

Adrian, Tobias; Westerfield, Mark M.

Working Paper

Disagreement and learning in a dynamic contracting model

Staff Report, No. 269

Provided in Cooperation with:
Federal Reserve Bank of New York

Suggested Citation: Adrian, Tobias; Westerfield, Mark M. (2006) : Disagreement and learning in a dynamic contracting model, Staff Report, No. 269, Federal Reserve Bank of New York, New York, NY

This Version is available at:

<https://hdl.handle.net/10419/60533>

Standard-Nutzungsbedingungen:

Die Dokumente auf EconStor dürfen zu eigenen wissenschaftlichen Zwecken und zum Privatgebrauch gespeichert und kopiert werden.

Sie dürfen die Dokumente nicht für öffentliche oder kommerzielle Zwecke vervielfältigen, öffentlich ausstellen, öffentlich zugänglich machen, vertreiben oder anderweitig nutzen.

Sofern die Verfasser die Dokumente unter Open-Content-Lizenzen (insbesondere CC-Lizenzen) zur Verfügung gestellt haben sollten, gelten abweichend von diesen Nutzungsbedingungen die in der dort genannten Lizenz gewährten Nutzungsrechte.

Terms of use:

Documents in EconStor may be saved and copied for your personal and scholarly purposes.

You are not to copy documents for public or commercial purposes, to exhibit the documents publicly, to make them publicly available on the internet, or to distribute or otherwise use the documents in public.

If the documents have been made available under an Open Content Licence (especially Creative Commons Licences), you may exercise further usage rights as specified in the indicated licence.

Federal Reserve Bank of New York
Staff Reports

Disagreement and Learning in a Dynamic Contracting Model

Tobias Adrian
Mark M. Westerfield

Staff Report no. 269
December 2006
Revised May 2008

This paper presents preliminary findings and is being distributed to economists and other interested readers solely to stimulate discussion and elicit comments. The views expressed in the paper are those of the authors and are not necessarily reflective of views at the Federal Reserve Bank of New York or the Federal Reserve System. Any errors or omissions are the responsibility of the authors.

Disagreement and Learning in a Dynamic Contracting Model

Tobias Adrian and Mark M. Westerfield

Federal Reserve Bank of New York Staff Reports, no. 269

December 2006; revised May 2008

JEL classification: D0, D8, G0

Abstract

We present a dynamic contracting model in which the principal and the 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, a shift that changes incentives. In our dynamic setting, the interaction between incentive provision and learning creates an intertemporal source of “disagreement risk” that alters optimal risk sharing. An endogenous regime shift between economies with small and large belief differences is present, and an early shock to beliefs can lead to large persistent differences in variable pay even after beliefs have converged. Under risk-neutrality, “selling the firm” to the agent does not implement the first-best outcome because it precludes state-contingent trades.

Key words: dynamic contracts, heterogeneous beliefs, learning, hidden action, principal-agent, continuous time

Adrian: Federal Reserve Bank of New York (e-mail: tobias.adrian@ny.frb.org).

Westerfield: University of Southern California (e-mail: mwesterf@usc.edu). The authors thank Daron Acemoglu, Markus Brunnermeier, Douglas Gale, Ken Garbade, Zhiguo He, Bengt Holmström, Dirk Jenter, Oguzhan Ozbas, Stephen A. Ross, Berk Sensoy, David Sims, Costis Skiadas, Michael Steinberger, Jianfeng Yu, Jan Zabojsnik, and seminar participants at Northwestern University, Duke University, the University of California, Los Angeles, Princeton University, the University of Southern California, the Western Finance Association, the Society for Economic Dynamics, and the Federal Reserve Bank of New York for their helpful insights and comments. The views expressed in this paper are those of the authors and do not necessarily reflect the position of the Federal Reserve Bank of New York or the Federal Reserve System.

1 Introduction

In organizations such as financial institutions, economic agents with potentially heterogeneous beliefs interact in hierarchical and non-market-mediated environments. Those belief differences have been shown to have substantial effects on compensation and financing arrangements. For example, Landier and Thesmar (2006) show that a one standard deviation increase in an entrepreneur's expectations error results in an average increase in short term credit line usage of 20% – half the sample standard deviation. Ben-David, Graham, and Harvey (2007) show that overconfidence on the part of CFOs drives firms to invest more, pay out less in dividends, change debt ratios, and engage in market timing. In addition, the compensation packages of those CFOs contain more bonus and performance based pay, but not higher average compensation, than those of their peers. While there is a theoretical literature on the static effects of belief differences and on the effects of heterogeneous beliefs on asset markets, there is very little work on the dynamic effects of differing beliefs on contracts, particularly involving learning. The purpose of our paper is to analyze a dynamic contracting model with disagreement and learning and to show how the addition of the dynamics adds to and changes static predictions.

In a static model, the principal shifts the agent's consumption allocation to states that the agent considers relatively more likely, changing incentives. We show that in a dynamic model such consumption shifting – side-bets – creates a new source of risk – disagreement risk – that the principal optimally shares with the agent. This risk exists because the project's payoff co-varies with the value of consumption shifting; the covariance exists because of the learning process, and the risk is relevant even when the principal is risk neutral.

In a homogeneous beliefs contracting model, the contract exhibits a trade-off between in-

centives and insurance. Dynamic heterogeneous beliefs add a second trade-off between side-bets/incentives and disagreement risk. This trade-off is structured like a reward-risk trade-off, and we show that it is of first-order importance, just as learning is. The consumption shifting/disagreement risk trade-off leads to an endogenous regime shift between moderately and extremely heterogeneous beliefs and to path dependence in the contract and divergence over time in response to small initial shocks.

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 way, our model is one of heterogeneous beliefs, where the disagreement is over the project's profitability, rather than overconfidence, where one participant believes his actions are more effective than they actually are. The moral hazard part of our model is standard: we begin with the canonical exponential-normal setup of Holmström and Milgrom (1987) and generalize to the intermediate consumption model of Sannikov (2007b). The project pays its value as the agent's unobservable effort plus noise; the agent pays a cost for his effort and the principal can only observe noisy project output. The principal's payment to the agent takes place either at the end of the economy or over intermediate periods. We use these familiar contracting models so that our results can be explicitly traced back to our assumptions on beliefs and their interactions with the contracting environment.¹

In our model of beliefs, 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 each is a rejection of one of the 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 can lead to counter-factual predictions, such as the impossibility of trade in common value assets (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 model differences in beliefs 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. In addition to Landier and Thesmar (2006) and Ben-David, Graham, and Harvey (2007), mentioned above, Ito (1990) finds belief heterogeneity among currency traders within and between firms. Odean (1999) demonstrates that investors with brokerage accounts tend to trade too often and interprets the finding as evidence of belief heterogeneity. Jenter (2004) demonstrates that some managers have systematically contrarian views that affect their decision making.

The first result from our model is that there is a risk-reward tradeoff between side-bets and disagreement risk. This tradeoff exists because the total value of shifting the agent's consumption to states the agent consider relatively more likely – side-bets – is based on the level of disagreement; that level of disagreement changes over time as the participants learn. The reward to consumption shifting is an increase in overall (static) efficiency as each participant consumes in states for which that participant has relatively high probability assessments. The

risk is that the value to these side-bets changes (dynamically) in a way that is correlated with the risk from the project. As a result, there is a conflict between static and dynamic benefits.

The purely dynamic covariance effect – disagreement risk – drives many of our results, and it is present even when the principal is risk neutral. Differences in opinion co-vary with both the project’s outcome and the principal’s certainty equivalent wealth through the learning process. When the principal is risk averse but the cost of employing the agent is constant, the principal’s wealth level co-varies with both the principal’s learned assessment of the project’s future profitability and with the changing value of side-bets based on beliefs. When the principal is risk neutral but the cost of employing the agent varies with the agent’s wealth, the cost of employment co-varies with the project’s outcome which drives belief differences, the perceived future profitability of the project, and the future value of side bets. In either case, the principal’s opportunities vary jointly with the portion of variable pay offered to the agent and with beliefs, and so the principal has to consider changing beliefs as an additional risk factor. This is based on covariance and is present even in states of the economy in which the principal and agent momentarily agree.

Disagreement risk creates an endogenous regime shift in the economy as it moves from large to moderate belief differences. When belief differences are moderate, the agent faces moderate and declining (in time) variable pay, whereas when belief differences are large, the agent faces large and increasing (in time) variable pay. When belief differences are moderate, the important covariance is between the project’s output and the perceived future profitability of the project. This covariance is positive, and the principal responds to the increased risk by distorting the contract so as to push more of the project onto the agent. This effect is strongest near the beginning of the economy when future profitability varies most. Over time, the principal relaxes

the distortion, causing variable pay to decline. In contrast, when belief differences are large, the agent faces large and increasing variable pay because the important covariance is between the project's output and the future value to side-bets. This covariance is negative, and the principal responds to the reduced risk by taking more of the project for himself, reducing variable pay. This effect is also strongest near the beginning of the economy because more of the side-bets are then in the future.

In some economies, the contractual divergence over time created by disagreement risk can be larger than the convergence of beliefs over time that occurs when the participants learn. The result is that small shocks to beliefs near the beginning of the contract can push the economy across the regime shift and result in vastly different contracts later in the economy. This is the case if the underlying beliefs converge so as to undo the direct effects of the shock. Thus, there is a fundamental instability in contract forms when the participants learn.

