_t \neq 0 \tag{25a}$$

$$dv_t = rv_t dt - u(c_t, \mu_t^*)dt + \beta_t \sigma_t dB_t \qquad \text{if } dw_t = 0 \tag{25b}$$

where $v_0 = \hat{U}$. Using (24) and the fact that $W_t = v^{-1}(v_t)$, we can re-write the principal's objective function as

$$\begin{aligned} & \mathrm{E}\left[Y_0 - S_0 + \int_0^T e^{-rt}\left(dY_t - dS_t - c_t dt\right)\right] \\ = \; & \mathrm{E}\left[Y_0 + \int_0^T e^{-rt}\left(\mu_t - c_t + g(\mu_t)\right)dt - e^{-rT} v^{-1}(v_T)\right] \end{aligned}$$

These simplifications allows me to re-write the principal's problem:

Theorem 5 [The Principal's Problem]: *When $dw_t \neq 0$, a contract $\{S, c\}$ is a solution* *to the principal's problem (8) if and only if*

$$\begin{aligned}\{\beta^*, \gamma^*, c^*\} \quad &\in \quad \arg\max_{\{\beta,\gamma,c\}} \mathrm{E}\left[Y_0 + \int_0^T e^{-rt}\left(\mu_t^* - c_t + g(\mu_t^*)\right)dt - e^{-rT}v^{-1}(v_T)\right] \qquad (26)\\ \text{s.t.} \quad &\text{(i)} \quad dv_t = rv_t dt - u(c_t, \mu_t^*)dt + \beta_t\sigma_t dB_t + \gamma_t\phi_t dZ_t \text{ with } \hat{U} = v_0\\ &\text{(ii)} \quad \mu_t^* = \arg\max_{\mu_t} u(c_t, \mu_t) + \beta_t\mu_t + \gamma_t g(\mu_t)\end{aligned}$$

where v^{-1} is defined as the inverse of the utility function v, and S_t is equivalent to (13) with $\{\beta, \gamma, c\}$ given in (26).

When $dw_t = 0$, a contract $\{S, c\}$ is a solution to the principal's problem (8) if and only if

$$\begin{aligned}\{\beta^*, c^*\} \quad &\in \quad \arg\max_{\{\beta,c\}} \mathrm{E}\left[Y_0 + \int_0^T e^{-rt}\left(\mu_t^* - c_t\right)dt - e^{-rT}v^{-1}(v_T)\right] \qquad (27)\\ \text{s.t.} \quad &\text{(i)} \quad dv_t = rv_t dt - u(c_t, \mu_t^*)dt + \beta_t\sigma_t dB_t \text{ with } \hat{U} = v_0\\ &\text{(ii)} \quad \mu_t^* = \arg\max_{\mu_t} u(c_t, \mu_t) + \beta_t\mu_t\end{aligned}$$

where S_t is equivalent to (14) with $\{\beta, c\}$ given in (27).

Theorem 5 re-writes the principal-agent contracting problem as a conceptually simple (and numerically simple) optimization problem. There are existence theorems for optimal controls in such problems (Ex: Fleming and Rishel (1975) and Yong and Zhou (1999)), and the numerical methods for solving (26) and (27) are well known.

I will now proceed to describe some of the properties of optimal contracts, including how to reduce the number of choice variables in (26) from three to two.

4.1 Minimum Variance Controls

The two principal's problems in theorem 5 can be simplified further by realizing the principal will always use controls that minimize the variance of the agent's utility process. A more volatile utility process implies a more volatile payment process, and, since the agent is risk averse, the extra risk has to be compensated:

Proposition 6 [**Minimum Variance Controls**]: *If $dw_t \neq 0$, and $\{S, c\}$ solves the principal's problem (26) and implements μ^*, then $\{\beta_t^*, \gamma_t^*\}$ will jointly minimize $\beta_t^2\sigma_t^2 + \gamma_t^2\phi_t^2$ under the constraint (26ii).*

If $dw_t = 0$ and $\{S, c\}$ solves the principal's problem (27) and implements μ^*, then β_t^* will be the minimum value of β_t such that the constraint (27ii) is met.

This theorem is useful because there is frequently not a one-to-one mapping from the agent's control choice, μ, to the set of contracts that implement's it. For example, if $dw_t = 0$ and $\mu_t \in M = \{0, H\}$, then any level of β_t such that $\beta_t \geq \frac{1}{H}(u(c_t, 0) - u(c_t, H))$ will serve to implement $\mu_t^* = H$. The theorem tell us that the minimum $\beta_t(\mu_t = H) = \frac{1}{H}(u(c_t, 0) - u(c_t, H))$ will be used.

There is a similar situation when $dw_t \neq 0$ because for any desired μ_t^*, there is more than one choice of $\{\beta_t, \gamma_t\}$ that will implement it. Consider again the case in which M is continuous, so that (26ii) becomes $\beta_t = -u_\mu(c_t, \mu_t^*) - \gamma_t g'(\mu_t^*)$. Then total volatility is minimized when

$$\gamma_t^* = \frac{\sigma_t^2 g'(\mu_t^*)}{\phi_t^2}\beta_t^* \tag{28a}$$

$$\beta_t^* = -\frac{u_\mu(c_t, \mu_t^*)\phi_t^2}{g'(\mu_t^*)^2\sigma_t^2 + \phi_t^2} \tag{28b}$$

These calculations also show that when μ_t^* is in the interior of a feasible interval, $\beta_t^* > 0$. $\gamma_t^* \leq 0$ if and only if g is a cost function ($g' \leq 0$).

The sign of γ_t^* comes from the incentives (indirect costs) inherent in insuring the agent against idiosyncratic shocks (dZ_t) to the agent's wealth. Examining the salary process (22), the principal's direct cost from γ_t comes only from the agent's risk aversion $a(W_t)\frac{1}{2}\left(\frac{\gamma_t\phi_t}{v'(W_t)}\right)^2$. That value depends only the the absolute value of γ_t, not it's sign. However, the principal also has to incentivize the agent to choose the desired level of μ_t (26ii). When g is a cost function, lowering γ_t reduces the level of β_t required to motivate the agent. When the principal over-subsidizes the agent's private wealth cost to choosing an effort level ($\gamma_t < 0$), the principal can use lesser *direct* incentives (β_t) to motivate the agent and decrease the total amount of risk to which the agent is subject.

Conversely, when g is a benefit function – the agent's effort level complements his outside activities – then the principal offers less than full insurance ($\gamma_t > 0$) against outside shocks.

This effect is because the agent's outside wealth can substitute for direct incentives based on the project itself (β_t).

