Disagreement and Learning in a Dynamic Contracting Model*

Tobias Adrian
Federal Reserve Bank of New York

Mark M. Westerfield
University of Southern California

March, 2007

Abstract

We present a dynamic contracting model in which the principal and agent disagree about the resolution of uncertainty, and we illustrate the contract design in an application with Bayesian learning. The disagreement creates gains from trade that the principal realizes by transferring payment to states that the agent considers relatively more likely, which changes incentives. The principal's value function is convex in the underlying belief differences because the more optimistic the agent relative to the principal, the sharper the incentives and the lower the agent's required compensation. In our dynamic setting, the interaction between incentive provision and learning creates an intertemporal source of "disagreement risk" that changes second-best risk sharing. Under risk-neutrality, "selling the firm" to the agent does not implement the first-best because it precludes state-contingent trades.

JEL Classifications: D0, D8, G0

Keywords: Dynamic Contracts, Heterogeneous Beliefs, Hidden Action, Principal-Agent, Continuous Time

---

*Adrian can be reached at Capital Markets Research, Federal Reserve Bank of New York, 33 Liberty Street, New York, NY 10045; tobias.adrian@ny.frb.org; http://nyfedeconomists.org/adrian/. Westerfield can be reached at Department of Finance and Business Economics, University of Southern California, Hoffman Hall-701, MC-1427, 701 Exposition Blvd., Los Angeles, CA 90089-1427; mwesterf@usc.edu; http://www-rcf.usc.edu/~mwesterf/.

We are grateful to Daron Acemoglu, Markus Brunnermeier, Douglas Gale, Ken Garbade, Bengt Holmstrom, Dirk Jenter, Oguzhan Ozbas, Stephen A. Ross, David Sims, Michael Steinberger, Jianfeng Yu, and Jan Zabojnik for their helpful insights and comments.

1 Introduction

In organizations such as financial institutions, economic agents with potentially heterogeneous beliefs interact in non-market-mediated environments. Incentives are provided in contracts that can amplify or dampen the effects of belief heterogeneity. For example, incentive contracts for traders consist of a fixed salary and variable pay that depends on the profitability of the trader's portfolio. In the determination of the variable pay, should the trader's principal attempt to neutralize the effect of his disagreement with the trader, or can the principal take advantage of such disagreements?

Our model consists of a principal who hires an agent to manage a project, and the two do not share beliefs about the underlying evolution of the project. The principal and agent have heterogeneous beliefs about the probability distribution of the random innovations in the economy. In this model, we are explicitly assuming that the principal and agent can agree to disagree, so that after signing the contract they still do not agree on the project's evolution. There are two main justifications for this assumption, and they amount to a rejection of two conditions for the agreement theorem of Aumann (1976). The theorem states that economic agents cannot agree to disagree if 1) they have common priors *and* 2) they use Bayes rule to update beliefs. The first requirement leads to undesirable economic outcomes, such as no trade (see Milgrom and Stokey (1982)); it has come under theoretical attack in Gul (1998) and models such as Acemoglu, Chernozhukov, and Yildiz (2007); and loosening it has led to interesting results (see, for example, Yildiz (2003) in the context of bargaining). The second requirement of Bayesian learning has been widely questioned in the behavioral finance and economics literature (see Mullainathan and Thaler (2001) and Baker, Ruback, and Wurgler (2006) for surveys).

We specify a general dynamic contracting environment that includes moral hazard (hidden action) where differences in beliefs are modeled as subjective probabilities that can be generated either by Bayesian learning with heterogeneous priors *or* by non-Bayesian updating. This notion of subjective probabilities dates back to at least Savage (1954) and Anscombe and Aumann (1963). Empirical work documenting the importance of heterogeneous beliefs abounds. For example, Ito (1990) finds belief heterogeneity among currency traders. Odean (1999) demonstrates that investors with brokerage accounts tend to trade too often and interprets the finding as evidence of belief heterogeneity. Belief heterogeneity has also been shown to influence economic outcomes. Jenter (2004) demonstrates that some managers have systematically contrarian views, which in turn affect their decision making. Landier and Thesmar (2006) find that biased beliefs of entrepreneurs affect both the type of financing the entrepreneurs receive and their future profitability.

Our key insight is that disagreement creates potential gains from trade and from risk shifting. The principal and agent would like to place side-bets on the resolution of the economy's underlying uncertainty. Because they are limited to non-market interaction, the participants are constrained to incorporate state-contingent trades in the optimal contract. In addition to the usual trade-off between insurance and incentive provision, the principal incorporates side-bets over states of the world into the contract. Moreover, when the participants learn through time, it induces a correlation between the total gains from side-bet trade and the project's outcome. This correlation is a new source of risk and it changes the optimal degree of risk sharing between the principal and agent.

The gains from trade generated by differences in beliefs create a convexity in the principal's expected payoff. Consider a principal that is pessimistic relative to the agent. When the resolution of the underlying uncertainty is low, the principal becomes more pessimistic, and he increasingly desires to "sell" consumption in high states to the agent. The principal and agent place a series of side-bets by making the agent's payment more sensitive to the underlying project. This has the additional effect of strengthening the agent's incentives in the hidden action problem. As belief heterogeneity increases, the existing trade becomes more valuable and the principal will optimally choose to increase the quantity of trade. The result is that the value the principal places on his relationship with the agent is convex in the agent's relative optimism.

Beliefs also affect the optimal contract through a purely dynamic correlation effect: disagreement risk. Differences in opinion are correlated with both the project's outcome and the principal's certainty equivalent wealth through the learning process. As a result, when disagreement is moderate, the principal will try to shift some of the project's risk onto the agent. In doing so, the principal will implement a higher level of effort. This effect is only a result of the dynamic nature of the contract; It is based on correlation and is present even in states of the economy in which the principal and agent momentarily agree.

One important implication of these intuitions regards the first-best outcome. In a standard risk-neutral principal-agent model, the first-best can be achieved by "selling the firm" to the agent who will then choose to maximize total surplus. This is no longer the case with disagreement because such a solution precludes the principal from selling consumption to the agent in extreme states. In fact, if the principal and agent are both risk-neutral, the first-best solution does not exist because the participants would prefer an unbounded number of side-bets in addition to any other transactions. If the agent is risk-averse, then the first-best sharing rule might be steeper than in the second-best when beliefs are very different: effort is dictated in the first-best, and so there is no penalty from motivating the very high level of effort that a steep contract in the second-best would create. However, when beliefs are only moderately different, the sharing rule in the second-best can be steeper because there is a need to incentivize the agent.

An important feature of our setup is that the optimal contract allows for flexibility in commitment. This can be economically important when contract participants change their beliefs over time. The optimal contract allows for early termination by the principal and provides sufficient incentives that the agent never wants to quit, no matter how his beliefs evolve. Thus, neither side ever becomes "disappointed" with the contract, even if one realizes that the project under management is much less valuable than originally believed. This works because the principal is able to keep the agent on the edge of indifference to termination as the agent's beliefs evolve, and the principal can always offer the agent a terminal payment. We do not require the project to run to completion.

We present an illustration of our results in the context of a Bayesian model with heterogeneous priors, in which the principal learns about the underlying profitability of the project over time. Because of the convexity in the principal's problem, the principal is better off when his own priors are uncertain and when the agent is highly optimistic. This maximizes the amount of disagreement, and the principal will offer the agent a contract with high-powered incentives.

We conduct our analysis of the principal-agent model in continuous time, as it is more easily solved than a discrete time model. In our setting, dynamic contracts can be found in closed or semi-closed form and are easy to interpret. We also believe that heterogeneous beliefs naturally live in a dynamic setting because beliefs change over time. Learning is an intrinsically dynamic mechanism, and in our setup the dynamics of beliefs enter naturally into the optimal contract design.

Our setup is closely related to the traditional principal-agent literature, especially the seminal continuous-time model of Holmstrom and Milgrom (1987). Our derivation of the optimal contract follows Westerfield (2006). Related continuous-time principal-agent models are presented in Schättler and Sung (1993), Sannikov (2006), and Cvitanić and Zhang (2006). While, to our knowledge, the study of optimal contracting with heterogeneous beliefs in a dynamic principal-agent setting is new to this paper, static principal-agent models have been used to analyze situations with heterogeneous beliefs in papers such as Gervais, Heaton, and Odean (2006) and Benabou and Tirole (2002). Another example of conceptually related work is Van den Steen (2001), who studies a variety of managerial problems in the presence of differing beliefs. In addition, Bolton, Scheinkman, and Xiong (2006) study a dynamic model of CEO compensation under heterogeneous investor beliefs. Our primary contributions relative to these papers are 1) a general problem, 2) a simple and intuitive solution, and 3) flexibility of application.

Section 2 is an intuitive preview of the main results within the context of a learning model. In section 3, we lay out the details of the model and continue with the full solution in section 4. Section 5 concludes.

2 Contracting with Learning

In this section, we consider an example that illustrates the main results and intuitions discussed more generally in later sections.[1] We assume that the principal and agent start their contractual relationship with heterogeneous priors about the underlying evolution of a given project. While they know each other's beliefs, they still disagree, and that disagreement affects the optimal contract. Over time, Bayesian learning leads to a convergence of beliefs and hence to changes in incentive provision. We do not need to take a stance on who is wrong and who is right – we will see that it is belief differences that drive economic behavior.

2.1 The Setup

The incentive problem arises from an information asymmetry. The agent works on a project that pays Y_T at time T. While the principal can observe the output of the project, he cannot observe the agent's input (effort). The principal makes a payment C_T to the agent at time

[1]The results in this section are special cases of the general results in section 4, and so we leave the rigorous verification of our statements to that section. Technical requirements for variables and processes are in section 3.

T that depends on the entire observable history of Y.

We assume that the project's growth under some reference probability measure depends on the agent's effort, μ, so that

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

σ is a positive constant. The agent's choice of effort can be any real number.