As we show, disagreement and learning also affect the first-best outcome. In a standard risk-neutral principal-agent model, the first-best can be implemented by “selling the firm” to the agent who then chooses 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.

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.

Our paper differs significantly from papers that analyze dynamic adverse selection, such as Sannikov (2007a). In that paper, the agent has private information about the project's profitability that the principal obtains by offering the agent a menu of contracts. Once the contract is signed, the agent's information is known to the principal and the information does not change. In contrast, our model assumes that the disagreement about the project remains after the contract is signed because of heterogeneous priors or non-Bayesian learning. Furthermore, the level of disagreement changes over time because of the learning process. The changes in this residual disagreement are what generate disagreement risk.

The paper is organized as follows: Section 2 presents the basic CARA-Terminal Consumption model, while section 3 describes the optimal contract. We specialize the economy to Bayesian learning in section 4 and analyze the results. In section 5 we generalize the model to allow for intermediate consumption and more general preferences. Section 6 concludes.

2 The Basic Model

In this section, we lay out a simple model of the principal-agent economy. As a baseline, we use the well-known CARA model of Holmström and Milgrom (1987) to which we add a general disagreement or difference in beliefs between the principal and agent. The Holmström and Milgrom (1987) model is ideal for our purposes because it allows us to trace the effects we find directly back to our assumptions about beliefs. In section 5, we allow for heterogeneous beliefs in an economy with intermediate consumption and more general preferences.

2.1 Opportunities

There are two participants, a principal and an agent. The principal owns a project that he hires the agent to manage. The project's output is on $[0, T]$ and is driven by a standard Brownian motion B^Γ on a complete probability space $(\Omega, \mathcal{B}, \Gamma)$. The project pays Y_T at time T , where Y_0 is a constant and

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

where μ_t is the agent's effort level and $\sigma > 0$ is the project's (constant) volatility. The path of Y is observable to both the principal and the agent, but the path of B^Γ is observable only to the agent. \mathcal{B}_t is the augmented filtration generated by B^Γ and represents the agent's information set. \mathcal{Y}_t is the augmented filtration generated by Y and represents the principal's information set.

The agent's effort level is μ_t , which is unobservable to the principal. We restrict μ so that $\mu_t \in \mathbb{R}$ is \mathcal{B}_t -measurable and $\mu \in \mathcal{L}_1$.² The agent faces an opportunity cost for his effort which we model as a financial cost paid at T :

$$G_T = \int_0^T g(\mu_t) dt \tag{2}$$

for some three-times differentiable function g for which the integral in (2) exists. We also define

$$j(x) = g'^{-1}(x) \tag{3}$$

which is the inverse marginal cost of effort, and we assume that $g(0) = 0$, $g'(\cdot) \geq 0$, $g''(\cdot) > 0$,

and $g(j(\cdot))$ is convex.

In return for the agent's labor, the principal offers the agent a contract that specifies a terminal payment C_T , payable at time T . The principal can only observe the path of Y , and so the principal can only use observations of Y to offer the agent incentives to put forth effort. 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.

The principal and agent have exponential utility functions over terminal consumption, with differing levels of risk aversion. Thus the principal's utility function is

$$U(Y_T - C_T) = -\exp(-A(Y_T - C_T))$$

and the agent's utility function is

$$u(C_T - G_T) = -\exp(-a(C_T - G_T))$$

2.2 Beliefs

In addition to the reference probability measure Γ , we also introduce two other probability measures, \mathbb{P} for the principal and \mathbb{A} for the agent. Γ , \mathbb{P} , and \mathbb{A} are all mutually absolutely continuous (meaning they agree on zero-probability events to rule out arbitrage-like behavior), with $\Gamma = \mathbb{A}$. Let $\xi_t \equiv (d\mathbb{A}/d\mathbb{P})_t$ denote the Radon-Nikodym derivative of the probability measure \mathbb{A} , with respect to \mathbb{P} .³ $B_t^{\mathbb{P}}$ and $B_t^{\mathbb{A}}$ are Brownian motions under \mathbb{P} and \mathbb{A} respectively,

with

$$dB_t^{\mathbb{P}} = \delta_t dt + dB_t^{\mathbb{A}} \quad (4a)$$

$$\xi_t = \exp \left[-\frac{1}{2} \int_0^t \delta_s^2 ds + \int_0^t \delta_s dB_s^{\mathbb{P}} \right] \quad (4b)$$

We assume that $\delta \in \mathcal{L}_2$, and ξ_t is a martingale on $[0, T]$. We also assume that δ_0 and the functional form of the evolution of δ are known to all parties. We do not require any particular type of learning, so we can allow

$$d\delta_t = f(t, \cdot) dt + g(t, \cdot) dB_t^{\Gamma} = f(t, \cdot) dt + g(t, \cdot) dB_t^{\mathbb{A}} \quad (5)$$

for any $f(t, \cdot)$ and $g(t, \cdot)$ that are \mathcal{B}_t -measurable, integrable, and fulfill the martingale requirement for ξ_t . Later, when we assume a particular type of learning in an application, we will restrict the evolution of δ . However, as currently stated, δ simply describes the difference between the principal's and agent's priors over paths of the Brownian motion.

The statement $\Gamma = \mathbb{A}$ assigns the reference probability measure to the agent and means that the agent does not “learn in secret”, and so the agent's beliefs are not a hidden state variable.⁴ We make this assumption for both technical reasons (hidden state variable contracting models even with homogenous beliefs are not fully understood) and because it allows for more direct exposition. We interpret this to mean that either the agent is an expert who does not need to update in secret or a noise trader who is not able to do so.

From (1) and (4) we have

$$dY_t = \mu_t dt + \sigma dB_t^{\mathbb{A}} \quad (6a)$$

$$dY_t = (\mu_t - \sigma \delta_t) dt + \sigma dB_t^{\mathbb{P}} \quad (6b)$$

where the first equation expresses the agent’s beliefs about the project and the second expresses the principal’s beliefs. 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.

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 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.

2.3 The Principal’s Problem

We use subjective welfare analysis: the agent’s actions and associated objective function are taken with respect to the agent’s probability measure \mathbb{A} , while the principal’s actions and associated objective function are taken with respect to the principal’s probability measure \mathbb{P} . This reflects the fact any learning that takes place is incorporated into the subjective probability measures and the participants cannot observe any objective probability measure.

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

$$\max_{\mu} E^{\mathbb{A}} [-\exp(-a(C_T - G_T))] \geq \hat{U} \tag{7}$$

Assuming the agent accepts the contract with payment C_T , his problem is to find $\mu[C]$ so that

$$\begin{aligned} \mu[C] &\in \arg \max_{\mu} \mathbb{E}^{\mathbb{A}} [-\exp(-a(C_T - G_T))] \\ \text{s.t. } &dY_t = \mu_t dt + \sigma dB_t^{\mathbb{A}} \end{aligned} \quad (8)$$

If $\mu[C]$ solves the agent's problem for C_T , then we say that C_T *implements* $\mu[C]$.

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

$$\begin{aligned} C_T^* &\in \arg \max_{C_T} \mathbb{E}^{\mathbb{P}} [-\exp(-A(Y_T - C_T))] \\ \text{s.t. } & \text{(i) } dY_t = (\mu_t[C] - \sigma\delta_t) dt + \sigma dB_t^{\mathbb{P}} \\ & \text{(ii) } \mathbb{E}^{\mathbb{A}} [-\exp(-a(C_T - G_T))] \Big|_{\mu=\mu[C]} \geq \hat{U} \end{aligned} \quad (9)$$

Constraint (9i) is the dynamic incentive compatibility constraint, while (9ii) is a participation constraint. As written, the agent's participation constraint is only enforced at time 0, when the contract is signed. We show in future sections how this concept can be extended to allow to voluntary termination of the contract at times $t > 0$.

3 The Optimal Contract

In this section, we describe the optimal contract as a function of the agent's choice of effort and differences in beliefs. The result is a flexible contract form that reduces the principal's problem (9) to a basic dynamic programming problem.

3.1 The First-Best

Before analyzing the second-best contracting problem, let us consider the (static) first-best in which there is no information asymmetry and the principal can simply dictate the agent's choice of μ_t . The principal then chooses a sharing rule C_T and an optimal effort level μ so as to maximize a weighted sum of expected utilities (the central planner's problem):

$$C_T^*, \mu^* \in \arg \max_{C_T, \mu} \mathbb{E}^{\mathbb{P}} [-\exp(-A(Y_T - C_T))] + \lambda \mathbb{E}^{\mathbb{A}} [-\exp(-a(C_T - G_T))] \quad (10)$$

λ is chosen so as to meet the agent's participation constraint (7). From the definition of \mathbb{P} and \mathbb{A} , maximizing (10) is the same as maximizing

$$\mathbb{E}^{\mathbb{P}} [-\exp(-A(Y_T - C_T)) - \lambda \xi_T \exp(-a(C_T - G_T))]$$

which can be done state-by-state. The first order conditions are

$$g'(\mu_t^*) = 1 \quad (11a)$$

$$C_T^* = \tilde{\lambda} + G_T^* + \frac{A}{A+a} (Y_T^* - G_T^*) + \frac{1}{A+a} \ln(\xi_T) \quad (11b)$$

With the exception of the term with ξ_t , both expressions above are completely standard in homogeneous beliefs models: the principal chooses the effort level that maximizes social surplus, and the agent's consumption has a constant term from the participation constraint as well as first-best risk sharing of the total surplus from the project, $Y_T - G_T$.