In addition, the principal is not made unambiguously worse off by the presence of g as a cost function, compared to the agent having no outside opportunities. To see this, consider the case in which σ_t is very large and ϕ_t is very small. When $dw_t = 0$, the principal must use β_t to provide incentives, and he must pay a huge insurance premium to do so $(a(W_t)\frac{1}{2}\left(\frac{\beta_t\sigma_t}{v'(W_t)}\right)^2)$. By contrast, when $dw_t \neq 0$, (28) shows that the principal chooses β_t very close to zero, and motivates the agent primarily through the choice of γ_t. The insurance cost of this is much less because ϕ_t is so much smaller than σ_t. As long as g itself is not too large, the principal is made better off with an additional means of motivating the agent.

Theorem 6 also reduces the set of potential controls the principal can use to minimum variance controls. Since $\{\beta_t^*, \gamma_t^*\}$ can be written as functions of μ_t^* and c_t^*, the principal's problem can be stated as a maximization over what actinos to implement $\{\mu^*, c^*\}$ rather than $\{\beta, \gamma, c\}$. Formally, this means that the contract can now be written entirely as a function of the agent's wealth (W_t or $v_t = v(W_t)$) and the agent's actions ($\{\mu_t^*, c_t^*\}$).

4.2 Real Options

The most general result regarding optimal effort choice is one regarding real options. I consider the case in which $\mu_t \in M = \{0, H\}$, so the principal is deciding at any point in time whether the agent should be "on" ($\mu_t = H$) or "off" ($\mu_t = 0$).[11] I will also set σ_t and ϕ_t equal to constants so as to remove an extra source of variation. One can think of this as a decision about when to hire and fire managers. As one might expect, the principal will more often turn the agent "off" when the direct and indirect costs of employing the agent are high. However, the principal will systematically employ the agent at a (short term) loss near the beginning of the contract while he will not do so at the end. The principal's reasoning can be understood in a real-options framework: If the agent is "on", his wealth is fluctuating and may increase or decrease, and so the principal's cost of employment will fluctuate as

[11] In this section, I will write as as setting $\mu_t = 0$ is equivalent to terminating the agent. This is true in equilibrium in the sense that theorem 7 shows that once the principal sets $\mu_t = 0$, he does not later set $\mu_t = H$. The discussion after theorem 2 shows that this can be equivalent to the principal terminating the agent with no further payment. The restriction that there is only one feasible positive effort level (H) is not necessary for the results – theorem 7 would be true for any set M that contained 0 – but the restriction does simplify the interpretation.

well. If the principal is losing money on the agent now, he may still profitably employ the agent in the future if the agent's wealth changes so as to decrease the cost of employment.

Thus, the principal has a real option – which he exercises by terminating the agent – in which the underlying variable is the agent's cost of employment, determined by the agent's current level of wealth. The possibility that the agent's wealth may change in an advantageous way is only present when the agent is "on" (the volatility of W_t is zero when the agent is "off"), and so the principal will keep the option open – keep the agent "on" – even when doing so results in a short term loss. As the contract reaches its termination, the future ability of the principal to profitably employ an agent declines, and so the principal is less willing to bear a short term loss near T.

Theorem 7 [Real Options]: *Assume the optimal contract $\{S^*, c^*\}$ implements μ^* and solves the principal's problem. There exists an r^* so that if $0 \leq r \leq r^*$ and, for some history (path of Y and w on $[0,t]$), $\mu_t^* = 0$, then for all $s \in [t, T]$ it is the case that $\mu_s^* = 0$. If, in addition, $c_t \in C$, a continuum, then c_s^* is constant as well.*

Theorem 7 proves both parts of the real options intuition. First, it is a direct proof of the fact that once the agent has been turned "off", he is not re-started. To see the second – that the agent's is employed at a loss more often when t is far from T – I need a definition of the "short term cost" of employment over $[t, t+h]$:

$$\begin{aligned}\pi_{t,h} &= \max_c \mathrm{E}\left[\int_t^T e^{-rs}\left(\mu_s^* ds - dS_s - c_s^* ds\right) | \mathcal{B}_t,\ \mu^*_{[t,t+h]} = H,\ \mu^*_{[t+h,T]} = 0\right] \\ &\quad - \max_c \mathrm{E}\left[\int_t^T e^{-rs}\left(\mu_s^* ds - dS_s - c_s^* ds\right) | \mathcal{B}_t,\ \mu^*_{[t,T]} = 0\right]\end{aligned}$$

This represents the expected amount the principal gains by setting $\mu^* = H$ on $[t, t+h]$ and then $\mu^* = 0$ on $[t+h, T]$ over what the principal expects to gain by setting $\mu^* = 0$ on the entire $[t, T]$. Since termination is permanent, if $\pi_{t,h} > 0$, the principal will choose to set $\mu_t^* = H$: if the principal's short term gain from employing the agent is positive, his full gain from employing the agent must also be positive.

However, if $\pi_{t,h} < 0$, it is not true that the principal always sets $\mu_t = 0$ (termination). The reason is that $\pi_{t,h}$ does not represent the full value of employing the agent on $[t, t+h]$ because it does not take the maximizing value of μ after time $t+h$. Thus, the principal's full gain from employment might be greater than the short term gain, $\pi_{t,h}$, and so the principal may choose to employ the agent with a short term loss.

Thus, termination only takes place when the principal expects that his direct value of continued employment is negative. Since, short term and full values of employment must converge as $t \to T$ (since at T there is no distinction between the two), it must be that, at least weakly, the principal is employing the agent at a loss more often earlier in the contract.

In figure 1, I have simulated an economy for which $u(c_t, \mu_t) = -\frac{1}{2}c_t^{-2} - \frac{1}{2}\mu_t^2$, $v(W_T) = -\frac{1}{2}W_T^{-2}$ and $dw_t = 0$. The shaded areas represent parts of the state space in which $\mu_t^* = H$. The three dotted lines represent paths of v_t for which $dB_t = dZ_t = 0$ (which gives accurate values for v_t if $\mu_t^* = 0$ since then v_t is deterministic). As proved, the paths of v_t are such that once they leave the $\mu_t = H$ area, they do not return. Moreover, the $\mu_t = H$ region shrinks as t approaches T.

In figure 2, I have simulated a more exotic economy with a non-monotonic cost of employing the agent and discrete consumption choices. While there is a stopping time at which μ_t goes from H to 0, it cannot be represented as a simple function of either t or W_t. The key point is that while the curve that separates $\mu_t = H$ from $\mu_t = 0$ bends backwards, the paths of v_t are such that they never re-enter the H region after leaving. The stopping time curve bends backwards because in this particular specification, the principal's cost of employing the agent is not monotonic in W_t (I analyze this behavior in more detail in section 4.4). However, in this plot the $\mu_t = H$ region again shrinks as t approaches T.