The principal and the agent disagree with each other regarding the stochastic evolution of Y_t. We will assume that the agent assesses the project under the reference probability measure (but *not* necessarily the objective measure), while the principal believes the rate of growth for the project is $\mu_t - \sigma\delta_t$. So, if $\mathbb{P}$ is the principal's probability measure, and $\mathbb{A}$ is the agent's probability measure, we have

$$dY_t = (\mu_t - \sigma\delta_t)\, dt + \sigma dB_t^{\mathbb{P}} \tag{2a}$$

$$dY_t = \mu_t dt + \sigma dB_t^{\mathbb{A}} \tag{2b}$$

where $B_t^{\mathbb{P}}$ and $B_t^{\mathbb{A}}$ represent Brownian motions under the principal's and agent's probability measures, respectively. In (2), we see that δ_t parameterizes the differences in beliefs between the principal and the agent; δ_t is the agent's relative optimism, which can evolve over time.

In this learning example, the difference in beliefs about the growth rate of the project stems from a difference in priors that is not resolved by signing the contract. We assume that the principal has a Gaussian prior at time 0 with prior mean δ_0 and uncertainty γ_0. The principal uses Bayesian updating, so the prior means and uncertainties evolve as

$$d\delta_t = -\frac{\gamma_t}{\sigma^2}\left(dY_t - \mu_t^* + \sigma\delta_t\right) = -\frac{\gamma_t}{\sigma} dB_t^{\mathbb{P}} = \frac{\gamma_t}{\sigma}\delta_t dt - \frac{\gamma_t}{\sigma} dB_t^{\mathbb{A}} \tag{3a}$$

$$d\gamma_t = -\frac{{\gamma_t}^2}{\sigma^2} dt \tag{3b}$$

according to the Kalman-Bucy filter, as presented in Liptser and Shiryaev (2000). We have used μ_t^* instead of μ_t because the principal will solve his problem assuming that the agent uses his optimal control.[2]

[2]In this section, we use heterogeneous priors and Bayesian updating to generate agreement to disagree. However, we could have easily substituted a non-Bayesian process for δ. It is also not required that we specify beliefs that converge over time. Our key requirement is that the principal and agent do not resolve their belief differences when the contract is signed.

In this setup, one can view the agent as an expert who fully understands the economy, or as a noise trader who does not learn. We do not take any stand on which of these is correct, and, indeed, it will not affect the optimal contract.

It is the case that over any dt interval, the agent can only disagree with the principal about the drift in dY_t. However, this is not true over any finite period. For example, since $-\delta_t$ is positively correlated with $B_t^{\mathbb{P}}$, then a series of positive shocks will make the principal increasingly optimistic relative to the agent. This will imply that the agent, using $\mathbb{A}$, will disagree with the principal about the variance of B_T. In fact, with general δ, there is no reason that $B_T^{\mathbb{P}}$ need even be normally distributed under $\mathbb{A}$, and vice-versa.

The fundamental question that this paper addresses is how the optimal contract depends on the dynamics of the disagreement between the principal and the agent. A priori, one might think that the principal would want to write a contract that undoes differences in beliefs, so that the agent has the incentive to act "as if" he shared beliefs with the principal. We will see that in the optimal contract, the principal does not attempt to neutralize belief differences, but rather takes advantage of those differences in providing incentives.

The solution to this problem, and indeed the general heterogeneous beliefs problem, exhibits a separability between contracting and learning. While the learning problem influences the optimal contract, the contract choice does not change the learning process. In finding the optimal contract, we can treat the evolution of the difference in beliefs as exogenous. This makes economic sense because the contract affects only the agent's effort level, not outside business conditions or whatever else might drive the Brownian innovation term.

In this example, we assume that the agent faces a quadratic financial cost of effort[3], $\frac{1}{2}\mu_t^2$, and exponential utility with a coefficient of absolute risk aversion a. The principal will also

[3]In order to keep the example simple, we follow Holmstrom and Milgrom (1987) and make several simplifying assumptions such as CARA utility, constant volatility σ, and quadratic cost of effort. One additional simplification is that the cost function extends over the entire real line. We can justify this in two ways. First, the CARA utility can be viewed as an approximation to a particular section of the agent's utility function, and the cost functional as an approximation to the true cost function for a region in which the agent is not too pessimistic. Similar justifications are given in the affine term structure literature or the CARA-normal asymmetric information asset-pricing literature. A second justification is that the principal monitors the agent and makes it costly for the agent to sabotage the project. Under this interpretation, the agent pays a cost of effort when $\mu_t > 0$ and a cost to avoid sabotage monitoring when $\mu_t < 0$. Then, assuming no monitoring (free sabotage) is equivalent to imposing the constraint that $\beta_t \geq 0$. It is still possible to solve for the principal's value function with this constraint, but it makes the Hamilton-Jacobi-Bellman equation significantly more complicated. The additional constraint does not change the nature of the solution qualitatively.

have exponential utility, but with coefficient of absolute risk aversion A. The agent solves the following program:

$$\begin{aligned} \mu^* \quad &\in \quad \arg\max_{\mu} \; \mathrm{E}^{\mathbb{A}}\left[-\exp\left(-a\left(C_T - \frac{1}{2}\int_0^T \mu_t^2 dt\right)\right)\right] \\ &\text{s.t.} \quad dY_t = \mu_t dt + \sigma dB_t^{\mathbb{A}} \end{aligned}$$

where the expectation and the Brownian motion are both taken under the agent's probability measure/beliefs $\mathbb{A}$. We also assume that the agent has a participation constraint: he requires a certainty equivalent wage of $\hat{w}$ in order to accept the contract.

2.2 The Contract

Using the unravelling argument in section 4, we can show that the contract has a "locally linear" form. The time T value of the process C_t is C_T with

$$dC_t = \left(\frac{1}{2}\mu_t^{*2} + a\frac{1}{2}\beta_t^2\sigma^2\right) dt + \beta_t \left(dY_t - \mu_t^* dt\right) \tag{4}$$

where β_t is a process chosen by the principal and $C_0 = \hat{w}$ is the agent's certainty equivalent reservation wage. In equilibrium, the last term $(dY_t - \mu_t^* dt)$ is an accelerated Brownian motion under the agent's measure, $\sigma dB_t^{\mathbb{A}}$. A key fact about (4) is that it represents the agent's actual payment at time T as a function of what the agent is supposed to do (through μ_t^*) and what the agent actually does (through μ_t in dY_t).

The contract (4) is easy to interpret. The first term $(\frac{1}{2}\mu_t^{*2})$ is the agent's direct financial cost of effort for which he must be reimbursed. The second term $(a\frac{1}{2}\beta_t^2\sigma^2)$ represents the agent's insurance payment: it is the cost the principal must pay to insure the agent against the unobservable risk in the project. The last term $(\beta_t (dY_t - \mu_t^* dt))$ is the incentive term in the contract: when dY_t is high, the agent is paid a fraction of the increase above what was expected.

While μ^* is the agent's optimal control, we must relate the slope of the contract (β_t) to the agent's choice. In other words, we must find the value of β_t that makes μ_t^* incentive compatible. Solving for the agent's optimal control, as in section 4, we find that μ_t^* is incentive compatible if and only if

$$\mu_t^* = \arg\max_{\mu_t} \beta_t \mathrm{E}^{\mathbb{A}} \left[dY_t \right] - \frac{1}{2}\mu_t^2 dt \tag{5}$$

so that $\mu_t^* = \beta_t$. The agent equalizes the marginal cost of effort μ_t^* with the marginal benefit of an additional unit of effort β_t.

The next step is to analyze the contract under the principal's measure. Since $dB_t^{\mathbb{A}} = -\delta_t dt + dB_t^{\mathbb{P}}$, we find from (4) and (5) that

$$\begin{aligned} dC_t &= \left(\frac{1}{2}\mu_t^{*2} + a\frac{1}{2}\beta_t^2\sigma^2\right) dt + \beta_t \left(dY_t - \mu_t^* dt\right) & \text{(6a)} \\ &= \left(\frac{1}{2}\beta_t{}^2 + a\frac{1}{2}\beta_t^2\sigma^2\right) dt + \beta_t \sigma dB_t^{\mathbb{A}} & \text{(6b)} \\ &= \left(\frac{1}{2}\beta_t{}^2 + a\frac{1}{2}\beta_t^2\sigma^2 - \beta_t\sigma\delta_t\right) dt + \beta_t \sigma dB_t^{\mathbb{P}} & \text{(6c)} \end{aligned}$$

where we used incentive compatibility to move from (6a) to (6b) and the definitions of the probability measures to move from (6b) to (6c).

The key term of the contract under the principal's measure is $(\beta_t\sigma\delta_t)$, representing the gains from trade from heterogeneous beliefs. When the agent is more optimistic than the principal ($\delta_t > 0$), the agent believes that his portion of the project ($\beta_t dY_t$) is more valuable than the principal does. However, to meet the agent's participation constraint and provide incentives, the principal needs to pay the agent only part of the excess in $\beta_t dY_t$ according to the *agent's* measure: $\beta_t \left(dY_t - \mu_t^* dt\right)$. To the principal, the true excess is $\beta_t \left(dY_t - \mu_t^* dt + \sigma\delta_t dt\right)$, and the difference is a gain to the principal of $(\beta_t\sigma\delta_t)$. The principal increases this gain by increasing β_t, i.e. steepening incentives.

It is important to note that the objective probability measure does not appear in the optimal contract, nor is any objective probability measure necessary to even define the problem. In fact, δ_t appears in the contract only as the difference in beliefs between the principal and agent (the change from $\mathbb{A}$ to $\mathbb{P}$). This illustrates a key point – *what matters is the difference in beliefs, not the belief levels.* The level of beliefs fixes the level of the value functions of the principle and the agent, but their actions are determined by the difference in their beliefs.