When the principal and agent disagree, however, there is a new term that represents *side-bets* between the principal and agent. Because they disagree, consumption in a given state is optimally shifted towards the participant who thinks that state is more likely. When $\xi_T > 1$, it

means that the observed path of B or Y had a higher probability in the agents's beliefs than in the principal's beliefs, and so consumption is shifted towards the agent. The reverse holds when $\xi_T < 1$. Intuitively, differences in beliefs create gains from trade that the principal realizes by transferring consumption based on relative beliefs.

One important implication of the first-best sharing rule is that one can never achieve the first-best by "selling the firm to the agent". This is because doing so implies that the principal and agent are not engaging in any side-bets or belief based trade and so the gains from that trade are never realized.

It is also the case that the first-best fails to exist when both the principal and agent are risk neutral ($A = a = 0$). In this case, the optimal level of side-bets is unbounded. In fact, we can see that the level of side-bets (risk sharing based on beliefs) declines in both the risk aversion of the principal and the risk aversion of the agent; but, the risk sharing payment to the agent based on project output increases in the principal's risk aversion. Total project risk is exogenously fixed by production technology and can only be transferred, but "belief risk" is endogenous to the sharing rule and can be created and destroyed.

While the first-best formulas in this section are simple, they cannot answer two important questions: How does the sharing rule respond to shocks in a dynamic model, and how does the information asymmetry in the second-best affect optimal risk sharing? For example, given the learning process, is there a difference between an innovation at time $t = T/3$ and $t = 2T/3$? We address these questions in the next sections.

3.2 The Dynamic Contract

To look at the dynamics of C_T , we need to express C_T as a function of innovations in the project's output. We make C_T the terminal value of the process C_t , guided by the agent's certainty equivalent wealth: $-e^{-a(C_t - G_t)} = E_t^{\mathbb{A}} [-e^{-a(C_T - G_T)}]$. Then C_t is an "accrual" process that can be thought of as the time t value of a fund the agent receives at time T and to which the principal adds or subtracts as he observes innovations in Y .

Proposition 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 $\hat{U} = -\exp(-aC_0)$ and*

$$dC_t = g(\mu_t^*)dt + a\frac{1}{2}\beta^2(t, \mathcal{Y}_t)\sigma^2dt + \beta(t, \mathcal{Y}_t)(dY_t - \mu_t^*dt) \quad (12)$$

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

$$\mu_t^* = j(\beta(t, \mathcal{Y}_t)) \quad (13)$$

Furthermore,

$$E^{\mathbb{A}} [-\exp(-a(C_T - G_T)) | \mathcal{B}_t, \mu = \mu^*] = -\exp(-a(C_t^* - G_t^*)) \quad (14)$$

When μ_t is drawn from a discrete or bounded set, then (13) becomes

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

which can be substituted into (12). This changes μ_t^* from an explicit to an implicit function of β_t .

The contract is entirely stated in terms of variables observable to the principal: the agent is reimbursed for the equilibrium cost of effort ($g(\mu_t^*)$) and receives an insurance payment based on the level of risk in the contract ($\frac{1}{2}a\beta_t^2\sigma^2$). In addition, β_t represents the principal's choice of variable pay – how much of any surprise innovation in Y_t under \mathbb{A} is paid to the agent – and it implements agent's incentive compatible level of effort through (13). As is standard, the more variable pay the agent receives, the more effort the agent puts forth.

3.2.1 Commitment: Disappointment and Firing

A key feature of the optimal contract of proposition 1 is that it does not require as much commitment as the model in section 2 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, (14) 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. In addition, there is never a motivation for the agent to quit.

This flexibility in commitment is possible because C_t rewards the agent for costs *at the moment they are incurred*. Thus, potentially changing state variables enter the principal's problem, but not the agent's incentive compatibility constraint: (13) treats the agent's maximization decision as a repeated static problem. In addition, without a separation payment, either the principal or the agent would desire early termination with probability 1. Specifically, whenever $C_t > 0$ for t near T , the principal could gain by firing the agent without compensation, and

whenever $C_t < 0$, the agent could gain by quitting.

3.3 The Principal's Problem

Proposition 1 shows how to construct any optimal contract around the principal's choice of β_t , but only as a function of shocks $(dY_t - \mu_t^* dt)$ that are a surprise to the agent. Since the principal and the agent disagree, the principal believes that part of those shocks are predictable. 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 (12) is an agreed upon quantity. Since the principal assumes the agent's actions are optimal, we substitute $\mu = j(\beta_t)$ and $dY_t = (\mu_t^* - \sigma\delta_t) dt + \sigma dB_t^{\mathbb{P}}$ into (12), which becomes

$$dC_t = g(j(\beta_t))dt + a\frac{1}{2}\beta_t^2\sigma^2dt - \beta_t\sigma\delta_tdt + \beta_t\sigma dB_t^{\mathbb{P}} \quad (15)$$

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 (9i); so, for the purposes of the principal's problem, $\mu = j(\beta_t)$ and $dB_t^{\mathbb{P}}$ is observable from dY_t .⁵

It is important to note that the objective probability measure does not appear in (15); the objective probability measure is not even necessary to define the principal's problem. In fact, δ_t appears in the contract only as the difference in beliefs between the principal and agent – the portion of $dY_t - \mu_t^* dt$ that the principal believes to be predictable. 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 principal and the agent, but their actions are determined by the difference in their beliefs.

Belief heterogeneity is reflected in the term $-\beta_t\sigma\delta_t$. 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 these items have.

The principal's and agent's relative beliefs about the value of the project allow for side-bets, 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. The principal allocates a fixed payment (over dt) to himself so that the agent's participation constraint binds exactly. In effect, the principal and agent engage in a constrained trade of consumption in different states, the gains from which are entirely captured by the principal.

Because the amount of value created by these side-bets is endogenous – it depends on the principal's choice of β_t – the principal has reason to alter the contract to increase the value of these bets. In fact, from the contract in (15), one can see that as the agent becomes more optimistic (δ_t increases), the marginal cost of implementing a particular choice of μ^* declines. This leads us to the principal's relaxed problem:

Proposition 2 [The Principal's Problem]: *A contract C_T is a solution to the principal's*

problem (9) if and only if

$$\beta^* \in \arg \max_{\beta} \mathbb{E}^{\mathbb{P}} [-\exp(-A(Y_T - C_T))] \quad (16)$$

$$\text{s.t. (i) } dY_t = (j(\beta_t) - \sigma\delta_t) dt + \sigma dB_t^{\mathbb{P}}$$

$$\text{(ii) } dC_t = g(j(\beta_t))dt + a\frac{1}{2}\beta_t^2\sigma^2dt - \beta_t\sigma\delta_tdt + \beta_t\sigma dB_t^{\mathbb{P}}$$

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 several different ways, and numerical techniques are well developed. We solve one version of it – the case of Bayesian learning from heterogeneous priors – in the next section.

4 Bayesian Learning

In this section, we look at the model of sections 2 and 3 under the specific assumptions of one-sided Bayesian learning and quadratic costs.

4.1 The Principal's Problem

We assume that 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 use a model of one-sided Bayesian learning after the contract is signed. This means that the principal is a Bayesian learner whose priors about project profitability at time 0 are Gaussian with a mean of $\sigma\delta_0$ and a variance of $\sigma^2\gamma_0$. The agent, in contrast, does not update his beliefs. This may be because the agent is perfectly informed or because the agent believes that he is perfectly informed. In the former case, the agent has no need to update, while in the latter case the agent's bias prevents updating. In either case, however, the agent's effort adds value to the project.

In addition, we assume that the agent faces a quadratic financial cost of effort⁶, $g(\mu_t) = \frac{1}{2}\mu_t^2$.

The agent's incentive compatibility constraint (13) implies that

$$\mu_t^* = \beta_t$$

The principal uses Bayesian updating, so the prior means and uncertainties evolve as

$$d\delta_t = -\frac{\gamma_t}{\sigma}(dY_t - \mu_t^* + \sigma\delta_t) = -\gamma_t dB_t^{\mathbb{P}} \quad (17a)$$

$$d\gamma_t = -\gamma_t^2 dt \quad (17b)$$

according to the Kalman-Bucy filter, as presented in Liptser and Shiryaev (2000).

To solve the principal's relaxed problem (16), we use dynamic programming. After substituting in $\mu_t^* = \beta_t$, we have

$$dY_t - dC_t = \left[\beta_t - \sigma\delta_t - \frac{1}{2}\beta_t^2 + \beta_t\sigma\delta_t - a\frac{1}{2}\beta_t^2\sigma^2 \right] dt + (1 - \beta_t)\sigma dB_t^{\mathbb{P}}$$

The principal's value function, which we later verify, has the form

$$V(t, Y_t - C_t, \delta_t) = -\exp\left(-A(Y_t - C_t + F(t) + G(t)\delta_t + H(t)\delta_t^2)\right) \quad (18)$$

with boundary condition $V(T, Y_T - C_T, \delta_T) = -\exp(-A(Y_T - C_T))$.