The intuitions about real-options hold if r is near zero, but they can fail if r moves too far from zero. The reason is that the discount rate may mean that there are more gains from trade at the end of the contract than at the beginning. For an example, consider the case where r is large. Then the evolution of dv_t (26i) implies

$$v_T = e^{rT} v_0 + e^{rT} \int_0^T e^{-rs} \left(-u(c_s, \mu_s) ds + \beta_s \sigma_s dB_s + \gamma_s \phi_s dZ_s \right)$$

which has a time t "instantaneous variance" of

$$e^{2r(T-t)} \left(\beta_t^2 \sigma_t^2 + \gamma_t^2 \phi_t^2 \right)$$

The point is that the instantaneous variance is decreasing over t for a fixed β and γ. The promise-keeping constraint (the rv_t term in the evolution of v_t) means that the agent's

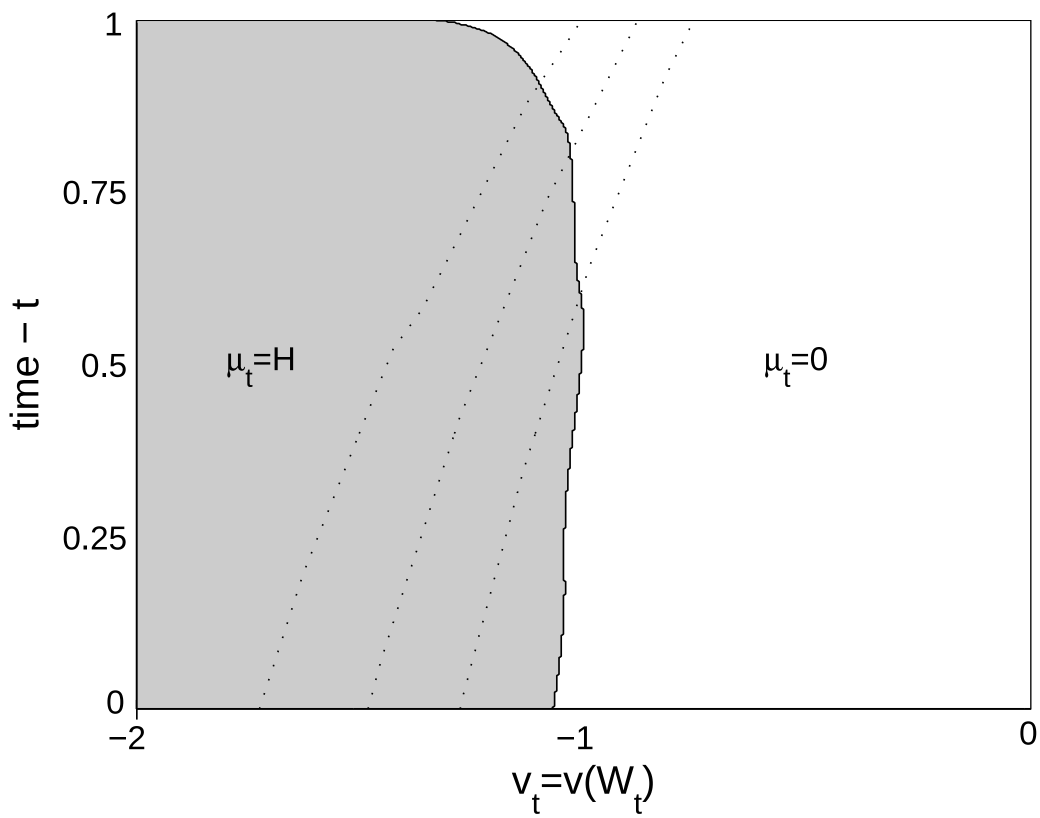

Figure 1: This plot was generated by setting $u = -\frac{1}{2}c_t^{-2} - \frac{1}{2}\mu_t^2$, $v = -\frac{1}{2}W_T^{-2}$, $\sigma = 5$, $T = 1$, and $dw_t = 0$. The principal was restricted to setting $\mu_t \in \{0, 1\}$ and $c_t \in [0, K]$ for K large. The shaded region represents the portion of (t, v) space in which $\mu_t^* = H = 1$. The three dotted lines represent paths of v_t for which $dB_t = dZ_t = 0$. The principal sets $\mu_t = H$ at the beginning of the economy and continues to keep the agent "on" until some optimal stopping time is reached, at which point the principal sets $\mu_t^* = 0$.