2.3 The Principal's Problem

We can now replace the contract and the optimal effort into the principal's problem:

$$\begin{aligned}\beta_t^* &= \max_{\beta_t} \mathrm{E}^{\mathbb{P}}\left[-\exp\left(-A\left(W_T = Y_T - C_T\right)\right)\right] \\ \text{s.t.} \quad dW_t &= \left[\beta_t - \sigma\delta_t - \frac{1}{2}\beta_t^2 + \beta_t\sigma\delta_t - a\frac{1}{2}\beta_t^2\sigma^2\right]dt + \left(1-\beta_t\right)\sigma dB_t^{\mathbb{P}}, \quad W_0 = Y_0 - \hat{w}\end{aligned}$$

The drift of the wealth process has five components. The first and second components ($\beta_t - \sigma\delta_t = \mu_t^* - \sigma\delta_t$) are the drift of Y_t under the principal's probability measure. The third component ($\frac{1}{2}\beta_t^2 = \frac{1}{2}\mu_t^{*2}$) is the compensation that the principal pays the agent for his effort. The fourth component ($\beta_t\sigma\delta_t$) is due to belief heterogeneity – when the agent has positive beliefs relative to the principal ($\delta_t > 0$), the principal gains from the difference in the drift of Y under $\mathbb{P}$ and $\mathbb{A}$. The last component ($-a\frac{1}{2}\beta_t^2\sigma^2$) of the principal's wealth growth is the compensation for the agent's risk-aversion. Finally, a fraction β_t of the innovations to the project goes to the agent, and the remainder goes to the principal.

To solve the principal's problem, we will use dynamic programming. The principal's value function, which we later verify, has the form

$$V(t, W_t, \delta_t) = -\exp\left(-A\left(W_t + F(t) + G(t)\delta_t + H(t)\delta_t^2\right)\right)$$

with boundary condition $V(T, W_T, \delta_T) = -\exp\left(-AW_T\right)$.

The Hamilton-Jacobi-Bellman equation is

$$\begin{aligned}0 = \max_{\beta_t} \quad & -AV(t, W_t, \delta_t)\left[\beta_t - \sigma\delta_t - \frac{1}{2}\beta_t^2 + \beta_t\sigma\delta_t - a\frac{1}{2}\beta_t^2\sigma^2 - \frac{1}{2}A\left(1-\beta_t\right)^2\sigma^2\right. \\ & +F'(t) + G'(t)\delta_t + H'(t)\delta_t^2 + A(1-\beta_t)\left(G(t) + H(t)2\delta_t\right)\gamma_t \\ & \left. + \left(H(t) - \frac{1}{2}A\left(G(t) + 2H(t)\delta_t\right)^2\right)\left(\frac{\gamma_t}{\sigma}\right)^2\right]\end{aligned}$$

and so the principal's optimal choice is

$$\beta_t^* = \frac{1 + A\sigma^2 + \sigma\delta_t - A\gamma_t\left(G(t) + H(t)2\delta_t\right)}{1 + a\sigma^2 + A\sigma^2} \tag{7}$$

which we will interpret later.

Plugging β_t^* in the HJB equation yields the following set of ODEs:

$$\begin{aligned}
H'(t) &= 2\frac{\gamma(t)}{\sigma}\frac{A\sigma^2 H(t)}{1+a\sigma^2+A\sigma^2}+2\frac{\gamma(t)^2}{\sigma^2}A\frac{(1+a\sigma^2)H(t)^2}{1+a\sigma^2+A\sigma^2}-\frac{\sigma^2}{2(1+a\sigma^2+A\sigma^2)} \\
G'(t) &= 2A\frac{\gamma(t)^2}{\sigma^2}\frac{1+a\sigma^2}{1+a\sigma^2+A\sigma^2}G(t)H(t)+\sigma\frac{a\sigma^2-2\sigma Aa\gamma(t)H(t)+A\gamma(t)G(t)}{1+a\sigma^2+A\sigma^2} \\
F'(t) &= -H(t)\frac{\gamma(t)^2}{\sigma^2}+\frac{-2AG(t)\gamma(t)a\sigma^4-G(t)^2A\gamma(t)^2(1+a\sigma^2)}{2\left(1+a\sigma^2+A\sigma^2\right)}-\frac{1+A\sigma^2-Aa\sigma^4}{2\left(1+a\sigma^2+A\sigma^2\right)}
\end{aligned}$$

where $\gamma(t)=\frac{\gamma_0\sigma^2}{\gamma_0 t+\sigma^2}$. The ODEs can be easily solved numerically with the boundary condition $F(T)=G(T)=H(T)=0$ (and they can be solved analytically when the principal is risk-neutral: $A=0$). We interpret the solution to the principal's optimal control choice and welfare function in the next section.

2.4 Interpretation

The key finding is that both dynamic and static differences in beliefs determine the incentive contract. Relative to the Holmstrom and Milgrom (1987) baseline (no differences of opinion), our expression for the slope of the contract (7) and the optimal level of effort (μ_t^*) contains two extra terms:

$$\sigma\delta_t - A\gamma_t\left(G(t)+H(t)2\delta_t\right)$$

The first term, $\sigma\delta_t$, is a direct effect: when the agent is relatively more optimistic ($\delta_t>0$), the principal grants steeper incentives so as to maximize the value of the contract and the side-bets on the project's outcome. When the agent is relatively pessimistic, β_t^* declines, again because the slope of the contract creates both incentives and side-bets.

The second term, $-A\gamma_t\left(G(t)+H(t)2\delta_t\right)$, is a correlation effect – disagreement risk. Differences of opinion (δ_t) are correlated with the project's outcome (Y_t) and with the principal's certainty equivalent value from the remainder of the contract. This correlation creates a desire to shift risk between the participants, and this can only be accomplished through the contract and hence through β. In doing so, the principal changes the effort level that the contract implements.

The principal's certainty equivalent wealth is made up of two parts: one from the value of the project itself and one from the value of the side-bets in the contract. When δ_t is small, project profitability dominates, the correlation between beliefs and certainty equivalent wealth is positive ($G(t) < 0$), and the principal pushes more of the project onto the agent. When δ_t is large, the contract dominates, the correlation is negative ($H(t) > 0$), and the principal keeps more of the project for himself.

Notice that the correlation effect occurs whether or not the principal and agent actually disagree at any particular time. If $\delta_t = 0$ for some t, the direct effect vanishes, but the correlation effect does not because it is based on future changes. For the same reason, the correlation effect is weaker as one moves closer to the completion of the contract. The correlation effect is illustrated in figures 1 and 2.

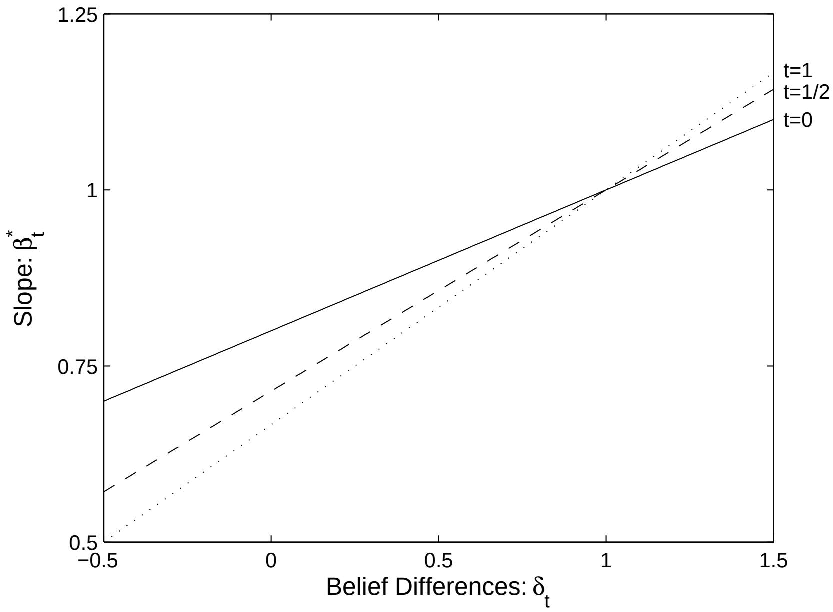

Figure 1: β^*_t as a function of δ_t for $t = 0$, $t = \frac{1}{2}$, and $t = 1$. We set $A = a = \sigma = T = 1$ and $\gamma_0 = 2$. For reference, the baseline value of β^*_t under homogeneous beliefs for these parameters is 2/3.

We also examine the principal's certainty equivalent wealth from the differences in beliefs and the contracting process:

$$-\frac{1}{A} \ln\left(-V(0, W_0 = 0, \delta_0)\right) + \sigma\delta_0 T$$

The latter term is added to cancel out the expected profitability of the project under the

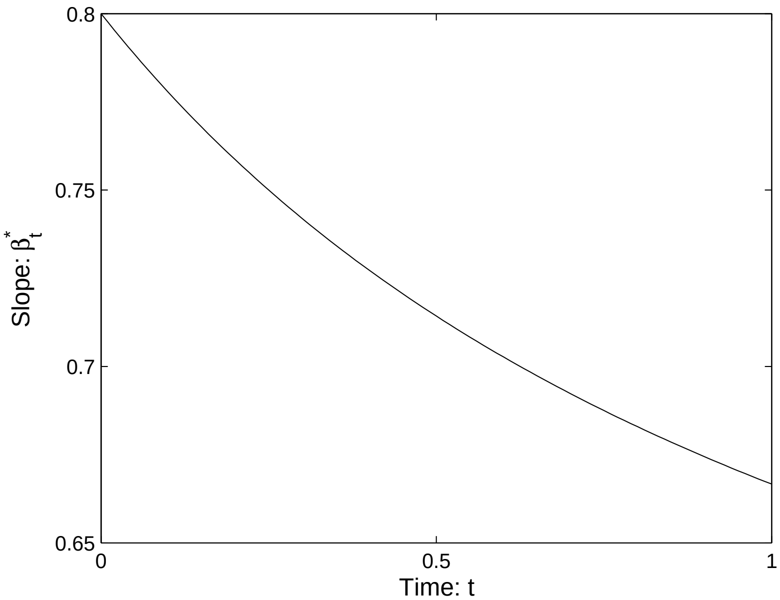

Figure 2: β_t^* as a function of time when $\delta_t = 0$. We set $A = a = \sigma = T = 1$ and $\gamma_0 = 2$.

principal's measure. We wish to examine gains from contracting under heterogeneity, rather than the expected payoff from the project itself. This quantity is plotted in figure 3 .