The Hamilton-Jacobi-Bellman equation is

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

and so the principal's optimal choice of variable pay is

$$\beta_t^* = \frac{1 + A\sigma^2}{1 + a\sigma^2 + A\sigma^2} + \frac{\sigma\delta_t - A\sigma\gamma_t(G(t) + H(t)2\delta_t)}{1 + a\sigma^2 + A\sigma^2} \quad (19)$$

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

$$\begin{aligned} H'(t) &= 2\gamma(t)\frac{A\sigma^2 H(t)}{1 + a\sigma^2 + A\sigma^2} + 2\gamma(t)^2 A\frac{(1 + a\sigma^2)H(t)^2}{1 + a\sigma^2 + A\sigma^2} - \frac{\sigma^2}{2(1 + a\sigma^2 + A\sigma^2)} \\ G'(t) &= 2A\gamma(t)^2\frac{1 + a\sigma^2}{1 + a\sigma^2 + A\sigma^2}G(t)H(t) + \sigma^2\frac{a\sigma - 2\sigma Aa\gamma(t)H(t) + A\gamma(t)G(t)}{1 + a\sigma^2 + A\sigma^2} \\ F'(t) &= -H(t)\gamma(t)^2 + \frac{-2AG(t)\gamma(t)a\sigma^3 + G(t)^2 A\gamma(t)^2(1 + a\sigma^2)}{2(1 + a\sigma^2 + A\sigma^2)} - \frac{1 + A\sigma^2 - Aa\sigma^4}{2(1 + a\sigma^2 + A\sigma^2)} \end{aligned}$$

where $\gamma(t) = \frac{\gamma_0}{\gamma_0 t + 1}$. The ODEs can be easily solved with the boundary condition $F(T) = G(T) = H(T) = 0$. We interpret the solution to the principal's optimal control choice and welfare function in the next section.

4.2 The Solution

4.2.1 Side-bets, Disagreement Risk and Effort

The key finding is that both dynamic and static differences in beliefs determine the incentive contract. Relative to Holmström and Milgrom's (1987) baseline model (no differences of opinion), our expression for the slope of the contract (19) contains two extra terms:

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

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, $-A\sigma\gamma_t(G(t) + H(t)2\delta_t)$, is a covariance effect – disagreement risk. Differences of opinion (δ_t) co-vary with the project’s outcome (Y_t) and with the principal’s certainty equivalent value from the remainder of the contract. This covariance 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.

Notice that the covariance 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 covariance effect does not because it is based on future changes. For the same reason, the covariance effect is weaker as one moves closer to the completion of the contract. The covariance effect is illustrated in Figures 1 and 2.

< INSERT FIGURE 1 ABOUT HERE >

< INSERT FIGURE 2 ABOUT HERE >

The key fact illustrated in Figures 1 and 2 is that there is a regime shift when ones moves from moderate belief differences to large ones. When belief differences are moderate (δ_t is small), the agent’s variable pay is small and decreases over time. When belief differences are large, the agent’s variable pay is large and increases over time.

The regime shift happens because 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, and a positive shock to the project increases the principal’s assessment of future profitability. Because the covariance between beliefs and certainty equivalent wealth is positive ($G(t) < 0$), the covariance reflects additional risk, and so the principal pushes more of the project onto the agent in response.

When δ_t is large, the contract dominates, and a positive shock to the project pushes beliefs closer together and reduces the principal’s value to future consumption shifting. Because the covariance between beliefs and certainty equivalent wealth is negative ($H(t) > 0$), the covariance reflects a reduction in risk, and so the principal keeps more of the project for himself. Both effects decline over time because future covariances becomes less important as the contract nears completion.

4.2.2 Disagreement Risk and Non-Convergence

Under our model of Bayesian learning, beliefs converge over time; however, the regime shift created by disagreement risk can cause the contract to fail to converge. The rationale is illustrated in Figure 2: in the cross section of economies, large belief differences lead to large and increasing variable pay, while small belief differences lead to small and declining variable pay. In some economies, this divergence in variable pay as a function of risk and time can actually dominate the convergence in beliefs that is simultaneously taking place. The result is illustrated in Figures 3, 4 and 5. These plots are impulse response functions: they plot the path of the expected value of β_t or δ_t given a shock to δ_0 under the agent’s measure \mathbb{A} and the principal’s measure \mathbb{P} .

< INSERT FIGURE 3 ABOUT HERE >

< INSERT FIGURE 4 ABOUT HERE >

< INSERT FIGURE 5 ABOUT HERE >

In Figure 3, beliefs converge under the agent’s measure because the agent believes the principal will “learn to agree”. Under the principal’s measure, there is no such convergence because the principal knows that the agent is not updating his beliefs.

In Figures 4 and 5 we see the effects on the contract from a shock to δ_0 . Under the agent's measure, if we held time constant and allowed beliefs to converge, the level of variable pay would converge as well. However, the regime shift from disagreement risk that we illustrated earlier means that if we held beliefs constant and varied time only, the level of variable pay would diverge. Putting these two effects together can cause them to roughly cancel, giving the long term persistence pattern that we see under the agent's measure. Under the principal's measure, beliefs do not converge in expectation, and so we see only the effect of contractual divergence over time.

These plots indicate that in some simple and apparently well understood contracting models, early shocks to the economy can cause contracts to be radically different even much later. This result indicates that the pattern of beliefs can have large impacts on contractual relationships even after these belief differences have declined.

One might think that the non-convergence result would be diminished by taking T to be very large, but this is not the case. It is true that when T is very large and t is near T that we may see some convergence. However when t is small, the non-convergence is stronger, and the overall pattern is a more extreme version of Figures 4 and 5. This result can be understood from its source: disagreement risk exists because the principal's certainty equivalent wealth is made up in part from the project's future profitability and in part from the future value of the contract, and these both co-vary with the principal's existing wealth. When the economy becomes longer, the magnitude of the principal's certainty equivalent wealth and its co-variances become larger. Thus, disagreement risk is more extreme and the dispersion in β_t is larger as well.

4.2.3 Convexity and Uncertainty

We now examine the principal's certainty equivalent wealth from the contract. To do so, we include value from differences in beliefs and from the project but exclude value from the level of the agent's participation constraint and from the level of Y_0 . Since the level of the agent's participation constraint is captured by $C_0 = -\frac{1}{a} \ln(-\hat{U})$, the desired certainty equivalent wealth (\hat{W}_t) is defined by

$$-\exp(-A(\hat{W} + Y_0 - C_0)) = \max_{C_T} \mathbb{E}[-\exp(-A(Y_T - C_T))] = V(0, Y_0 - C_0, \delta_0)$$

Because of the exponential form of the value function (18), we can write this as

$$\hat{W} = -\frac{1}{A} \ln(-V(0, Y_0 - C_0 = 0, \delta_0))$$

which is plotted in Figure 6.

< INSERT FIGURE 6 ABOUT HERE >

Two facts emerge from Figure 6. 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 flow increases as well. However, as δ_t increases, (19) and the fact that⁷ $H(t) > 0$ imply 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 attaches 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} (17a). Because the principal's value function is convex in δ_0 , he is better off with relative priors that can easily change.

4.3 Comparison to the First-Best Contract

The first-best contract has many of the interesting features of the second-best – disagreement risk and non-convergence – although we require a dynamic analysis to bring them out. Simply observing the apparently static first-best (11) is not enough. Instead, we analyze C_t as in the case of the second-best, but without imposing the incentive compatibility constraint. Then, the principal's first-best dynamic problem is to solve (9) without the incentive compatibility constraint.

Unlike in the second-best, the principal can control μ_t and β_t separately and so wealth evolves as

$$dY_t - dC_t = \left[\mu_t - \sigma\delta_t - \frac{1}{2}\mu_t^2 + \beta_t\sigma\delta_t - a\frac{1}{2}\beta_t^2\sigma^2 \right] dt + (1 - \beta_t)\sigma dB_t^{\mathbb{P}}$$

and the value function has the form

$$\hat{V}(t, Y_t - C_t, \delta_t) = -\exp\left(-A\left(Y_t - C_t + \hat{F}(t) + \hat{G}(t)\delta_t + \hat{H}(t)\delta_t^2\right)\right)$$

with boundary condition $\hat{V}(T, Y_T - C_T, \delta_T) = -\exp(-A(Y_T - C_T))$. The Hamilton-Jacobi-

Bellman equation is

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

and so the principal chooses

$$\mu_t^* = 1 \tag{20a}$$

$$\beta_t^* = \frac{A}{a + A} + \frac{\sigma\delta_t - A\sigma\gamma_t \left(\hat{G}(t) + \hat{H}(t)2\delta_t \right)}{a\sigma^2 + A\sigma^2} \tag{20b}$$

The functions $\hat{F}(t)$, $\hat{G}(t)$, and $\hat{H}(t)$ form a solvable system of ODEs with boundary condition $\hat{F}(T) = \hat{G}(T) = \hat{H}(T) = 0$, which verifies our claim about the form of the value function.