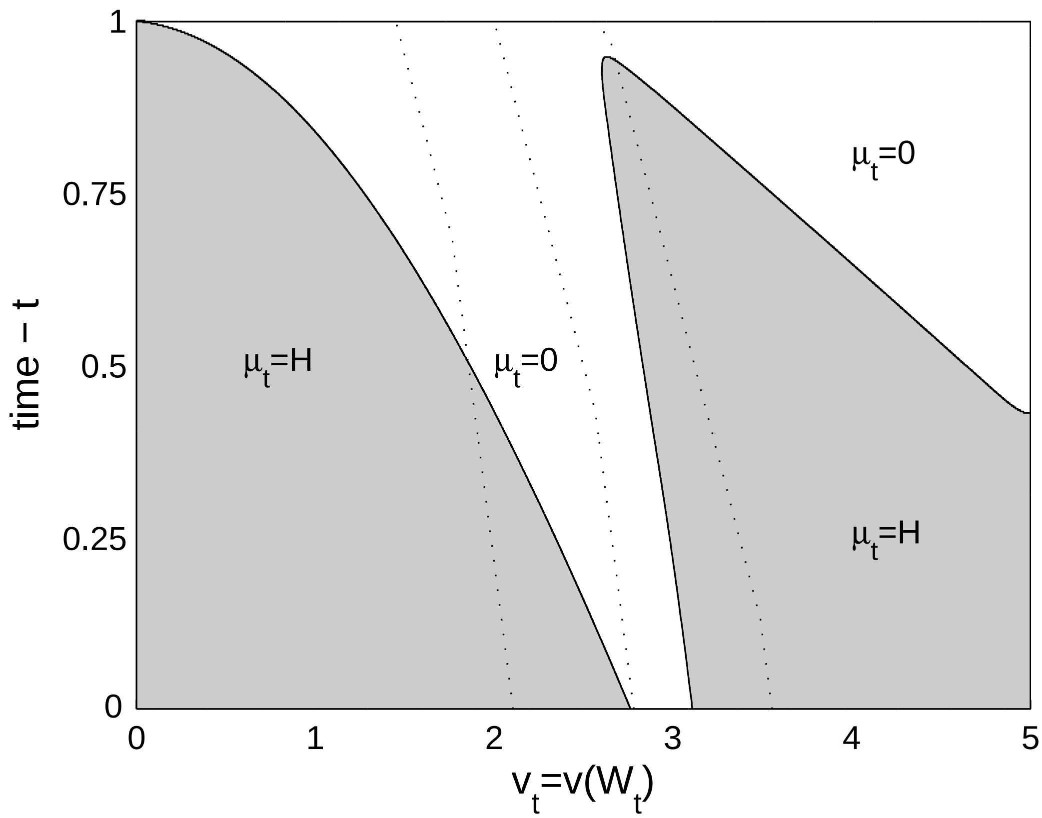

Figure 2: This plot was generated by setting $u = 2\sqrt{c_t} - \frac{1}{2}\mu_t^2$, $v = 2\sqrt{W_T}$, $\sigma = 5$, $T = 1$, and $dw_t = 0$. The principal was restricted to setting $\mu_t \in \{0, 1\}$ and $c_t \in \{\frac{1}{4}, 4\}$. The shaded region represents the portion of (t, v) space in which $\mu_t^* = H = 1$. The three dotted lines represent paths of v_t for which $dB_t = dZ_t = 0$. The principal sets $\mu_t = H$ at the beginning of the economy and continues to keep the agent "on" until some optimal stopping time is reached, at which point the principal sets $\mu_t^* = 0$.

promised utility must grow over time to keep up with the discount rate. Thus, any random innovations in B at time $t < T$ have a larger impact on v_T than similar innovations in B and time T.

Since the instantaneous variance of v_t is decreasing over time, the cost the principal pays to compensate the agent for it (from $\mathrm{E}\left[e^{-rT}v^{-1}(v_T)\right]$) also decreases over time. Thus, it can be less costly to implement effort at the end of the contract when the discount rate is sufficiently large.

4.3 Intermediate Consumption

While the principal-agent model described in section 2 is extremely general, parts of it cannot be easily solved without making functional form assumptions. In fact, if one is willing to assume a separation in the utility function, then the agent's inverse marginal utility is a martingale.

A standard cost-minimization argument, like that of Rogerson (1985b), shows us that

Theorem 8 [Intermediate Consumption]*: If the contract* $\{S, c\}$ *solves the principal's problem (26 or 27),* $u(c_t, \mu_t) = u(c_t) - j(\mu_t)$*, and* $c_t \in C$ *is a continuum, then it must be the case that*

$$\frac{1}{u'(c_t)} = \mathrm{E}\left[\frac{1}{u'(c_s)}|\mathcal{B}_t\right] = \mathrm{E}\left[\frac{1}{v'(W_T)}|\mathcal{B}_t\right] \tag{29}$$

for all $t \leq s \leq T$*.*

Moreover, if u *is not separable in* c_t *and* μ_t*, but* $\mu^* = 0$ *on* $[t, s]$*, then* c *is constant on* $[t, s]$ *as well.*

Theorem 8 shows that, as a result of cost minimization, the principal will keep the agent's inverse marginal utility as a martingale. However, this result requires separability of consumption and effort because of how the agent forms his optimal control. In the general model, the agent sets μ_t^* as a function of both the slope of the contract (β_t and γ_t) and his current level of consumption (c_t) (20). The principal's cost of implementing a particular level of effort will change over time and directly as a function of c_t. Thus, the principal's choice of c_t and μ_t are not separate, and so a cost minimization argument would have to take into account the complicated interactions between these variables. On the other hand, when the agent's utility function is separable, the agent's choice of μ_t^* now depends only on β_t and γ_t, but not on c_t. Under that specification, a cost minimization procedure can be done without regard to the level of effort being implemented.

4.4 Comparative Statics

There are two comparative statics that one might like to be true in a principal-agent problem: first, one might want the principal's optimal choice of the agent's consumption (c_t) to be monotonically related to agent's continuation value (W_t). This intuition comes from a deterministic optimization problem – if the principal owes the agent a certain amount of utility, then the concavity of the agent's utility function implies that the principal should try to pay the debt off over time, with increasing intermediate payments for increasing amounts of debt. Second, one might want the principal's optimal choice of effort to be monotonically related to the agent's level of expected utility. If the agent has a standard DARA utility function (decreasing absolute risk aversion, like power utility), then more wealth means that the agent is less risk averse and cheaper to insure. The insurance/effort tradeoff would then indicate that the principal should implement a higher level of effort when the agent is wealthier.

Neither of the above intuitions is correct, and both intuitions fail because of subtle interactions in dynamic problems. I will demonstrate this using generic examples when $dw_t = 0$.

Consumption

For the first (consumption) intuition, I will examine the principal's choice at time T. At T, the economy is being terminated and so the principal has no forward looking inputs into his decision making. To simplify the calculations, let us consider a variation of the financial cost of effort setting applied to intermediate consumption: $r = 0$, $\mu_t \in M = \{0, H\}$, $\sigma = 1$, $u(c, \mu) = u(c - \mu)$, and $c_t \in C = [0, K]$ for K large. I will also assume that $a(\cdot)$ is very small. Then, as given in (27), the principal acts so as to maximize

$$\max_{\mu_t, c_t} \mathrm{E}_t \left[\mu_t dt - dS_t | t = T \right]$$

while setting β_T according to theorem 6:

$$\beta_T(c_T, \mu_T = 0) = 0 \text{ and } \beta_T(c_T, \mu_T = H) = \tfrac{1}{H}\left(u(c_T) - u(c_T - H)\right)$$

So, the principal acts to maximize

$$\mu_T - c_T + \frac{u(c_T - \mu_T)}{v'(W_T)} - a(W_T)\frac{1}{2}\left(\frac{\beta_T(c_T, \mu_T)}{v'(W_T)}\right)^2$$

If a is small, then principal will set $v'(W_T) \approx u'(c_T - \mu_T)$, equalizing marginal payoffs. If we fix μ_T and vary W_T, then one can see that the principal increases the intermediate payment to the agent (c_t) as the required terminal payment increases, matching the deterministic trade-off intuition. However, μ_T is not constant: the principal will set $\mu_T = H$ if $H + \frac{u(c_T - H)}{v'(W_T)} - \frac{u(c_T)}{v'(W_T)} \gtrapprox 0$, which will be the case when W_T is small (and so v' is large). Consider the region around the point at which μ_T goes from H to 0. As W_T rises to the critical value, c_T increases with it, and at the critical value μ_T jumps down. The principal, however, is setting $v'(W_T) \approx u'(c_T - \mu_T)$, so when μ_T jumps down, c_T must jump down as well.