We observe two facts in figure 3. The first observation is that the principal's value function is convex in δ_0. This results from combining income and substitution effects. The principal's direct cash flow from belief differences is $\beta_t \sigma \delta_t$. Therefore, as δ_t increases, the principal's direct cash flows increase as well. However, as δ_t increases, (7) and $H(t) > 0$ show us that the principal increases the steepness of the agent's incentives. The first is an income effect, the second a substitution effect, and the sum is a more than linear increase in the principal's utility as a function of δ_0.

The convexity of the value function drives the second observation: The principal has a higher expected value to the contract when γ_0 is large. This is because a series of surprise innovations in B_t will cause δ_t to vary as if by adding a mean-preserving spread under $\mathbb{P}$ (3a). Because the principal's value function is convex in δ_0, he is better off with relative priors that can easily change.

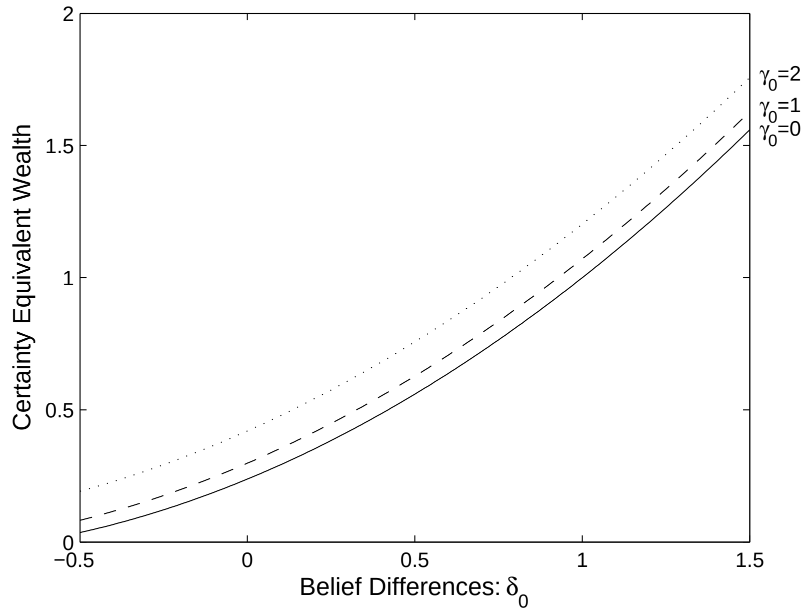

Figure 3: The principal's value to contracting is the function $-\frac{1}{A}\ln\left(-V(0, W_0 = 0, \delta_0)\right) + \sigma\delta_0 T$. We plot this as a function of δ_0 for three values of γ_0. The plot sets $T = a = \sigma = 1$ and $A = .1$.

2.5 Comparison to the First-Best

We now compare the second-best results of the previous section to the situation under the first-best assumption that the principal can costlessly dictate the agent's choice of μ_t^* and that the principal is risk-neutral.[4] In the standard principal-agent problem, it would be the case that the principal pays the agent a fixed fee to meet the participation constraint, but that is not what happens here. Keeping the same notation of C_T for the sharing rule, the principal's first-best problem is to find

$$\begin{aligned} \{C_T^{FB}, \mu^{FB}\} \quad &\in \quad \arg\max_{C_T, \mu} \mathrm{E}^{\mathbb{P}}\left[Y_T - C_T\right] \qquad (8) \\ \text{s.t.} \quad (\text{i}) \quad & dY_t = (\mu_t - \sigma\delta_t)\, dt + \sigma dB_t^{\mathbb{P}} \\ (\text{ii}) \quad & \mathrm{E}^{\mathbb{A}}\left[-\exp\left(-a\left(C_T - \frac{1}{2}\int_0^T \mu_t^2 dt\right)\right)\right] \geq -\exp\left(-a\hat{w}\right) \end{aligned}$$

[4]We set $A = 0$ to simplify the algebra and to allow for a purely analytical comparison between first and second best.

To continue, we also need to define the Radon-Nikodym derivative ξ_T between the probability measures $\mathbb{P}$ and $\mathbb{A}$:

$$\xi_T = \exp\left[-\frac{1}{2}\int_0^T \delta_t^2 dt + \int_0^T \delta_t dB_t^{\mathbb{P}}\right] \tag{9}$$

This means that if Z_T is an appropriate random variable, $\mathrm{E}^{\mathbb{A}}\left[Z_T\right] = \mathrm{E}^{\mathbb{P}}\left[\xi_T Z_T\right]$.

To solve the principal's problem, we will use the planner's approach[5] and maximize

$$\begin{aligned} & \mathrm{E}^{\mathbb{P}}\left[Y_T - C_T\right] + \lambda\left[\mathrm{E}^{\mathbb{A}}\left[-\exp\left(-a\left(C_T - \frac{1}{2}\int_0^T \mu_t^2 dt\right)\right)\right] + \exp\left(-a\hat{w}\right)\right] \\ = \; & \mathrm{E}^{\mathbb{P}}\left[\int_0^T\left(\mu_t - \sigma\delta_t - \frac{1}{2}\mu_t^2\right)dt - \hat{C}_T - \lambda\xi_T\exp\left(-a\hat{C}_T\right)\right] + \lambda\exp\left(-a\hat{w}\right) \end{aligned}$$

where $\hat{C}_T = C_T - \frac{1}{2}\int_0^T \mu_t^2 dt$ and λ is the Lagrange multiplier on the agent's participation constraint. Since δ_t and ξ_T are exogenous, the expectation can be maximized point-wise to find that $\mu_t^{FB} = 1$ and $\lambda\xi_T a\exp(-a\hat{C}_T^{FB}) = 1$. The equilibrium value for μ_t^{FB} is the standard first-best solution, but the equilibrium payment to the agent is new. Without heterogeneous beliefs, $\xi_T = 1$ and the principal would make a fixed payment to the agent. With heterogeneous beliefs, the principal makes a payment

$$C_T^{FB} = \frac{1}{2}\int_0^T \mu_t^{FB^2} dt + \frac{1}{a}\ln\left(\lambda a\xi_T\right) \tag{10}$$

As in the second-best, while individual beliefs (δ_t) enter the principal's welfare function, their levels do not affect the form of the solution: ξ_T is purely a function of belief differences.

While (10) is a full description of the sharing rule, we want to put it into a form in which we can directly compare it to the sharing rule in the second-best case (6). First, we solve for λ using the agent's participation constraint:

$$-\exp\left(-a\hat{w}\right) = \mathrm{E}^{\mathbb{A}}\left[-\exp\left(-a\hat{C}_T^{FB}\right)\right] = \mathrm{E}^{\mathbb{A}}\left[-\frac{1}{a\lambda\xi_T}\right] = \mathrm{E}^{\mathbb{P}}\left[-\frac{1}{a\lambda}\right] = -\frac{1}{a\lambda}$$

[5] Formally, one might think that C_T must still be based on Y because that is the only economic variable that the principal and agent can observe and agree on. However, because both the principal and agent can observe μ_t in the first-best, they can both observe $B_t^{\mathbb{A}}$ and $B_t^{\mathbb{P}}$ as well. Thus, equilibrium payments can be directly a function of the Brownian motions.

So $\lambda = \frac{1}{a}\exp(a\hat{w})$. Then, we set

$$C_t^{FB} = \frac{1}{2}\int_0^t {\mu_s^{FB}}^2 ds + \frac{1}{a}\ln(\lambda a \xi_t) = \frac{1}{2}\int_0^t {\mu_s^{FB}}^2 ds + \frac{1}{a}\ln(\xi_t) + \hat{w}$$

and $C_0^{FB} = \hat{w}$, which is the same as in the second-best case (as required by the participation constraint). Using Ito's lemma, we find that

$$dC_t^{FB} = \frac{1}{2}{\mu_t^{FB}}^2 dt - \frac{1}{2a}\delta_t^2 dt + \frac{1}{a}\delta_t dB_t^{\mathbb{P}} \tag{11}$$

which we can directly compare to the second-best case in (6) and (7).

Comparison of (11) to (6) and (7) reveals one key difference: the volatility. The first-best "slope" corresponds to $\beta_t^{FB} = \frac{1}{a\sigma}\delta_t$, compared to $\beta_t^{SB} = \frac{1+\sigma\delta_t}{1+a\sigma^2}$ in the second-best. However, the contract is unchanged *as a function of* β_t *and* μ_t: the first-best drift from (11) is

$$\frac{1}{dt}\mathrm{E}\left[dC_t^{FB}|\mathcal{B}_t\right] = \frac{1}{2}{\mu_t}^2 - \frac{1}{2a}\delta_t^2 = \frac{1}{2}{\mu_t}^2 + a\frac{1}{2}{\beta_t^{FB}}^2\sigma^2 - \delta_t\beta_t^{FB}\sigma$$

which is equal to the second-best drift from (6). The principal must still reimburse the agent for his direct cost of effort, he must still pay the agent an insurance payment to compensate him for risk, and he still receives gains from trade based on beliefs. The difference is that β_t^{FB} captures only side-bets and not incentives, while β_t^{SB} captures both. These differences are illustrated in figure 4.