The difference between the first- and second-best can be seen from comparing the two optimal policies, (19) and (20). The first-best also has the direct and indirect effects of heterogeneous beliefs through the term $\sigma\delta_t - A\sigma\gamma_t \left(\hat{G}(t) + \hat{H}(t)2\delta_t \right)$, and the effects of disagreement risk and the various covariances run in the same direction. However, because the principal no longer has to choose μ_t^* and β_t^* in an incentive compatible way, the principal is free to offer the agent either much more or much less variable pay than would be required to motivate a desired level of effort. This is reflected in Figure 7, which shows how disagreement risk changes the level of variable pay over time.

< INSERT FIGURE 7 ABOUT HERE >

The key fact in Figure 7 is that the second-best contract is more moderate than the first-best. In the first-best, the level of variable pay diverges fast over time and the regime shift between

small and large belief differences is more extreme. This happens because when belief differences are large in the second-best, the principal would like undertake a very large number of side-bets but implement only a moderate level of effort. Simultaneously, when belief differences are small, the principal does not want to engage in side-bets, but must still implement effort. When effort is de-coupled from side-bets in the first-best, the level of variable pay becomes more extreme.

< INSERT FIGURE 8 ABOUT HERE >

< INSERT FIGURE 9 ABOUT HERE >

The de-coupling of side-bets and incentives in the first-best makes the principal's value function more convex, as illustrated in Figure 8, and it makes the non-convergence result more extreme, as illustrated in Figure 9, which can be compared to Figures 4 and 5 in the second-best. In total, the first-best has a similar dynamic structure to the second-best, but because the principal does not have to link effort and variable pay, the effects of dynamic and changing beliefs on variable pay and value are greater.

5 Intermediate Consumption

In this section we take our general model of disagreement and learning and apply it to a contracting model with intermediate consumption and effort cost. We have three purposes here: 1) demonstrate that risk aversion of the principal is *not* required for our results, particularly disagreement risk⁸, 2) show that the commitment result under heterogeneous beliefs can be generalized to allow for intermediate consumption, and 3) demonstrate that the methods we use with exponential utility can be used in the more general setting.

5.1 The Model

We maintain the same assumptions about the probability space, the project, and beliefs as in section 2. We change only preferences, the presence of intermediate consumption and costs, and the structure of the contract.

In return for the agent's labor, the principal offers the agent a contract that specifies a terminal payment C_T , payable at time T , along with a flow of intermediate consumption c . We assume that c is restricted so that $c_t \in \mathbb{R}$ and $c \in \mathcal{L}_1$. The principal is restricted to offering a contract for which C_T is \mathcal{Y}_T -measurable and c_t that are \mathcal{Y}_t -measurable.

The principal is risk neutral and acts to maximize

$$\mathbb{E}^{\mathbb{P}} \left[\int_0^T e^{-rt} (dY_t - c_t dt) - e^{-rT} C_T \right] \quad (21)$$

while the agent acts to maximize

$$\mathbb{E}^{\mathbb{A}} \left[\int_0^T e^{-rt} (u(c_t) - g(\mu_t)) dt + e^{-rT} U(C_T) \right] \quad (22)$$

and r is the shared rate of time discounting. u and U are standard three-times differentiable concave utility functions.

5.2 The Optimal Contract

We now present the optimal dynamic contract using the utility-based “accrual” process that we used in part 3, generalized to allow for the addition of intermediate consumption:

Proposition 3 [Optimal Contracts]: *Assume a given contract $\{C_T, c\}$ 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 $\hat{U} = -\exp(-aC_0)$ and

$$dC_t = -\frac{u(c_t) - g(\mu_t^*)}{U'(C_t)}dt + \frac{rU(C_t)}{U'(C_t)}dt + a(C_t)\frac{1}{2}\left(\frac{\beta_t\sigma}{U'(C_t)}\right)^2 dt + \frac{\beta_t}{U'(C_t)}(dY_t - \mu_t^*dt) \quad (23)$$

for which $\beta(t, \mathcal{Y}_t)$ and μ_t^* are related by⁹

$$\mu_t^* = j(\beta(t, \mathcal{Y}_t)) \quad (24)$$

Furthermore,

$$\begin{aligned} & \mathbb{E}^{\mathbb{A}} \left[\int_0^T e^{-rt} (u(c_t) - g(\mu_t)) dt + e^{-rT} U(C_T) \mid \mathcal{B}_t, \mu_t = \mu_t^* \right] \\ &= \int_0^t e^{-rs} (u(c_s) - g(\mu_s^*)) ds + e^{-rt} U(C_t) \end{aligned} \quad (25)$$

As in the case of terminal consumption, when μ_t is drawn from a discrete or bounded set, then (24) becomes $\mu_t^* = \arg \max_{\mu_t} \beta(t, \mathcal{Y}_t)\mu_t - g(\mu_t)$ which can be substituted into (23). This changes μ_t^* from an explicit to an implicit function of β_t .

The form of the contract is very similar to the terminal consumption case, with two exceptions. First, because the cost of effort is now a utility cost rather than a financial cost, the agent is motivated by variable utility rather than variable cash payment. We can use an alternate formulation that takes the agent's remaining expected utility as the state variable: $\mathcal{U}_t = U(C_t)$ instead of C_t . Then

$$d\mathcal{U}_t = r\mathcal{U}_t dt - u(c_t)dt + g(\mu_t^*)dt + \beta_t (dY_t - \mu_t^*dt) \quad (26)$$

and β_t is the agent's utility share of the project rather than the cash share.

Equation (26) also shows that the difference in beliefs now most directly affects the agent's expected level of utility. As the contract is written, the agent believes that his level of utility is a martingale; because the principal believes that part of the project output net of effort ($dY_t - \mu_t^* dt$) is predictable, the principal believes that the agent's level of expected utility declines over time when the agent is optimistic:

$$d\mathcal{U}_t = r\mathcal{U}_t dt - u(c_t) dt + g(\mu_t^*) dt - \beta_t \sigma \delta_t dt + \beta_t \sigma dB_t^{\mathbb{P}} \quad (27)$$

By the principal's measure, the agent overvalues variable utility by an amount $\beta_t \sigma \delta_t$, which is positive when the agent is optimistic. So, as the contract persists, the principal anticipates having to reward the agent with progressively lower levels of utility and consumption.

The second exception is that C_t now includes the agent's intermediate consumption and effort costs. In fact, this is what allows us to generalize our commitment result. Because the potential terminal payment C_t now includes the agent's intermediate consumption and cost of effort, it is still the case that C_t rewards the agent for costs and penalizes the agent for benefits at the moment the costs or benefits are incurred. So, as shown in (25), the agent is always indifferent to quitting, being fired, or continuing to manage the project, as long as C_t can be promised in the event of termination and the agent receives his continuation utility $e^{-rt}U(C_t)$ from this payment. As in the terminal consumption case, the maximization condition (24) treats the agent's maximization decision as a repeated static problem.

In the utility formulation, taking $\mathcal{U}_t = U(C_t)$ as the state variable instead of C_t , the principal acts to maximize

$$\mathbb{E}^{\mathbb{P}} \left[\int_0^T e^{-rt} (dY_t - c_t dt) - e^{-rT} U^{-1}(\mathcal{U}_T) \right]$$

while \mathcal{U}_t evolves as in (27).

We can now state the principal's problem analogously to the case with only terminal consumption:

Proposition 4 [The Principal's Problem]: *A contract C_T is a solution to the principal's problem (9) if and only if*

$$\beta^* \in \arg \max_{\beta, c} \mathbb{E}^{\mathbb{P}} \left[\int_0^T e^{-rt} (dY_t - c_t dt) - e^{-rT} U^{-1}(\mathcal{U}_T) \right] \quad (28)$$

$$\text{s.t. (i) } dY_t = (j(\beta_t) - \sigma \delta_t) dt + \sigma dB_t^{\mathbb{P}}$$

$$\text{(ii) } d\mathcal{U}_t = r\mathcal{U}_t dt - u(c_t) dt + g(j(\beta_t)) dt - \beta_t \sigma \delta_t dt + \beta_t \sigma dB_t^{\mathbb{P}}$$

where $Y_0 = 0$ and $\hat{U} = U(C_0)$.

The principal's problem is concave over \mathcal{U}_t .

The concavity of the principal's problem in \mathcal{U}_t creates a disagreement risk in the intermediate consumption model very similar to the disagreement risk in the exponential model. The reason the principal's problem is *concave* in the agent's utility is because the principal always makes *payments* that are a *convex* function of the agent's utility. If we re-label $u_t = u(c_t)$, then the principal's cost optimization problem is to maximize

$$\mathbb{E}^{\mathbb{P}} \left[\int_0^T -e^{-rt} u^{-1}(u_t) dt - e^{-rT} U^{-1}(\mathcal{U}_T) \right]$$

such that

$$d\mathcal{U}_t = r\mathcal{U}_t dt - u_t dt + g(j(\beta_t))dt - \beta_t \sigma \delta_t dt + \beta_t \sigma dB_t^{\mathbb{P}}$$

This problem amounts to maximizing a concave function ($-u^{-1}$ plus $-U^{-1}$) subject to a linear budget constraint ($d\mathcal{U}_t$).¹⁰

Economically, this concavity can be interpreted as the existence of an optimal level of wealth for the agent at which the agent can still be punished with lower utility and it is not too expensive to grant the agent utility through consumption.