So, c_T is increasing in W_T throughout most of it's range, but it is non-monotonic in the area in which the principal adjusts incentives. It is the interaction of consumption and effort that drives the effect: When not motivating the agent to choose high effort, the agent is on a higher level of intermediate utility and so there is diminished return from giving the agent intermediate consumption. Instead, the principal will choose to pay the agent with terminal consumption. In other words, the agent's marginal utility of consumption does decline as W_T increases, but the interaction of consumption and effort means that the declining marginal utility does not imply a consumption increase. This example is illustrated using power utility in figure 3a.

Effort

The second potential comparative statics result is that effort is monotonic in the agent's wealth. Again, I will examine this intuition at time T so as to simplify the principal's problem. To simplify the calculations, I will assume that $r = 0$, $u(c, \mu) = c^{\frac{3}{4}} - \frac{1}{2}\mu^2$, $v(W) = W^{\frac{3}{4}}$, $\sigma = 1$, $\mu_t \in M = [0, K]$, and $c_t \in C = [0, K]$ for K very large. Then, at time T the principal sets $\beta_T = \mu_T$ and acts to maximize

$$\mu_T - c_T + \frac{4}{3}W_T^{\frac{1}{4}}(c_T^{\frac{3}{4}} - \frac{1}{2}\mu_T^2) - \frac{2}{9}W_T^{-\frac{1}{2}}\mu_T^2$$

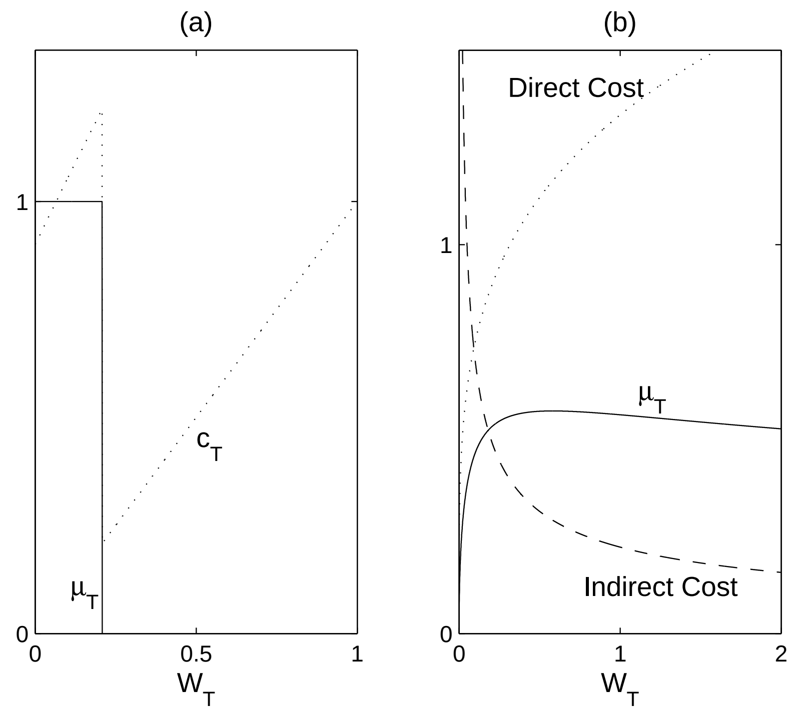

Figure 3: (a) Optimal controls at time T for the principal when $u = -\frac{1}{2}(c - .9\mu)^{-2}$, $\sigma = 2$, $v = -\frac{1}{2}W^{-2}$, $r = 0$, $\mu_t \in M = \{0, 1\}$, and $c_t \in C = [0, K]$ for K very large. (b) Optimal controls at time T for the principal when $u = c^{\frac{3}{4}} - \frac{1}{2}\mu^2$, $\sigma = 1$, $r = 0$, $v = W^{\frac{3}{4}}$, $\mu_t \in [0, K]$, and $c_t \in C = [0, K]$ for K very large. The direct cost is $\frac{4}{3}W_T^{\frac{1}{4}}$ and the insurance cost is $\frac{2}{9}W_T^{-\frac{1}{2}}$

by setting

$$\mu_T^* = \min\left[\frac{1}{\frac{4}{3}W_T^{\frac{1}{4}} + \frac{4}{9}W_T^{-\frac{1}{2}}}, K\right]$$

which is not monotonic. The key here is that the two different costs that the principal faces – the direct cost of effort and the cost of insuring the agent – work in opposite directions. The direct cost of effort (from $-\frac{u}{v'} \to -\frac{2}{3}W_T^{\frac{1}{4}}\mu_T^2$) is highest when the agent has a high level of expected utility (high W_T) because the agent's utility function is relatively flat there and the cost of effort is measured in units of utility rather than consumption. The cost of insuring the agent, however, is lowest when expected utility is high because the agent has decreasing absolute risk aversion (from $a\frac{1}{2}\frac{\beta^2}{v'^2} \to -\frac{2}{9}W_T^{-\frac{1}{2}}\mu_T^2$). Those two costs add so the agent is most profitably employed when he has a moderate level of wealth. So, it is the interaction of insurance costs and direct costs that provide for non-monotonicity of μ in W. This effect is illustrated in figure 3b.

5 Conclusion

The methods I use to derive the optimal contract and policy are applicable to more general problems than the one studied in this paper. In particular, one can add state variables, even correlated with the Y_t process, without difficulty. My formulation's ability to address additional state variables stems directly from the fact that I have not used the weak method (Girsanov's theorem) to determine the optimal contract.

For example, there may be an economic variable that enters into the agent's utility function directly: if we specify the dividend returns process to be geometric, as $dY_t = \mu_t Y_t dt + \sigma_t Y_t dB_t$, then the agent's utility function might read $u(c_t, \frac{\mu_t}{Y_t})$ or something similar. A similar state variable X_t, uncorrelated with Y_t might be added in the same way. More generally, any state variable can be added to the agent's objective function (including cost function) or the dividend's evolution, and those variables will simply be carried through to the optimal contract. In the geometric example, we can simply place $u(c_t, \frac{\mu_t}{Y_t})$ into the evolution of S_t.

The exception to this rule is variables that change the principal's inference problem. For example, if the process A_t is observable by the principal and correlated with B_t or Z_t, then the innovations in A_t will enter the contract directly. Instead of a contract of the form $dS_t = \alpha_t dt + \beta_t dY_t + \gamma_t dW_t$, one will obtain $dS_t = \alpha_t dt + \beta_t dY_t + \gamma_t dW_t + \delta_t dA_t$. This is derived from the fact that, following the discussion in section 3.2, $\mathcal{V}$ will be represented as a $\{\mathcal{Y}_t, \mathcal{A}_t\}$ martingale, rather than just a $\mathcal{Y}_t$ martingale. However, the optimal contract can be re-derived by simply making the appropriate substitutions into the equations in section 3.