The key variation is in how β_t – the "slope" – is different in the first- and second-best cases. When δ_t is near zero, the second-best contract is steeper: the incentive provision dominates differences in beliefs in the optimal contract. When δ_t is large, the first-best sharing is steeper because enormous side-bets can be implemented with no cost penalty associated with a too-high effort level. Consider the case in which the agent is more optimistic than the principal ($\delta_t > 0$). This means that the agent believes that the high Y_t states (high $B_t^{\mathbb{P}}$ or $B_t^{\mathbb{A}}$ states) are relatively more likely than the principal does. There are gains from trade in consumption with the principal selling the agent consumption when Y_T is very high. From the definition of ξ_T, when Y_T (or $B_T^{\mathbb{P}}$ or $B_T^{\mathbb{A}}$) is realized to be very positive, ξ_T will be very large and the agent's terminal payment (C_T) will be very large. Conversely, if Y_T is realized to be relatively low, then the corresponding agent's payment will be small.

This result holds because, in the first-best, the contract is constructed so as to maxi-

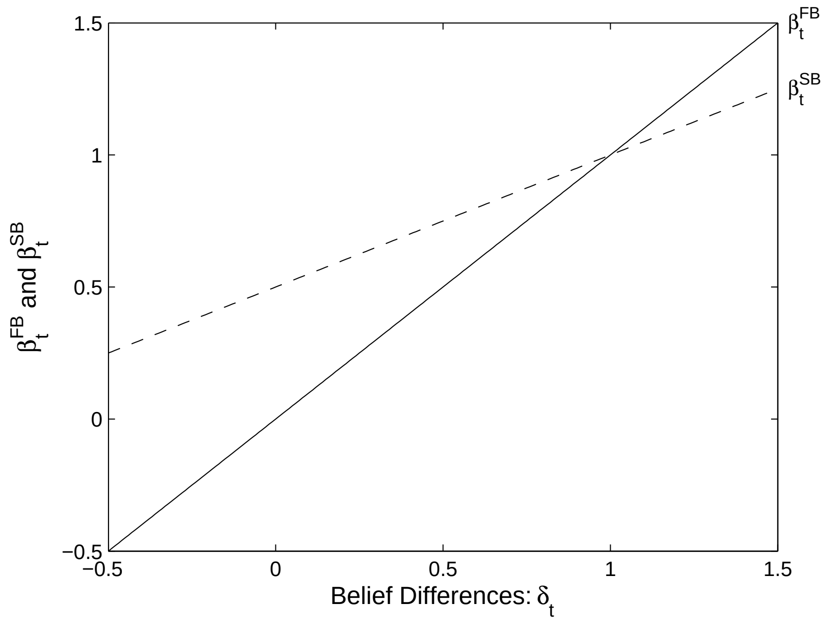

Figure 4: β_t^{FB} and β_t^{SB} as functions δ_t. We set $a = \sigma = 1$ and $A = 0$.

mize its joint value. However, the principal and agent differ in their valuation of various states because they differ in their assessment of the probabilities. So, to maximize the joint valuation of the contract, the principal allocates more consumption to the agent in states that the agent believes are more likely and more valuable, and the principal allocates more consumption to himself in the states that he believes are more likely and more valuable. Therefore, the first-best sharing rule must be a function of the different beliefs.

2.6 The Risk-Neutral Case

Standard principal-agent moral-hazard models, such as that in Holmstrom (1979), impose a trade-off between insurance and incentives. However, with homogenous beliefs, that trade-off disappears when both the principal and agent are risk-neutral. The first-best can then be implemented by selling the firm to the agent. With heterogeneous beliefs, that result no longer holds. The reason is that the principal and agent have different valuations of the project and, as a result, different valuations of the contract. The difference between this case and the risk-averse case is that the potential benefits from trade are unbounded. This can be seen by considering the central planner's problem in the first-best case. The following objective function needs to be maximized with respect to C_T and μ:

$$\begin{aligned} & \mathrm{E}^{\mathbb{P}}\left[Y_T - C_T\right] + \mathrm{E}^{\mathbb{A}}\left[C_T - \frac{1}{2}\int_0^T \mu_t^2 dt\right] \\ = \; & \mathrm{E}^{\mathbb{P}}\left[\int_0^T \left(\mu_t - \sigma\delta_t - \frac{1}{2}\mu_t^2\right) dt - \hat{C}_T + \xi_T \hat{C}_T\right] \end{aligned}$$

where $\hat{C}_T = C_T - \frac{1}{2}\int_0^T \mu_t^2 dt$. Because the linearity of the objective function, a solution to this problem does not exist. For states in which $\xi_T > 1$, states the agent believes are more likely, the principal would like to transfer an infinite amount of consumption to the agent. Conversely, for states in which $\xi_T < 1$, the principal would like to demand an infinite amount of consumption from the agent. Another way to see this is to consider the limit of $\beta_t^{FB} = \frac{\delta_t}{a\sigma}$ as $a \to 0$: risk-neutrality creates unbounded consumption shifting.

While the first-best does not exist under risk-neutrality, the second-best does. In fact, it can be solved by replacing $a = 0$ in (4) and (7):

$$\begin{aligned} dC_t^* &= \frac{1}{2}\mu_t^{*2} dt + \beta_t \left(dY_t - \mu_t^* dt\right) \\ \beta_t^* &= 1 + \sigma\delta_t \end{aligned} \tag{12}$$

The result holds because of the quadratic cost of effort: if the principal were to try to implement a very large number of side-bets, he would have to do so through a very high β_t – a very steep contract. This would have the by-product of implementing such a high level of effort at such a high cost to the agent that the principal would have to pay too much to meet the agent's participation constraint.

Equation (12) shows that in the second-best under risk-neutrality, it is not optimal to sell the firm to the agent. Intuitively, the principal and agent disagree about the value of the project at time zero. The principal can extract more or less effort from the agent than the agent would exert by himself. The agent disagrees with the principal's assessment of the world, but responds optimally to incentives.

3 The General Model

Information Structure

Uncertainty is described by a Brownian motion, $B(t, \omega)$, for $0 \le t \le T$, defined on a complete probability space $(\Omega, \mathcal{F}, \Gamma)$, where Γ is the reference probability measure. $\mathcal{B}_t$ is the augmented filtration generated by B_t^Γ, which is a Brownian motion with respect to Γ. The probability space fulfills the usual conditions. All processes we consider are appropriately adapted to $\mathcal{B}_t$. All expectations will contain a superscript indicating which probability measure they are being taken under.

We introduce the process Y, such that

$$dY_t = \mu_t dt + \sigma_t dB_t^\Gamma$$

where μ_t will be defined later and Y_0 is a constant. We define $\mathcal{Y}_t$ to be the augmented filtration generated by Y_t. It will turn out that $\mathcal{B}_t$ represents the agent's information set, while $\mathcal{Y}_t$ represents the principal's information set. σ_t is an exogenous $\mathcal{Y}_t$-measurable process such that $0 < s \le \sigma_t \le S < \infty$.

We will also introduce two other probability measures, $\mathbb{P}$ and $\mathbb{A}$. Γ, $\mathbb{P}$, and $\mathbb{A}$ are all mutually absolutely continuous (meaning they agree on zero-probability events), with $\Gamma = \mathbb{A}$. The zero-probability events assumption is important because it rules out costless arbitrage opportunities. It will turn out that $\mathbb{P}$ is the principal's probability measure, and $\mathbb{A}$ is the agent's probability measure. Let $\xi_t \equiv (d\mathbb{A}/d\mathbb{P})_t$ denote the Radon-Nikodym derivatives of the probability measure $\mathbb{A}$, with respect to $\mathbb{P}$. Then

$$dB_t^{\mathbb{P}} = \delta_t^P dt + dB_t^{\mathbb{A}} \tag{13a}$$

$$\xi_t = \exp\left[-\frac{1}{2}\int_0^t {\delta_s}^2 ds + \int_0^t \delta_s dB_s^{\mathbb{P}}\right] \tag{13b}$$

We assume that δ_t is an element of $\mathcal{L}_2$[6] and that ξ_t is a martingale on $[0, T]$. We will also

[6]The spaces $\mathcal{L}_1$ and $\mathcal{L}_2$ are defined so that

$$\mathcal{L}_1 = \left\{X : \int_0^T |X_t| dt < \infty \quad \text{a.s.}\right\}$$

$$\mathcal{L}_2 = \left\{X : \int_0^T X_t^2 dt < \infty \quad \text{a.s.}\right\}$$

assume that the initial value, δ_0, and the functional forms of the evolution of δ_t are known to all parties. We do not require any particular type of learning, so we can allow

$$d\delta_t = f(t, \cdot)dt + g(t, \cdot)dY_t \tag{14}$$

for any $f(t, \cdot)$ and $g(t, \cdot)$ that are $\mathcal{Y}_t$-measurable and in $\mathcal{L}_2$ that also fulfill the martingale requirement for ξ_t.

It is important that we have not defined an objective probability measure, only a reference one. We take no stand on whether the principal, the agent, or both, are objectively wrong.[7]

We make no statement on the source of the disagreement between the principal and the agent, nor do we make any statement on the evolution of the differences in beliefs. Thus, the disagreement could be caused by heterogeneous priors, non-Bayesian learning, or a variation on "noise trading." We require only that both sides are aware of the initial magnitude of the difference and the functional form of its evolution. Similarly, δ_t may have any evolution, representing any type of learning process, as long as that process is known to both sides. It is not required that beliefs converge over time.

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 project are owned by the principal, but to undertake the project, the principle hires the agent. The agent exerts a control, $\mu(t, \omega)$, that is not observable to the principal. In choosing μ, the agent determines the evolution of Y_t:

$$dY_t = \mu_t dt + \sigma_t dB_t^{\mathbb{A}} \tag{15}$$

We will assume that the agent chooses μ_t so that the process μ must be an element in $\mathcal{L}_1$.