So, while the principal is risk neutral with respect to his own consumption, he is effectively risk averse with respect to the agent's utility, and it is the agent's utility, along with beliefs, that is the relevant state variable for the contract. The processes \mathcal{U}_t and δ_t co-vary in the intermediate consumption problem in the same way that $Y_t - C_t$ and δ_t co-vary in the exponential utility problem.

5.3 Discussion

5.3.1 Disagreement Risk

If we make the same assumptions as in section 4 – Bayesian learning and quadratic cost of effort – the solution to the principal's problem is characterized by the principal's value function

$$e^{-rt} F(t, \mathcal{U}_t, \delta_t) = \max_{\beta, c} \mathbb{E}^{\mathbb{P}} \left[\int_t^T e^{-rs} (dY_s - c_s dt) - e^{-rT} U^{-1}(\mathcal{U}_T) \mid \mathcal{B}_t \right]$$

The value function solves the Hamilton-Jacobi-Bellman equation

$$0 = \max_{\beta_t, c_t} \left[\beta_t - c_t - rF + F_t + F_{\mathcal{U}} \left(r\mathcal{U}_t - u(c_t) + \frac{1}{2}\beta_t^2 - \beta_t\sigma\delta_t \right) + \frac{1}{2}F_{\mathcal{U}\mathcal{U}}\beta_t^2\sigma^2 + \frac{1}{2}F_{\delta\delta}\gamma_t^2 + F_{\mathcal{U}\delta}\gamma_t\sigma\beta_t \right]$$

where $e^{-rT}F(T, \mathcal{U}_T, \delta_T) = -e^{-rT}U^{-1}(\mathcal{U}_T)$ is the terminal condition. Using the first order conditions, the optimal control for β_t is given by

$$\beta_t^* = \frac{1 - F_{\mathcal{U}}\sigma\delta_t + F_{\mathcal{U}\delta}\gamma_t\sigma}{-F_{\mathcal{U}} - F_{\mathcal{U}\mathcal{U}}\sigma^2} \quad (29)$$

The key point is that disagreement risk still exists in the economy, except that it is driven by covariances with the agent's remaining expected utility (\mathcal{U}_t) rather than covariances with the principal's wealth level. In the previous sections, the principal's value function was concave over wealth. Now the principal is risk neutral, but his value function is concave over the other relevant state variable: the agent's utility. In the model with intermediate consumption, the agent's utility captures the cost of employing the agent.

Because the principal is “effectively risk averse” with respect to the agent's utility, the principal views state variables, like beliefs, as additional risk factors. The principal's certainty equivalent wealth is now made up of two factors in addition to the cost of employing the agent, as denoted by the agent's expected utility. The first factor is the principal's assessment of the project's underlying profitability. Because the principal learns, his assessment of the project's profitability positively co-varies with unexpected shocks to output. But the principal must also pay the agent for these unexpected shocks, and so the project's assessed profitability positively co-varies with the agent's remaining expected utility. The second factor is the value the principal receives from the side-bets inherent in the contract. If the agent is optimistic and

there is a positive shock to the project, then the principal learns and beliefs converge. This makes future side-bets less valuable at the same time the principal must increase the agent's expected utility, so the value to side-bets negatively co-varies with the agent's remaining utility.

Together, these two effects mean that the principal still faces disagreement risk even if he is risk neutral. The presence of intermediate consumption and cost of effort simply means that the principal faces this as a risk factor affecting the agent's utility and cost of employment rather than the principal's own wealth directly. As a result, the optimal level of variable pay has both "direct" and "indirect" parts. The direct term, $-F_U\sigma\delta_t$, represents the direct effect of side-bets. The indirect term, $F_{U\delta}\gamma_t\sigma$, represents the impact of covariation between belief differences (δ) and the agent's utility (\mathcal{U}) on the value function of the principal.

5.3.2 A Note on 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 manipulating the agent's utility evolution (26), we can see that expected continuation utility evolves

according to a martingale along the optimal path:

$$\mathbb{E}_t^{\mathbb{A}} \left[\int_0^T e^{-rs} (u(c_s) - g(\mu_s)) ds + e^{-rT} U(C_T) | \mathcal{B}_t, \mu_t = \mu_t^* \right] = \hat{U} + \int_0^t e^{-rs} \beta_s \sigma dB_s^{\mathbb{A}}$$

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. Differences in beliefs only affect the agent's marginal utility through their impact on the slope of the incentive contract, which in turn determines the volatility of the agent's utility. The principal takes the differences in beliefs into account in designing the contract, but the agent's overall 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 incentives.

In our setup, as in most standard principal-agent models, 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.

6 Conclusion

We present a model of contracting under heterogeneous beliefs in continuous time. We first imbed a general model of beliefs into a simple and well-understood contracting model to illustrate our results. We derive a proposition that reduces the principal-agent model with belief

differences to a standard dynamic programming problem. We then generalized the model of contracting to allow for intermediate consumption and more general utilities.

Our main result is that there is a risk-reward tradeoff between incentives/side-bets and disagreement risk. As the principal engage in more side bets and consumption shifting, the principal takes more risk from the changing beliefs that allow the bets. The covariance between the project and beliefs can induce the principal to shift risk to or from the agent to induce changing levels of effort. There is a regime shift across moderate and large belief differences, and small shocks to beliefs can cause large variations in contract form well after beliefs have converged.

Together, our results show that dynamic considerations are important in understanding the effects of beliefs on contracts, and that these contracts can be sensitive to belief changes as well as levels.

A Proofs

Proof of Proposition 1. Define the agent's expected utility, given that he uses $\mu = \mu^*$ and measured with respect to the principal's information set, as

$$\mathcal{V}_t \equiv \mathbb{E}^{\mathbb{A}} [-\exp(-a(C_T - G_T)) | \mathcal{Y}_t, \mu = \mu^*] \quad (30)$$

Since the contract is assumed to solve the principal's problem, the participation constraint must bind exactly, and so $\mathcal{V}_0 = \hat{U}$.

Because \mathcal{V}_t is a martingale with respect to the information set $\{\mathcal{Y}_t, \mu = \mu^*\}$ (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 with $\phi_t \sigma \in \mathcal{L}_2$ such that

$$d\mathcal{V}_t = \phi_t (dY_t - \mu_t^* dt) \quad (31)$$

where $dY_t - \mu_t^* dt$ has zero drift under the agent's beliefs.¹¹

Notice that $\mathcal{V}_T = -\exp(-a(C_T - G_T^*))$. 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 (31) and using Ito's lemma, we find that

$$a \exp(-a(C_t - G_t^*)) \left[dC_t - g(\mu_t^*) dt - \frac{1}{2} a (\text{vol}(C_t))^2 dt \right] = \phi_t (dY_t - \mu_t^* dt) \quad (32)$$

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

$$dC_t = g(\mu_t^*)dt + a\frac{1}{2}\beta_t^2\sigma^2dt + \beta_t(dY_t - \mu_t^*dt) \quad (33)$$

This establishes the form of the contract. Notice that it depends on both μ (through dY_t) and μ_t^* .

We now use the incentive compatibility constraint to find the relationship between β and μ^* . This part of the proof is a standard verification theorem, and it is adapted from Vayanos and Wang (2006).

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

$$\hat{\mathcal{V}}_t = V(t, C_t, G_t) = -\exp(-a(C_t - G_t)) \quad (34)$$

for some general μ process. Here, C_t denotes the terminal consumption process for the control μ described by (33), and $\hat{\mathcal{V}}_T$ is the agent's final realized utility. Observe that the Hamilton-Jacobi-Bellman equation can be written as

$$0 = \max_{\mu} \mathbb{E}^{\mathbb{A}} \left[d\hat{\mathcal{V}}_t | \mathcal{B}_t \right] = \max_{\mu_t} \left[V_t + V_G g(\mu_t) + \frac{1}{2} V_{CC} (\beta_t \sigma)^2 + V_C \left(\beta_t (\mu_t - \mu_t^*) + a \frac{1}{2} (\beta_t \sigma)^2 + g(\mu_t^*) \right) \right] \quad (35)$$

Observe that the second-order conditions are met and that substituting in (34) shows that

$$\mu_t^* = \arg \max_{\mu_t} \beta_t \mu_t - g(\mu_t) \quad (36)$$

achieves the maximum ($\mu_t^* = j(\beta_t)$) and sets the right-hand side of (35) equals zero. Thus, the

drift of $\hat{\mathcal{V}}_t$ is less than or equal to zero for any μ_t , and so

$$\hat{\mathcal{V}}_T \leq \hat{\mathcal{V}}_t + \int_t^T \beta_s a \exp(-a(C_s - G_s)) \sigma dB_s^{\mathbb{A}} \quad (37)$$

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

$$\hat{\mathcal{V}}_t \geq \mathbb{E}^{\mathbb{A}} \left[\hat{\mathcal{V}}_T | \mathcal{B}_t \right] \quad (38)$$

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 (38) 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

$$\hat{\mathcal{V}}_t = \mathbb{E}^{\mathbb{A}} \left[\hat{\mathcal{V}}_T | \mathcal{B}_t, \mu = \mu^* \right] \quad (39)$$

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 is unique.