A Proofs

Proof of Proposition 1. Re-arrange the agent's budget constraint (3), multiply by e^{-rt}, and integrate to obtain

$$\begin{aligned} dW_t - rW_t dt &= dS_t + g(\mu_t)dt + \phi_t dZ_t \\ e^{-rT}W_T - W_0 &= \int_0^T e^{-rt} dS_t + \int_0^T e^{-rt}\left(g(\mu_t)dt + \phi_t dZ_t\right) \\ e^{-rT}W_T &= S_0 + \int_0^T e^{-rt} dS_t + \int_0^T e^{-rt}\left(g(\mu_t)dt + \phi_t dZ_t\right) \end{aligned} \tag{30}$$

where the last line follows from the fact that $W_0 = S_0$. So, for any two contracts $\{S, c\}$ and $\{\hat{S}, c\}$ for which $S_0 + \int_0^T e^{-rt} dS_t = \hat{S}_0 + \int_0^T e^{-rt} d\hat{S}_t$, the same choice of μ_t leads to the same value of W_T. Thus, the agent's feasible set is the same under both contracts, and both contracts must result in the same choices and utilities. With μ unchanged, the principal's payoff of $Y_0 - S_0 + \int_0^T e^{-rt} dY_t - \int_0^T e^{-rt} dS_t$ achieves the same value from both contracts.

For the second statement, assume by contradiction that the salaries are not equivalent. Since $W_T = e^{rT}\left(S_0 + \int_0^T e^{-rt} dS_t\right) + e^{rT}\int_0^T e^{-rt} dw_t$, then there exists a set of paths of $\{Y, w\}$ of positive measure for which the value of W_T under the two salaries must be different. Thus, if W_T is the same almost surely across paths of $\{Y, w\}$, then the two salary processes must be equivalent. This completes the proof. ■

Proof of Theorem 2. Very little is needed to make the discussion in section 3.2 rigorous.

Because a solution to the principal's problem is assumed to exist in the statement of the theorem, constraint (8i) implies that there exists a solution to the agent's problem, which I label μ^*. Since a solution to the agent's problem exists, so does the agent's expected utility under the optimal control. Since the agent's problem is taken with respect to $\mathcal{B}_t$, the agent's expected utility must also exist with respect to $\mathcal{Y}_t$ if $\mu = \mu^*$ (since taking a value of μ as given makes the two information sets the same). I label the agent's conditional $\mathcal{Y}_t$-measurable expected utility with $\mu = \mu^*$ as $\mathcal{V}_t$ in (9).

The martingale representation theorem of Davis and Varaiya (1973), suitably updated in Revuz and Yor (2005) shows that $d\mathcal{V}_t$ can be represented as in (10) for some $e^{-rt}\beta_t\sigma_t$ and $e^{-rt}\gamma_t\phi_t$ which are members of $\mathcal{L}_2$. The agent's participation constraint (8iv) implies that $\mathcal{V}_0 = \hat{U}$.

Using the martingale representation of $\mathcal{V}_t$ (10) and the definition of $\mathcal{V}_t$ (9), I have a value for W_T as a function of the agent's optimal control μ^*, the path of w_t, and the contract $\{S, c\}$:

$$\begin{aligned}\mathcal{V}_T &= \hat{U} + \int_0^T e^{-rt}\beta_t\sigma_t\left(dY_t - \mu_t^* dt\right) + \int_0^T e^{-rt}\gamma_t\phi_t\left(dw_t - g(\mu_t^*)dt\right) \\ &= \int_0^T e^{-rt}u(c_t, \mu_t^*)dt + v(W_T)\end{aligned}$$

where W_t is $\mathcal{Y}_t$-measurable. The discussion in section 3.2 does the algebra to invert this relationship to find a pair of processes $\{S_t, W_t\}$ for which the terminal value of W_t is equal to the required W_T. The evolution of S_t then takes this W_t as an argument. Proposition 1 shows that the optimal contract must be equivalent to the found S_t.

The last result (15) is a re-statement of the definitions of $\mathcal{V}_t$ and $\mathcal{U}_t = v(W_t)$. If μ is known to equal μ^*, then the conditioning the expectation on $\mathcal{B}_t$ takes the same value as conditioning on $\mathcal{Y}_t$. This completes the proof. ■

Proof of Theorem 3. I will solve for the agent's optimal control using a standard dynamic programming verification theorem. This proof is an adapted version of a similar proof that appears in Vayanos and Wang (2006). A canonical version can be found in Fleming and Rishel (1975).

I prove both the $dw_t \neq 0$ ($\phi_t \neq 0$) and the $dw_t = 0$ ($\phi_t = g(\mu_t) = 0$) results together using the $\{S, c\}$ given. The result is the same for any equivalent $\{S, c\}$ by proposition 1. Define the variable $\hat{\mathcal{V}}_t$ so that

$$\hat{\mathcal{V}}_t = \int_0^t e^{-rs}u(c_s, \mu_s)ds + e^{-rt}v(W_t) \tag{31}$$

for some general μ process. Here, W_t denotes the terminal wealth process for the control μ (16 or 18), and $\hat{\mathcal{V}}_T$ is the agent's final realized utility. Observe that the Hamilton-Jacobi-Bellman equations (17 and 19) can be re-written as

$$0 = \max_{\mu} \mathrm{E}\left[d\hat{\mathcal{V}}_t | \mathcal{B}_t\right] \tag{32}$$

Since the right hand side of (32) achieves the maximum at zero when $\mu_t = \mu_t^*$ (20 and

21), the drift of $\hat{\mathcal{V}}_t$ is less than or equal to zero for any μ_t. Thus,

$$\hat{\mathcal{V}}_T \leq \hat{\mathcal{V}}_t + \int_t^T e^{-rs}\beta_s\sigma_s dB_s + \int_t^T e^{-rs}\gamma_s\phi_s dZ_s$$

Theorem 2 showed that $e^{-rt}\beta_t\sigma_t \in \mathcal{L}_2$ and $e^{-rt}\gamma_t\phi_t \in \mathcal{L}_2$ (integrability), so we can take expectations:

$$\hat{\mathcal{V}}_t \geq \mathrm{E}\left[\hat{\mathcal{V}}_T|\mathcal{B}_t\right] \tag{33}$$

This shows that $\hat{\mathcal{V}}_t$ is an upper bound on the agent's expected utility at time t.