The Contract

In return for the agent's labor, the principal offers the agent a contract or sharing rule that specifies a terminal payment C_T, payable at time T. The principal cannot directly observe μ, but he can observe the history of Y. As a result, the principal offers a payment that depends

[7]In addition, because $\mathbb{P}$ and $\mathbb{A}$ are mutually absolutely continuous, their augmented filtrations, $\mathcal{B}_t^{\mathbb{A}}$ and $\mathcal{B}_t^{\mathbb{P}}$ agree, and so we will simply write $\mathcal{B}_t$.

on the entire path of Y 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 for which C_T is $\mathcal{Y}_T$-measurable and $\mathrm{E}^{\mathbb{A}}[\exp(-C_T)]$ exists and is finite.

Objectives and Beliefs: Agent

The agent's actions and associated objective function are taken with respect to the agent's probability measure $\mathbb{A}$. Under that measure, according to (13) and (15), the project evolves as

$$dY_t = \mu_t dt + \sigma_t dB_t^{\mathbb{A}} \tag{16}$$

and the agent's objective function is

$$\mathrm{E}^{\mathbb{A}}\left[-\exp\left(-a\left(C_T - G_T\right)\right)\right] \tag{17}$$

The agent faces an opportunity cost for his effort and for his time. The agent pays a financial cost for his effort

$$G_T = \int_0^T g(\mu_t)dt \tag{18}$$

for some function g for which $g(\mu) \in \mathcal{L}_1$. We assume that $g(0) = 0$, $g'(\cdot) \geq 0$, and $g''(\cdot) > 0$. The total cost of effort is paid entirely upon completion of the economy at time T.

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

$$\max_{\mu} \mathrm{E}^{\mathbb{A}}\left[-\exp\left(-a\left(C_T - G_T\right)\right)\right] \geq \hat{U} \tag{19}$$

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

$$\begin{aligned} \mu^* \quad &\in \quad \arg\max_{\mu}\ \mathrm{E}^{\mathbb{A}}\left[-\exp\left(-a\left(C_T - G_T\right)\right)\right] \\ &\text{s.t.} \quad dY_t = \mu_t dt + \sigma_t dB_t^{\mathbb{A}} \end{aligned} \tag{20}$$

If μ^* solves the agent's problem for C_T, then we say that C_T *implements* μ^*.

Objectives and Beliefs: Principal

The principal's actions and associated objective function are taken with respect to the principal's probability measure $\mathbb{P}$. Under that measure, according to (13) and (15), the project evolves as

$$dY_t = (\mu_t - \sigma_t \delta_t)\, dt + \sigma_t dB_t^{\mathbb{P}} \tag{21}$$

Here, δ_t parameterizes the principal's beliefs about the evolution of Y. In principle, δ_t is a $\mathcal{Y}_t$-measurable process, although the principal will be able to infer the motion of $B_t^{\mathbb{P}}$ in equilibrium.

The principal's objective function is

$$\mathrm{E}^{\mathbb{P}} \left[-\exp \left(-A \left(Y_T - C_T \right) \right) \right] \tag{22}$$

The principal's problem is to maximize his objective function subject to the constraints that the agent 1) accepts the contract and 2) behaves optimally. Thus, the principal's problem is to find C_T^* so that

$$\begin{aligned} C_T^* \quad & \in \quad \arg\max_{C_T} \mathrm{E}^{\mathbb{P}} \left[-\exp \left(-A \left(Y_T - C_T \right) \right) \right] \\ \text{s.t.} \quad \text{(i)} \quad & \mu[C] \text{ solves the agent's problem} \\ \text{(ii)} \quad & dY_t = (\mu_t[C] - \sigma_t \delta_t)\, dt + \sigma_t dB_t^{\mathbb{P}} \\ \text{(iii)} \quad & \mathrm{E}^{\mathbb{A}} \left[-\exp \left(-a \left(C_T - G_T \right) \right) \right] \Big|_{\mu = \mu[C]} \geq \hat{U} \end{aligned} \tag{23}$$

We will also 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 it. His certainty equivalent wealth from undertaking the project may be less than zero or fail to exist. In this case, we assume that the principal rejects the project.

Equilibrium

An equilibrium consists of a contract C_T^* and an implemented level of effort μ^*. μ^* must be adapted to the agent's information set (be $\mathcal{B}_T$-measurable) and solve the agent's problem (20). C_T^* must be adapted to the principal's information set (be $\mathcal{Y}_T$-measurable) and solve the principal's problem (23).

4 Optimal Contracting

In this section, we will outline and solve the more general problem of contracting under heterogeneous beliefs. In doing so, we will allow for all the generality of section 3, and we use the "unraveling" approach of Westerfield (2006) in deriving the optimal contract.

4.1 Solving the Agent's Problem

The Contract Form

The first step is to simplify the set of contracts the principal can offer the agent. To find the $\mathcal{Y}_T$-measurable terminal payment to the agent (C_T), we will unravel the principal's $\mathcal{Y}_t$-measurable estimate of the agent's utility:

$$\mathcal{V}_t = \mathrm{E}^{\mathbb{A}}\left[-\exp\left(-a\left(C_T - G_T\right)\right)|\mathcal{Y}_t, \mu_t = \mu_t^*\right] \tag{24}$$

Here, $\mathcal{V}_t$ is the principal's estimate of the agent's total expected utility. It is 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 behaving optimally, $\mu_t = \mu_t^*$ (which is $\mathcal{Y}_t$-measurable), because the principal assumes that the agent behaves optimally (23i) when determining payment.

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 $\mathcal{V}_0 = \hat{U}$.

We next 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 (from Davis and Varaiya (1973) and updated in Revuz and Yor (2005)) to show that there exists a ϕ_t such that

$$d\mathcal{V}_t = \phi_t\left(dY_t - \mu_t^* dt\right) \tag{25}$$

where $dY_t - \mu_t^* dt$ has zero drift under the agent's beliefs. In (25), ϕ_t represents the principal's control over the volatility of the agent's utility.

We use the principal's estimates of the agent's utility to find the value of the principal's terminal payment to the agent, C_T. Notice that $\mathcal{V}_T = -\exp\left(-a\left(C_T - G_T^*\right)\right)$, where G_T^* is

G_T evaluated at $\mu = \mu^*$. Let us define the process C_t so that $\mathcal{V}_t = -\exp(-a(C_t - G_t^*))$ and $\hat{U} = -\exp(-aC_0)$. Then C_T must be the time T value of the process C_t.

Substituting $\mathcal{V}_t = -\exp(-a(C_t - G_t^*))$ into (25) and using Ito's lemma, we find that

$$a \exp(-a(C_t - G_t^*)) \left[dC_t - g(\mu_t^*)dt - \frac{1}{2}a(\text{vol}(C_t))^2 dt\right] = \phi_t (dY_t - \mu_t^* dt)$$

where $\text{vol}(C_t)$ is the volatility of C_t. Matching volatility, solving for dC_t, and substituting $\phi_t = \beta_t a \exp(-a(C_t - G_t^*))$, we find

$$dC_t = g(\mu_t^*)dt + a\frac{1}{2}\beta_t^2\sigma_t^2 dt + \beta_t (dY_t - \mu_t^* dt) \tag{26}$$

The principal controls the volatility of the agent's expected utility process (25) ($\phi_t dY_t = \beta_t a \exp(-a(C_t - G_t^*)) dY_t$) by controlling the volatility of the agent's payment process (26) ($\beta_t dY_t$). So, the principal's terminal payment to the agent is the time T value of the process C_t, where β_t is the process the principal uses to determine the volatility of the agent's expected utility function.

The contract in (26) is stated as if the drift in dY_t were given under the agent's measure $\mathbb{A}$ (25) even though it uses only agreed upon observations (dY_t). This means that we have discovered how much the agent has to be paid under his own measure, and it is the sum of three terms. The first term represents the agent's direct cost for his optimal control choice. The third term represents the agent's payment as a function of the underlying economic project Y_t. The agent's steepness of incentives, or the "local slope" of his payment, is β_t. However, because the principal cannot observe the agent's choice, but instead only the project outcome, the agent faces payment risk based on random fluctuations in the project, and so the agent must be compensated for that risk. This compensation makes up the second term.

We now turn to the relationship between β_t and μ_t^*.

The Optimal Control

As stated, we have the optimal contract as a function of the "slope" of the contract, β_t, the agent's optimal control μ_t^*, and the agent's actual control μ_t. However, we do not yet know how to make the contract *incentive compatible*; we need to know how the principal sets β_t in order to control the agent's choice of μ_t^*. We will use dynamic programming techniques to solve this problem.

To proceed, we use the continuous-time version of the Hamilton-Jacobi-Bellman equation. If $V(t, C_t, G_t)$ represents the agent's value function, optimality requires

$$0 = \max_{\mu_t} \mathrm{E}^{\mathbb{A}} \left[dV(t, C_t, G_t) \right]$$

Using Ito's lemma to write out the Hamilton-Jacobi-Bellman equation gives us

$$\begin{aligned} 0 \;=\; & \max_{\mu_t} \left[V_t + V_G g(\mu_t) + \frac{1}{2} V_{CC} \left(\beta_t \sigma_t\right)^2 \right. \\ & \left. + V_C \left(\beta_t \left(\mu_t - \mu_t^*\right) + a \frac{1}{2} \left(\beta_t \sigma_t\right)^2 + g(\mu_t^*) \right) \right] \end{aligned} \tag{27}$$

A candidate value function for the agent, suggested by the analysis of the previous section, is $V(t, C_t, G_t) = -\exp\left(-a\left(C_t - G_t\right)\right)$. Using that candidate, we see that the second-order conditions are met by assumption, that

$$\mu_t^* = \arg\max_{\mu_t} \beta_t \mu_t - g(\mu_t) \tag{28}$$

and that the right-hand side of (27) equals zero at the optimum. This is both necessary and sufficient for optimality, and we make a formal statement of this below.