We now have two statements: First, if a contract C_T implements μ_t^* , then (33) must hold. Second, if (33) holds, then (36) is true if and only if C_T implements μ_t^* . Together, these imply that C_T implements μ_t^* if and only if (33) and (36) hold.

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

Proof of Proposition 2. Proposition 1 is sufficient to show that a solution to the principal's original problem (9) is also a solution to the principal's relaxed problem (16).

For the converse: the feasible set in the principal's relaxed problem (16) is (weakly) contained in the feasible set for the principal's original problem (9). Since we have shown that any optimum over the larger set (9) must be in the smaller set (16) (proposition 1) and the objective is the same, then any optimum over the smaller set must also an optimum over the larger set.

■

Proof of Proposition 3. Define the agent's expected utility under the contract $\{C_T, c\}$, given that the agent uses the control μ , as

$$\mathcal{V}_t^\mu \equiv \mathbb{E}^{\mathbb{A}, \mu} \left[\int_0^T e^{-rt} (u(c_t) - g(\mu_t)) dt + e^{-rT} U(C_T) \mid \mathcal{B}_t \right] \quad (40a)$$

$$= \int_0^t e^{-rs} (u(c_s) - g(\mu_s)) ds + e^{-rt} \mathcal{U}_t^\mu \quad (40b)$$

\mathcal{U}_t is the agent's remaining expected utility. The expectation $\mathbb{E}^{\mathbb{A}, \mu}$ is taken under the agent's beliefs and the probability measure induced by μ , so that the project evolves as in (6a). Since the contract is assumed to solve the principal's problem, the participation constraint must bind exactly, and so $\mathcal{V}_0 = \mathcal{U}_0 = \hat{U}$.

Since \mathcal{V}_t^μ is a martingale (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 β_t process with $\beta_t \sigma \in \mathcal{L}_2$ so that

$$d\mathcal{V}_t^\mu = e^{-rt} \beta_t (dY_t - \mu_t dt) \quad (41)$$

where $\frac{1}{\sigma} (dY_t - \mu_t^* dt)$ is a Brownian motion under the measure induced by \mathbb{A} and μ . Substituting (41) into (40b) and using Ito's lemma yields

$$d\mathcal{U}_t^\mu = r\mathcal{U}_t^\mu dt - u(c_t)dt + g(\mu_t)dt + \beta_t (dY_t - \mu_t dt) \quad (42)$$

A key fact is that the value of \mathcal{U}_t^μ depends only on the values of $\{C_T, c\}$ and μ after time t , since \mathcal{U}_t^μ is the agent's continuation value.

Next we show that the optimal strategy for the agent is μ^* with

$$\mu_t^* = \arg \max_{\mu} \beta_t \mu_t - g(\mu_t) \quad (43)$$

To see this, consider the value received by the agent for alternate strategy in which the agent uses μ before time t and then switches to μ^* after time t . This value is

$$\hat{\mathcal{V}}_t = \int_0^t e^{-rs} (u(c_s) - g(\mu_s)) ds + e^{-rt} \mathcal{U}_t^{\mu^*} \quad (44)$$

Using Ito's lemma and combining (44) and (42), we have

$$\begin{aligned} d\hat{\mathcal{V}}_t &= e^{-rt} (g(\mu_t^*) - g(\mu_t)) dt + e^{-rt} \beta_t (dY_t - \mu_t^* dt) \\ &= e^{-rt} (g(\mu_t^*) - g(\mu_t) - \beta_t \mu_t^* + \beta_t \mu_t) dt + e^{-rt} \beta_t (dY_t - \mu_t dt) \end{aligned} \quad (45)$$

Notice that $(dY_t - \mu_t dt)$ has drift zero when the agent uses the control μ . In addition, from the definition of μ^* (43), it is the case that $g(\mu_t^*) - g(\mu_t) - \beta_t \mu_t^* + \beta_t \mu_t < 0$. Thus

$$\mathcal{V}_0^\mu = \mathbb{E}^{\mathbb{A}, \mu} [\hat{\mathcal{V}}_T] = \hat{\mathcal{V}}_0 + \mathbb{E}^{\mathbb{A}, \mu} \left[\int_0^T d\hat{\mathcal{V}}_t \right] \leq \hat{\mathcal{V}}_0 = \mathcal{V}_0^{\mu^*}$$

Since $\mathcal{V}_0^{\mu^*} \geq \mathcal{V}_0^\mu$, μ^* is at least as preferred as any other control μ . Moreover, if $\mu \neq \mu^*$ on a set of positive measure, then the inequality is strict and μ^* is strictly preferred to μ .

To continue, notice that $\mathcal{V}_T^\mu = \mathcal{U}_T^\mu = U(C_T)$. Let us define the process C_t so that $\mathcal{U}_t^{\mu^*} = U(C_t)$ and $\hat{U} = U(C_0)$. Then C_T must be the time T value of the process C_t . Substituting $\mathcal{U}_t^{\mu^*} = U(C_t)$

into (42) and using Ito's lemma, find that

$$U'(C_t)dC_t + \frac{1}{2}U''(C_t)(\text{vol}(C_t))^2 dt = rU(C_t)dt - u(c_t)dt + g(\mu_t^*)dt + \beta_t(dY_t - \mu_t^*dt) \quad (46)$$

where $\text{vol}(C_t)$ is the volatility of C_t . Matching volatility, solving for dC_t , and substituting $a(C_t) = -\frac{U''(C_t)}{U'(C_t)}$, we find

$$dC_t = -\frac{u(c_t) - g(\mu_t^*)}{U'(C_t)}dt + \frac{rU(C_t)}{U'(C_t)}dt + a(C_t)\frac{1}{2}\left(\frac{\beta_t\sigma}{U'(C_t)}\right)^2 dt + \frac{\beta_t}{U'(C_t)}(dY_t - \mu_t^*dt) \quad (47)$$

We have now shown that if the contract $\{C_T, c\}$ solves the principal's problem, it implements μ^* if and only if (47) and (42) define the evolution of the state variables C_t and \mathcal{U}^{μ^*} . Moreover, it must be the case that μ_t^* and β_t are related as in (43), and the statement in the theorem follows from the definition of j as the inverse of g' . The final equation is a re-statement of (41b) with $\mu = \mu^*$ after substituting in $\mathcal{U}_t^{\mu^*} = U(C_t)$. ■

Proof of Proposition 4. The first part of the proposition follows analogously to the arguments given in the proof of proposition 2. We now demonstrate concavity:

Let $\lambda \in (0, 1)$. Let $\{\mathcal{U}_0^1, \delta_0, \sigma_0\}$ and $\{\mathcal{U}_0^3, \delta_0, \sigma_0\}$ be two initial states and $\{\beta^1, u^1\}$ and $\{\beta^3, u^3\}$ be the associated optimal policies. Consider the initial state $\{\mathcal{U}_0^2, \delta_0, \sigma_0\}$ with $\mathcal{U}_0^2 = \lambda\mathcal{U}_0^1 + (1 - \lambda)\mathcal{U}_0^3$ and the associated (possibly optimal) policy $\{\beta^2, u^2\} = \{\lambda\beta^1 + (1 - \lambda)\beta^3, \lambda u^1 + (1 - \lambda)u^3\}$, which is feasible. Then, since

$$\mathcal{U}_T = e^{rT}\mathcal{U}_0 + \int_0^T e^{r(T-t)}(-u_t + g(j(\beta_t)) - \beta_t\sigma\delta_t) dt + \int_0^T e^{r(T-t)}\beta_t\sigma dB_t^{\mathbb{P}}$$

and $g(j(\beta_t))$ is assumed to be convex, $\mathcal{U}_T^2 < \lambda\mathcal{U}_T^1 + (1 - \lambda)\mathcal{U}_T^3$. Then the concavity of $-u^{-1}$ and

$-U^{-1}$ and that $-U^{-1}$ is declining imply

$$\begin{aligned} & E \left[\int_0^T e^{-rt} (\beta_t^2 - u^{-1}(u_t^2)) dt - e^{-rT} U^{-1}(\mathcal{U}_T^2) \right] \\ > \lambda E \left[\int_0^T e^{-rt} (\beta_t^1 - u^{-1}(u_t^1)) dt - e^{-rT} U^{-1}(\mathcal{U}_T^1) \right] \\ & + (1 - \lambda) E \left[\int_0^T e^{-rt} (\beta_t^3 - u^{-1}(u_t^3)) dt - e^{-rT} U^{-1}(\mathcal{U}_T^3) \right] \end{aligned}$$

and so the value function must also be concave. ■

References

- Acemoglu, D., V. Chernozhukov, and M. Yildiz, 2008, “Fragility of Asymptotic Agreement under Bayesian Learning,” Working Paper, MIT.
- Anscombe, F., and R. Aumann, 1963, “A Definition of Subjective Probability,” *Annals of Mathematical Statistics*, 34, 199–205.
- Aumann, R., 1976, “Agreeing to Disagree,” *Annals of Statistics*, 4(6), 1236–1239.
- Baker, M., R. Ruback, and J. Wurgler, 2006, “Behavioral Corporate Finance: A Survey,” in *The Handbook of Corporate Finance*, ed. by E. Eckbo. Elsevier/North Holland, first edn.
- Ben-David, I., J. R. Graham, and C. R. Harvey, 2007, “Managerial Overconfidence and Corporate Policies,” Working Paper, University of Chicago.
- Bénabou, R., and J. Tirole, 2002, “Self-Confidence and Personal Motivation,” *Quarterly Journal of Economics*, 117(3), 871–915.
- Bolton, P., J. Scheinkman, and W. Xiong, 2006, “Executive Compensation and Short-Termist Behavior in Speculative Markets,” *Review of Financial Studies*, 73, 577–610.
- Cvitanić, J., and J. Zhang, 2007, “Optimal Compensation with Adverse Selection and Dynamic Actions,” *Mathematics and Financial Economics*, forthcoming.
- Davis, M. H., and P. Varaiya, 1973, “Dynamic Programming Conditions for Partially Observable Stochastic Systems,” *SIAM Journal on Control*, 11(2), 226–261.
- Dumas, B., A. Kurshev, and R. Uppal, 2008, “Equilibrium Portfolio Strategies in the Presence of Sentiment Risk and Excess Volatility,” *Journal of Finance*, forthcoming.