Now, we repeat equations (31) and (33) for μ_t^*. Since the μ^* solves the maximization in (32) with the right hand side equal to zero, the drift of $\hat{\mathcal{V}}_t$ is zero for $\mu = \mu^*$ and

$$\hat{\mathcal{V}}_t = \mathrm{E}\left[\hat{\mathcal{V}}_T|\mathcal{B}_t,\, \mu = \mu^*\right] \tag{34}$$

This shows that the upper bound on the agent's utility is realized when $\mu = \mu^*$, meaning μ^* is the (unique, up to a set of measure zero, because the solution to the HJB equation was unique) optimal control. This completes the proof. ■

Proof of Theorem 4. Theorem 2 shows that if a contract $\{S, c\}$ implements either μ^*, then the contract form (13 or 14) must hold (up to equivalence). Theorem 3 shows that if the contract form (13 or 14) holds (up to equivalence), then the optimality condition (20 or 21) is true if and only if $\{S, c\}$ implements μ^*. Together, these imply that $\{S, c\}$ implements μ^* if and only if the contract form (13 or 14) and the optimality condition (20 or 21) both hold. This completes the proof. ■

Proof of Theorem 5. Theorem 4 and the discussion in the text above the statement of theorem 5 are sufficient to show that a solution to the principal's original problem (8) is also a solution to the principal's revised problem (26 or 27).

For the converse: First, the integrability assumptions (μ_t and c_t are members of time-independent compact sets and additional assumption 2) are enough to show that under any feasible contract, the agent's expected utility process (9) exists. Then, the feasible set in the principal's revised problem (26 or 27) is the same as the feasible set for the principal's original problem (8) because contracts in the revised problem are those that give rise to the stated utility process. Since the objective is the same in both problems, the optimum must be the same as well. This completes the proof. ■

Proof of Proposition 6. Begin with the principal's relaxed problem (26). Because v is concave, $-v^{-1}$ is also concave. Then Jensen's Inequality shows that the principal's utility is decreasing in the variance of the utility process: $\beta_t^2\sigma_t^2 + \gamma_t^2\phi_t^2$. The statement of the theorem follows. This completes the proof. ■

Proof of Theorem 7. I will prove the result for the $dw_t \neq 0$ model. The $dw_t = 0$ result follows similarly. This proof will use the results of theorem 8, even though that theorem appears later in the paper.

Set $r = 0$ and consider any $\{\bar{\beta}, \bar{\gamma}, \bar{c}\}$ that implement $\bar{\mu}_t^* = 0$ for some positive measure of time (δ) along almost all the paths of B and Z, starting at some point (t', v').[12] Theorem 6 shows that if $\bar{\mu}_t^* = 0$, then, since the contract is optimal, $\bar{\beta}_t = \bar{\gamma}_t = 0$, and v_t follows a deterministic path on $[t', t' + \delta]$. Label the path of v for these controls as $\bar{v}$.

Consider an alternate set of controls $\{\hat{\beta}, \hat{\gamma}, \hat{c}\}$ which are defined so that $\{\hat{\beta}, \hat{\gamma}, \hat{c}\}(t, v) = \{\bar{\beta}, \bar{\gamma}, \bar{c}\}(t + \delta, v + \bar{v}_{t'+\delta} - \bar{v}_{t'})$ on $[t', T - \delta]$. (The value of the controls on $[T - \delta, T]$ will be given later.) This definition means that $d\hat{v}_t$ will have the same distribution as $d\bar{v}_{t+\delta}$.[13]

Forward induction shows that for all $t \in [t', T - \delta]$, it is the case that $\hat{v}_t \overset{d}{=} \bar{v}_{t+\delta} + \bar{v}_{t'+\delta} - \bar{v}_{t'}$. In particular, $\hat{v}_{T-\delta} \overset{d}{=} \bar{v}_T + \bar{v}_{t'+\delta} - \bar{v}_{t'}$. To show that I can construct a set of controls superior to $\{\bar{\beta}, \bar{\gamma}, \bar{c}\}$, I need only show that I can construct controls that will do "better" on $[T-\delta, T]$ than setting $\hat{\mu}^*_{[T-\delta,T]} = 0$, as $\{\bar{\beta}, \bar{\gamma}, \bar{c}\}$ did on $[t', t'+\delta]$. The results of theorem 6 show that $\bar{c}_{[t',t'+\delta]}$ is constant (since $\bar{\mu}_{[t',t'+\delta]} = 0$). So, set $\hat{c}_{[T-\delta,T]} = \bar{c}_{[t',t'+\delta]}$. Then, choose $\{\hat{\beta}, \hat{\gamma}\}$ so as to set $\mu_t = H$ if and only if $\mathrm{E}\left[H - dS_t(\mu_t = H)\right] > 0$. Since $\{\bar{\beta}, \bar{\gamma}, \bar{c}\}$ set $\mu_{[t',t'+\delta]} = 0$ unconditionally, it is clear from the principal's objective function ($\mathrm{E}\left[Y_0 - S_0 + \int_0^T e^{-rt}\left(dY_t - dS_t\right)\right]$), and the fact that the values of c are the same on the relevant intervals, that $\{\hat{\beta}, \hat{\gamma}, \hat{c}\}$ is at least weakly preferred to $\{\bar{\beta}, \bar{\gamma}, \bar{c}\}$.

To show that $\{\hat{\beta}, \hat{\gamma}, \hat{c}\}$ is strictly preferred to $\{\bar{\beta}, \bar{\gamma}, \bar{c}\}$, observe that unless $\bar{\mu}_t = \bar{\beta}_t = \bar{\gamma}_t = 0$ everywhere, the probability that $\{\hat{\beta}, \hat{\gamma}, \hat{c}\}$ sets $\mu_t = H$ on a set of t with measure greater

[12] The point (t', v') is sufficient to describe the history because the principal can always achieve the optimum with markovian controls, and t and v are the only state variables. For a proof of this, see Davis and Varaiya (1973) or Fleming and Rishel (1975).

[13] This does not violate the adapted requirement of the controls because the time t value of the controls is assessed using time t state variables (t, v) and a non-random quantity (change in $\bar{v}$ on $[t', t' + \delta]$).

than zero also exceeds zero.[14] Thus, $\{\hat{\beta}, \hat{\gamma}, \hat{c}\}$ is strictly preferred because it results in a positive probability of a gain and no probability of a loss relative to $\{\bar{\beta}, \bar{\gamma}, \bar{c}\}$. The result on r^* follows from the continuous nature of the sample paths and the fact that the principal's controls are markovian.

To show that $\mu^*_{[t,t'']} = 0$ implies $c_{[t,t'']}$ is constant, apply the results of theorem 6. This completes the proof. ■

Proof of Theorem 8. I will prove the result for the case in which $dw_t \neq 0$. The $dw_t = 0$ result can be proved using the same argument.