The optimal salary representation in (26) reduces the agent's dynamic problem to the repeated static problem given in (28). The agent simply maximizes his utility/contractual benefits $(\beta_t \mu_t)$ minus his cost of effort $g(\mu_t)$.

4.2 Implementation

If we put together the results above, we arrive at the following theorem:

Theorem 1 ***[Optimal Contracts]***: *Assume a given contract C_T solves the principal's problem. Then the contract implements μ_t^* if and only if C_T is the terminal value of the C_t process with $\hat{U} = -\exp\left(-aC_0\right)$ and*

$$dC_t = g(\mu_t^*)dt + a\frac{1}{2}\beta^2(t, \mathcal{Y}_t)\sigma_t^2 dt + \beta(t, \mathcal{Y}_t)\left(dY_t - \mu_t^* dt\right) \tag{29}$$

for which $\beta(t, \mathcal{Y}_t)$ *and* μ_t^* *are related by*

$$\mu_t^* = \arg\max_{\mu_t} \beta(t, \mathcal{Y}_t)\mu_t - g(\mu_t) \tag{30}$$

Furthermore,

$$\mathrm{E}^{\mathbb{A}}\left[-\exp\left(-a\left(C_T - G_T\right)\right)|\mathcal{B}_t,\ \mu = \mu^*\right] = -\exp\left(-a\left(C_t^* - G_t^*\right)\right) \tag{31}$$

The interpretations of theorem 1 are all contained in section 2.

Commitment: Disappointment and Firing

A key feature of the optimal contract of theorem 1 is that it does not rely on as much commitment as the model in section 3 specified. In particular, the general setup can accommodate early termination. As long as the principal can commit to paying the agent C_t at the time of termination, (31) says that the agent is always indifferent to staying, quitting, or being fired. Thus, there is no concern that the principal will be unable to fire the agent if the principal's valuation of the contract becomes negative over time.

4.3 The Principal's Problem

The principal's problem (23) is to find

$$\begin{array}{rcl} C_T^* & \in & \arg\max_{C_T} \mathrm{E}^{\mathbb{P}}\left[-\exp\left(-A\left(Y_T - C_T\right)\right)\right] \\ \text{s.t.} & \text{(i)} & \mu[C] \text{ solves the agent's problem} \\ & \text{(ii)} & dY_t = \left(\mu_t[C] - \sigma_t\delta_t^P\right)dt + \sigma_t dB_t^{\mathbb{P}} \\ & \text{(iii)} & \mathrm{E}^{\mathbb{A}}\left[-\exp\left(-a\left(C_T - G_T\right)\right)\right]\Big|_{\mu=\mu[C]} \geq \hat{U} \end{array}$$

Theorem 1 shows how to construct any optimal contract around the principal's choice of β_t, but only under the agent's measure $\mathbb{A}$. To proceed, we must evaluate the contract under the principal's measure $\mathbb{P}$. Fortunately, while the principal and agent disagree about the evolution of B_t, Y_t is observable to both. Thus the evolution of the payment given in (29) is an agreed upon quantity. Since the principal assumes the agent's actions are optimal, we substitute $\mu = \mu^*$ into (29), which becomes

$$dC_t = g(\mu_t^*)dt + a\frac{1}{2}\beta_t^2\sigma_t^2 dt - \beta_t\sigma_t\delta_t dt + \beta_t\sigma_t dB_t^{\mathbb{P}} \tag{32}$$

The principal can "observe" δ and $B_t^{\mathbb{P}}$ in equilibrium because the principal solves his problem under the assumption that the agent behaves optimally (23i); so, $\mu = \mu^*$ and $dB_t^{\mathbb{P}}$ is observable from dY_t.

Belief heterogeneity is reflected in the term $-\beta_t\sigma_t\delta_t$. This represents the difference in the principal's and agent's assessments of the payment process due to their different measures. When the agent is optimistic relative to the principal ($\delta_t > 0$), it means that the agent believes the underlying profitability of the project is high. Since the agent is paid a portion of that project ($\beta_t dY_t$), the agent believes that his payment will be high. The principal, who has a lower assessment of the project's profitability, has a correspondingly lower belief about the value the agent's final payment will take. So while the project's evolution and the agent's final payment are observable and agreeable, the principal and agent disagree about the expected value those items will have.

The principal's and agent's relative views on the value of their contracts allow a side-bet, through the contract, about the outcome of the project. When the agent is relatively optimistic ($\delta_t > 0$) and β_t is high, the agent receives a relatively high fraction of the project and a relatively high value from optimism about the project, and the principal allocates to himself to make the agent's participation constraint bind exactly. In effect, the principal and agent are engaging in the constrained trade of consumption in different states, with the principal being able to allocate all the gains from trade to himself. We follow the literature and refer to this as a "side-bet".

Because the amount of value created by these side-bets is endogenous – it depends on the principal's choice of β_t – the principal will have an incentive to alter the contract to increase the value of these bets. In that sense, differences in beliefs, from the principal's perspective, are similar to endogenous changes in the cost function. Consider an example in which $\mu \in M = [0, K]$ for K large. Then the agent will choose $g'(\mu_t^*) = \beta_t$ and the value to the principal of heterogeneous beliefs is $-g'(\mu_t^*)\sigma_t\delta_t$. The total cost (ignoring the cost of insuring the agent, which is present with homogeneous beliefs) to the principal is then $g(\mu_t^*) - g'(\mu_t^*)\sigma_t\delta_t$. So as the agent becomes more optimistic (δ_t increases), the convexity of

g implies that the marginal cost of implementing a particular choice of μ^* declines.[8]

Lastly, we deal with observability. The principal knows the agent's optimal control (as a function of $B_t^{\mathbb{A}}$) for any contract he writes, because the principal can solve the agent's problem. However, knowing the agent's optimal control allows the principal to use his observations of Y to infer the path and value of $B_t^{\mathbb{P}}$, under the assumption – which the principal makes and the contract ensures – that the agent actually uses the optimal control. Since the principal can "observe" $dB_t^{\mathbb{P}}$ in equilibrium, the principal can write an optimal contract based on any quantity, including δ_t that depends on $dB_t^{\mathbb{P}}$. This leads us to the principal's relaxed problem:

Theorem 2 [The Principal's Problem]: *A contract C_T is a solution to the principal's problem (23) if and only if*

$$\begin{aligned} \beta^* \quad &\in \quad \arg\max_{\beta} \mathrm{E}^{\mathbb{P}} \left[-\exp\left(-A \left(Y_T - C_T \right) \right) \right] \\ \text{s.t.} \quad \text{(i)} \quad & dY_t = \left(\mu_t^* - \sigma_t \delta_t \right) dt + \sigma_t dB_t^{\mathbb{P}} \\ \text{(ii)} \quad & dC_t = g(\mu_t^*)dt + a\frac{1}{2}\beta_t^2 \sigma_t^2 dt - \beta_t \sigma_t \delta_t dt + \beta_t \sigma_t dB_t^{\mathbb{P}} \\ \text{(iii)} \quad & \mu_t^* = \arg\max_{\mu_t} \beta_t \mu_t - g(\mu_t) \end{aligned} \tag{33}$$

where $Y_0 = 0$ and $C_0 = -\frac{1}{a} \ln(-\hat{U})$.

This is a standard optimal control problem with one choice variable: β_t. As such, it can be solved analytically in numerous well-understood ways, and numerical techniques are well developed. We solved one version of it in section 2.

4.4 Discussion

We now move on to two properties of the heterogeneous beliefs equilibrium. The first is a discussion of how to do welfare analysis with heterogeneity in beliefs. The second is the convexity of the principal's value function with respect to δ_t.

[8] This result is from the principal's perspective only. Simply substituting for g in the agent's problem will change the agent's optimum.

Welfare

When doing welfare analysis, one must be careful in choosing the probability measure used to calculate expected utility. Throughout this paper, we are consistently agnostic with respect to whether the principal and the agent are in fact correct in their evaluation of the world. Thus, we follow Savage (1954) and Anscombe and Aumann (1963) in using the participants' subjective probabilities to evaluate their welfare.

This is sensible because utility is about relative choices. In particular, we are interested in knowing the answer to this question: "what would the participants require – in terms of money – to give up the opportunity they now have?" This is a statement about how the participants value their opportunities, and so it must be taken under their own measures.

The agent obtains his reservation utility $\hat{U}$ at time 0 under his own measure. By replacing the agent's salary function (29) into the utility function and applying Ito's lemma, we can see that expected continuation utility evolves according to a martingale along the optimal path:

$$-\exp\left(-a\left(C_T^* - G_T^*\right)\right) = \hat{U} - \int_0^T a\exp\left(-a\left(C_t^* - G_t^*\right)\right)\beta_t\sigma_t dB_t^{\mathbb{A}}$$

So, while the evolution of the agent's utility function depends on the agent's beliefs, it does not directly depend on the principal's beliefs or on the difference between them. The principal takes the differences in beliefs into account in designing the contract, but the agent's welfare is pinned down by his outside option. The slope of the incentive contract β_t does generally depend on the differences in beliefs, but this does not affect the growth rate of the agent's utility; it affects only its variability. In equilibrium, the agent is exactly compensated for this risk induced by steeper incentive.

In our setup, the principal effectively has all the bargaining power with respect to the heterogeneity in beliefs. This is the source of the convexity of the principal's value function: the principal is able to allocate all the gains from trade (the gains from side-bets) to himself. Employing a more or less optimistic agent can change the principal's welfare drastically, but it does not change the agent's expected welfare.

Convexity

We now look across contracts to discover how optimal contracts under heterogeneous beliefs might differ from one environment to another. To do so, we are most interested in how the contract changes as δ rises and falls.

Consider a principal who is contracting with agent i with δ_t^i and agent j with δ_t^j. Assume that agent j is strictly more relatively optimistic (or the principal is relatively more pessimistic), so that $\delta_t^j - \delta_t^i = \Delta_t > 0$. Assume also that $\beta_t > 0$ and $\sigma_t > 0$. Then (6) shows us that for any given β_t, the principal will receive an additional cash flow from his contract with agent j over his contract with agent i in the amount of $\beta\sigma_t\Delta_t$. Assume that β_t^{i*} is optimal given agent i, so that changing β_t^{i*} by a small amount will result in no change in the principal's welfare from his contract with agent i. However, if the principal offers a contract with $\beta_t = \beta_t^{i*}$ to agent j, the principal is made strictly better off by increasing β_t because of the additional cash flow, $\beta_t\sigma_t\Delta_t$. We conjecture that this is true in a more general sense, so that if $\delta_t^i < \delta_t^j$, then $\beta_t^i < \beta_t^j$. The example in section 2 confirmed that this is true for at least one specification of δ_t.

What this conjecture says is that as the agent's relative optimism over the project increases, the principal will increase the severity of the agent's incentives. He does this, as noted in section 4.3, in order to increase the value of the side-bets between himself and the agent. What this conjecture also says is that the principal's gains *from contracting* with an agent are in some way convex as a function of δ. Clearly, the principal is better off as $\beta_t\sigma_t\delta_t$ increases: the principal prefers more to less. But the principal also changes his actions, so he is gaining from a substitution effect as well.

5 Conclusion

We have presented a model of contracting under heterogeneous beliefs in continuous time in a very general setting. We derive a theorem that reduces the principal-agent model with belief differences to a standard dynamic programming problem. Our main result is that the principal does not wish to implement a contract undoing belief differences. On the contrary, he takes advantage of disagreement in contract design, and differing priors can actually increase the principal's ex-ante welfare. In particular, the principal desires to sell the agent consumption in states the agent thinks are relatively more likely. In addition, the correlation between the project and beliefs can induce the principal to shift risk to the agent and induce a lower level of effort, even if current disagreement is zero.

We can go back to the question raised in the introduction of how the incentive pay of a trader should depend on his disagreement with the principal. A priori, one might have thought that the principal would want to undo disagreement by punishing the trader if he uses strategies contrary to the agent's beliefs. We do not observe such punishments in practice. In fact, traders on the same floor are free to hold contradictory beliefs. Our model provides an explanation for this observation, as we show that disagreement can increase the principal's expected payout.

A Proofs

Proof of Theorem 1. Very little besides a verification theorem is needed to make the discussion in section 4.1 rigorous. We will take the discussion and derivations in sections 4.1 and add details and missing steps.

First, the martingale representation theorem cited in section 4.1 has the additional implication that the volatility, $\phi_t \sigma_t$, is an element of $\mathcal{L}_2$. Nothing else is needed with regards to the optimal contract form in (26).

Next, we show that μ^* is indeed the agent's optimal control using a verification theorem. The form of this part of the theorem is quite standard; this version is adapted from Vayanos and Wang (2006).

Define the variable $\hat{\mathcal{V}}_t$ so that

$$\hat{\mathcal{V}}_t \quad = \quad -\exp\left(-a\left(C_t - G_t\right)\right) \tag{34}$$

for some general μ process.Here, C_t denotes the terminal consumption process for the control μ (26), and $\hat{\mathcal{V}}_T$ is the agent's final realized utility. Observe that the Hamilton-Jacobi-Bellman equation (27) can be re-written as

$$0 = \max_{\mu} \mathrm{E}^{\mathbb{A}}\left[d\hat{\mathcal{V}}_t | \mathcal{B}_t\right] \tag{35}$$

Since the right hand side of (35) achieves the maximum at zero when $\mu_t = \mu_t^*$ (28), the drift of $\hat{\mathcal{V}}_t$ is less than or equal to zero for any μ_t. Thus,

$$\hat{\mathcal{V}}_T \leq \hat{\mathcal{V}}_t + \int_t^T \beta_t a \exp\left(-a\left(C_t - G_t\right)\right)\sigma_s dB_s^{\mathbb{A}}$$

Since $\phi_t \sigma_t \in \mathcal{L}_2$, the above expression is integrable, and so we can take expectations:

$$\hat{\mathcal{V}}_t \geq \mathrm{E}^{\mathbb{A}}\left[\hat{\mathcal{V}}_T | \mathcal{B}_t\right] \tag{36}$$

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

Now, we repeat equations (34) and (36) for μ_t^*. Since the μ^* solves the maximization in

(35) with the right hand side equal to zero, the drift of $\hat{\mathcal{V}}_t$ is zero for $\mu = \mu^*$ and

$$\hat{\mathcal{V}}_t = \mathrm{E}^{\mathbb{A}}\left[\hat{\mathcal{V}}_T|\mathcal{B}_t,\, \mu = \mu^*\right] \tag{37}$$

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

The discussion in section 4.1 and the above arguments give us two statements. First, if a contract C_T implements μ_t^*, then (26) must hold. Second, if (26) holds, then (28) is true if and only if C_T implements μ_t^*. Together, these imply that C_T implements μ_t^* if and only if (26) and (28) hold.

The final equation is a re-statement of the value function and follows from the fact if $\mu = \mu^*$ is known, then $\mathcal{B}_t = \mathcal{Y}_t$. This completes the proof. ■

Proof of Theorem 2. Theorem 1 and the discussion in the text above the statement of theorem 2 are sufficient to show that a solution to the principal's original problem (23) is also a solution to the principal's revised problem (33).

For the converse: the feasible set in the principal's revised problem (33) is (weakly) contained in the feasible set for the principal's original problem (23) because contracts in the revised problem must be of the form (29) with optimal controls given by theorem 1. Since we have shown that any optimum over the larger set (23) must be in the smaller set (33) (theorem 1) and the objective is the same, then any optimum over the smaller set must also an optimum over the larger set. This completes the proof. ■

References

Acemoglu, Daron, Victor Chernozhukov, and Muhamet Yildiz, 2007, Learning and Disagreement in an Uncertain World, Working Paper, MIT.

Anscombe, Frank, and Robert Aumann, 1963, A Definition of Subjective Probability, *Annals of Mathematical Statistics* 34, 199–205.

Aumann, Robert, 1976, Agreeing to Disagree, *Annals of Statistics* 4, 1236–1239.

Baker, Malcolm, Richard Ruback, and Jeffrey Wurgler, 2006, Behavioral Corporate Finance: A Survey, in Espen Eckbo, eds.: *The Handbook of Corporate Finance* (Elsevier/North Holland,).

Benabou, Roland, and Jean Tirole, 2002, Self-Confidence and Personal Motivation, *Quarterly Journal of Economics* 117, 871–915.

Bolton, Patrick, Jose Scheinkman, and Wei Xiong, 2006, Executive Compensation and Short-Termist Behavior in Speculative Markets, *Review of Financial Studies* 73, 577–610.

Cvitanić, Jakša, and Jianfeng Zhang, 2006, Optimal Compensation with Adverse Selection and Dynamic Actions, Working Paper, SSRN.

Davis, Mark H., and Pravin Varaiya, 1973, Dynamic Programming Conditions for Partially Observable Stochastic Systems, *SIAM Journal on Control* 11, 226–261.

Gervais, Simon, John Heaton, and Terrance Odean, 2006, Overconfidence, Investment Policy, and Manager Welfare, Working Paper, Duke University.

Gul, Faruk, 1998, A Comment on Aumann's Baysian View, *Econometrica* 66, 923–927.

Holmstrom, Bengt, 1979, Moral Hazard and Observability, *RAND Journal of Economics* 10, 74–91.

Holmstrom, Bengt, and Paul Milgrom, 1987, Aggregation and Linearity in the Provision of Intertemporal Incentives, *Econometrica* 55, 303–328.

Ito, Takatoshi, 1990, Foreign Exchange Rate Expectations: Micro Survey Data, *American Economic Review* 80, 434–449.

Jenter, Dirk, 2004, Market Timing and Managerial Portfolio Decisions, *Journal of Finance* 60, 1903–1947.

Landier, Augustin, and David Thesmar, 2006, Contracting with Optimistic Entrepreneurs: Theory and Evidence, *Review of Financial Studies* forthcoming.

Liptser, Robert, and Albert Shiryaev, 2000, *Statistics of Random Processes*. (Springer-Verlag Berlin) second edn.

Milgrom, Paul, and Nancy Stokey, 1982, Information, Trade and Common Knowledge, *Journal of Economic Theory* 26, 17–27.

Mullainathan, Sendhil, and Richard Thaler, 2001, Behavioral Economics, in *International Encyclopedia of Social Sciences* (Pergamon Press,).

Odean, Terrance, 1999, Do Investors Trade Too Much?, *American Economic Review* 89, 1279–1298.

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

Sannikov, Yuliy, 2006, A Continuous Time Version of the Principal-Agent Problem, Working Paper, Stanford University.

Savage, Leonard, 1954, *The Foundations of Statistics*. (Wiley New York).

Schättler, Heinz, and Jaeyoung Sung, 1993, The First Order Approach to the Continuous-Time Principal-Agent Problem with Exponential Utility, *Journal of Economic Theory* 61, 331–371.

Van den Steen, Eric, 2001, Essays on the Managerial Implications of Differing Priors, Ph.D. Dissertation, Stanford University.

Vayanos, Dimitri, and Jiang Wang, 2006, Capital Market Theory, Manuscript in preparation, LSE and MIT.

Westerfield, Mark, 2006, Dynamic Contracts with Hidden Actions in Continuous Time, Working Paper, SSRN.

Yildiz, Muhamet, 2003, Bargaining Without a Common Prior: An Immediate Agreement Theorem, *Econometrica* 71, 793–811.