Assume that the contract $\{S, c\}$ implements μ^* and solves the principal's problem (26). Then (26ii) implies that the separability of the agent's utility function means that β_t and γ_t will not depend on c, only μ^*. Some algebra and (26i) show that this contract delivers the agent a terminal utility

$$v_T = e^{rT} v_0 + \int_0^T e^{-r(t-T)} \left(u(c_t) - j(\mu^*_t)\right) dt + \int_0^T e^{-r(t-T)} \left(\beta^*_t \sigma_t dB_t + \gamma^*_t \phi_t dZ_t\right)$$

Now, consider an alternate contract with the same values of β and γ, so that it implements the same value of μ^*. However, it uses a different level of consumption, $\hat{c}_t$. $\hat{c}_t$ is chosen so that $u(\hat{c}_t) = u(c_t) + e^{r(t-T)}\alpha$ on $t \in [t', t' + \epsilon)$, $u(\hat{c}_t) = u(c_t) - e^{r(t-T)}\alpha$ on $[t'', t'' + \epsilon)$, and $c_t = \hat{c}_t$ elsewhere. t' and t'' are chosen so that the two time intervals are separate and lie inside $[0, T]$. Then, v_T is the same under the old and new contracts.

Thus, we need only examine the principal's objective function and pick out the terms that depend on $\hat{c}_t$:

$$\mathrm{E}\left[-\int_0^T e^{-rt} \hat{c}_t dt\right]$$

Since we have assumed that the original contract (with $\alpha = 0$) is optimal, then $\alpha = 0$ must

[14]If $\bar{\mu}_t = \bar{\beta}_t = \bar{\gamma}_t = 0$ everywhere for all possible $\{\bar{\beta}, \bar{\gamma}, \bar{c}\}$ controls, then the result of the theorem follows trivially.

minimize

$$
\begin{aligned}
& \mathrm{E}\left[\int_0^T e^{-rt}\hat{c}_t dt|\mathcal{B}_\tau\right] \\
= & \mathrm{E}\left[\int_{t'}^{t'+\epsilon} e^{-rt}u^{-1}\left(u(c_t)+e^{r(t-T)}\alpha\right)dt|\mathcal{B}_\tau\right] \\
& +\mathrm{E}\left[\int_{t''}^{t''+\epsilon} e^{-rt}u^{-1}\left(u(c_t)-e^{r(t-T)}\alpha\right)dt|\mathcal{B}_\tau\right]
\end{aligned}
$$

for all t', t'', ϵ, and $\tau \leq t', t''$. A necessary condition for this is that for any t' and t''

$$\frac{1}{u'(c^*_{t'})} = \mathrm{E}\left[\frac{1}{u'(c^*_{t''})}|\mathcal{B}_{t'}\right]$$

If u is non-separable in c_t and μ_t, the same argument can be used to show that

$$\frac{1}{u_c(c^*_{t'}, \mu^*_{t'})} = \mathrm{E}\left[\frac{1}{u_c(c^*_{t''}, \mu^*_{t''})}|\mathcal{B}_{t'}\right]$$

as long as μ is zero on $[t', t'']$ (since then β_t and γ_t are zero and v_t is deterministic). In that case, the above condition shows that c is almost surely constant on $[t', t'']$ as well.

For time T, notice that the principal has no forward looking concerns, and so he must choose c_T so as to minimize $\mathrm{E}\left[dS_T|\mathcal{B}_T\right]$ and so $u'(c_T) = v'(W_T)$. This completes the proof.
■

References

Cadenillas, A., J. Cvitanić, and F. Zapatero (2005): "Dynamic Principal-Agent Problems with Perfect Information," Working Paper, SSRN.

Cvitanić, J., and J. Zhang (2006): "Optimal Compensation with Adverse Selection and Dynamic Actions," Working Paper, SSRN.

Davis, M. H. A., and P. Varaiya (1973): "Dynamic Programming Conditions for Partially Observable Stochastic Systems," *SIAM Journal on Control*, 11(2), 226–261.

DeMarzo, P. M., and Y. Sannikov (2004): "A Continuous-Time Agency Model of Optimal Contracting and Capital Structure," Working Paper, Stanford University.

Fleming, W. H., and R. W. Rishel (1975): *Deterministic and Stochastic Optimal Control*, vol. 1 of *Applications of Mathematics*. Springer-Verlag, New York, first edn.

Holmstrom, B., and P. Milgrom (1987): "Aggregation and Linearity in the Provision of Intertemporal Incentives," *Econometrica*, 55(2), 303–328.

Jewitt, I. (1988): "The First Order Approach to Principal-Agent Problems," *Econometrica*, 56(5), 1177–1190.

Kydland, F. E., and E. C. Prescott (1977): "Rules Rather than Discretion: the Inconsistency of Optimal Plans," *Journal of Political Economy*, 85(3), 473–492.

Marino, A. M., and J. Zábojník (2006): "Work-Related Perks, Agency Problems, and Optimal Incentive Contracts," Working Paper, University of Southern Calfornia.

Ou-Yang, H. (2003): "Optimal Contracts in a Continuous-Time Delegated Portfolio Management Problem," *Review of Financial Studies*, 16(1), 173–208.

Phelan, C., and R. M. Townsend (1991): "Computing Multi-period, Information-Constrained Optima," *Review of Economic Studies*, 58(5), 853–881.

Revuz, D., and M. Yor (2005): *Continuous Martingales and Brownian Motion*, vol. 293 of *A Series of Comprehensive Studies in Mathematics*. Springer-Verlag, Berlin, third edn.

Rogerson, W. P. (1985a): "The First Order Approach to Principal-Agent Problems," *Econometrica*, 53(6), 1357–1368.

——— (1985b): "Repeated Moral Hazard," *Econometrica*, 53(1), 69–76.

Sannikov, Y. (2006): "A Continuous Time Version of the Principal-Agent Problem," Working Paper, Stanford University.

Schättler, H., and J. Sung (1993): "The First Order Approach to the Continuous-Time Principal-Agent Problem with Exponential Utility," *Journal of Economic Theory*, 61(2), 331–371.

Spear, S. E., and S. Srivastava (1987): "On Repeated Moral Hazard with Discounting," *Review of Economic Studies*, 54(4), 599–617.

Sung, J. (1995): "Linearity with Project Selection and Controllable Diffusion Rate in Continuous-Time Principal-Agent Problems," *The RAND Journal of Economics*, 26(4), 720–743.

Vayanos, D., and J. Wang (2006): "Capital Market Theory," Manuscript in preparation.

Williams, N. (2004): "On Dynamic Principal-Agent Problems in Continuous Time," Working Paper, Princeton University.

Yong, J., and X. Y. Zhou (1999): *Stochastic Controls: Hamiltonian Systems and HJB Equations*, vol. 43 of *Applications of Mathematics*. Springer-Verlag, New York, first edn.