- Gervais, S., and I. Goldstein, 2007, “The Positive Effects of Biased Self-Perceptions in Firms,” *Review of Finance*, 11(3), 453–496.
- Gervais, S., J. Heaton, and T. Odean, 2006, “Overconfidence, Investment Policy, and Manager Welfare,” Working Paper, Duke University.
- Gul, F., 1998, “A Comment on Aumann’s Bayesian View,” *Econometrica*, 66(4), 923–927.
- Holmström, B., and P. Milgrom, 1987, “Aggregation and Linearity in the Provision of Intertemporal Incentives,” *Econometrica*, 55(2), 303–328.
- Ito, T., 1990, “Foreign Exchange Rate Expectations: Micro Survey Data,” *American Economic Review*, 80(3), 434–449.
- Jenter, D., 2004, “Market Timing and Managerial Portfolio Decisions,” *Journal of Finance*, 60(4), 1903–1947.
- Landier, A., and D. Thesmar, 2006, “Contracting with Optimistic Entrepreneurs: Theory and Evidence,” *Review of Financial Studies*, forthcoming.
- Liptser, R., and A. Shiryaev, 2000, *Statistics of Random Processes*. Springer-Verlag, Berlin, second edn.
- Milgrom, P., and N. Stokey, 1982, “Information, Trade and Common Knowledge,” *Journal of Economic Theory*, 26(1), 17–27.
- Mullainathan, S., and R. Thaler, 2001, “Behavioral Economics,” in *International Encyclopedia of Social Sciences*. Pergamon Press, first edn., pp. 1094–1100.

- Odean, T., 1999, “Do Investors Trade Too Much?,” *American Economic Review*, 89(5), 1279–1298.
- 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.
- Sannikov, Y., 2007a, “Agency Problems, Screening and Increasing Credit Lines,” Working Paper, New York University.
- , 2007b, “A Continuous Time Version of the Principal-Agent Problem,” *Review of Economic Studies*, forthcoming.
- Savage, L., 1954, *The Foundations of Statistics*. Wiley, New York.
- 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.
- Van den Steen, E., 2001, “Essays on the Managerial Implications of Differing Priors,” Ph.D. Dissertation, Stanford University.
- Vayanos, D., and J. Wang, 2006, “Capital Market Theory,” Manuscript in preparation, LSE and MIT.
- Westerfield, M., 2006, “Dynamic Contracts with Hidden Actions in Continuous Time,” Working Paper, University of Southern California.
- Yildiz, M., 2003, “Bargaining Without a Common Prior: An Immediate Agreement Theorem,” *Econometrica*, 71(3), 793–811.

Notes

¹Related continuous-time principal-agent models are presented in Schättler and Sung (1993), Cvitanović and Zhang (2007), and Westerfield (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), Gervais and Goldstein (2007), and Bénabou 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. A paper that examines how an agent might try to exploit heterogeneous beliefs outside a contracting environment is Dumas, Kurshev, and Uppal (2004). In addition, Bolton, Scheinkman, and Xiong (2006) study a dynamic model of CEO compensation under heterogeneous investor beliefs. Our primary contribution relative to these papers is our dynamic analysis of contracts.

²The spaces \mathcal{L}_1 and \mathcal{L}_2 are defined so that

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

³In addition, because Γ , \mathbb{P} and \mathbb{A} are mutually absolutely continuous, their augmented filtrations, \mathcal{B}_t , $\mathcal{B}_t^{\mathbb{A}}$ and $\mathcal{B}_t^{\mathbb{P}}$ agree, and so we simply write \mathcal{B}_t .

⁴This does not mean that the agent cannot mislead the principal by choosing an off-equilibrium level of effort, just that such actions do not give rise to persistent hidden states with respect to the *agent's* beliefs.

⁵Note that the principal knows the agent's optimal control (as a function of $B_t^{\mathbb{A}}$) for any contract 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. Thus, the principal can write a contract based on δ_t and $B_t^{\mathbb{P}}$ because these variables are treated as functions of dY_t and μ_t^* .

The general heterogeneous beliefs problem exhibits a separability between contracting and learning because the agent does not “learn in secret”. 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.

⁶In order to keep the example simple, we follow Holmström 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.

⁷For a proof that $H(t) > 0$ for $t < T$, observe that $H(T) = 0$ and that H is continuously differentiable. Then, since direct substitution into the ODE shows that $H(t) = 0$ implies $H'(t) < 0$, it must be that $H(t) > 0$.

⁸In the previous section, the CARA preferences meant that the cost of employing the agent was constant, and so we required $A > 0$ to obtain our results on disagreement risk. In a more general model, the cost of employing the agent is not constant, and so we can allow for a risk neutral principal and still obtain our results. As we show in this section, we only require that the principal be “effectively risk averse” over the relevant state variable. In the CARA model, that variable is wealth, whereas in a more general model that variable is the cost of employing the agent.

⁹Recall that j is the agent's inverse marginal cost of effort (3).

¹⁰See Sannikov (2007b) for a detailed explanation of how the value function behaves in continuous time models with homogenous beliefs and intermediate consumption.

¹¹Alternately, one can use the weak solution and martingale methods to obtain this result. To do so, define the probability measure \mathbb{A}^μ so that $\frac{1}{\sigma}(dY_t - \mu_t dt)$ is a Brownian motion. Then, $d\mathcal{V}_t$ is a martingale under the agent's optimal control by the principal of optimality (proved in Schättler and Sung (1993) for models with exponential utility and Sannikov (2007b) for models with intermediate consumption).

Figures

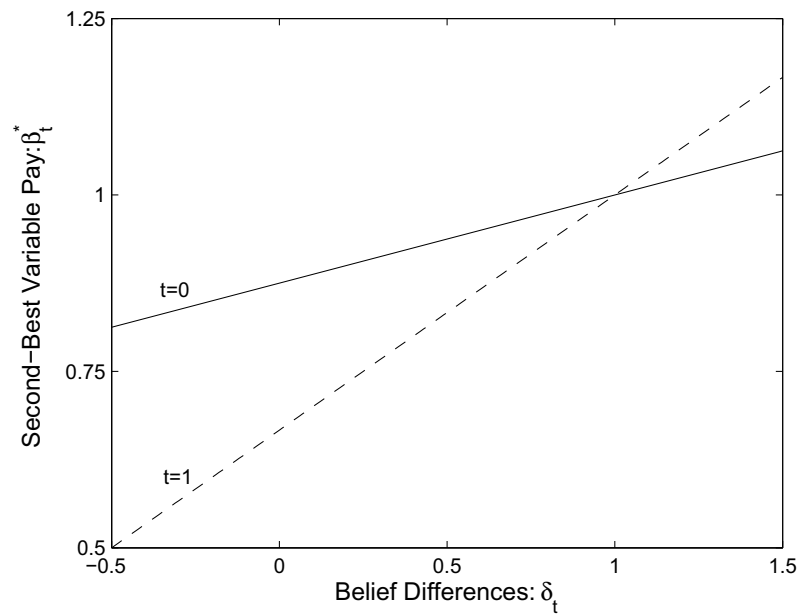


Figure 1: β_t^* as a function of δ_t for $t = 0$ and $t = 1$. We set $A = a = T = \sigma = 1$, and $\gamma_0 = 5$.




Figure 2: β_t^* as a function of time for $\delta_t = 0$ and $\delta_t = 3/2$. We set $A = a = T = \sigma = 1$ and $\gamma_0 = 5$. For reference, the value of β_t^* under homogeneous beliefs for these parameters is 0.67.




Figure 3: An impulse response function for belief differences given a shock at time 0 to beliefs.

We plot $E^A[\delta_t|\delta_0]$ and $E^P[\delta_t|\delta_0]$ for $\delta_0 = -1, 0, 1, 2$. We set $\gamma_0 = 5$.




Figure 4: An impulse response function to variable pay (β_t) given a shock at time 0 to beliefs. We plot $E^A[\beta_t|\delta_0]$ for $\delta_0 = -1, 0, 1, 2$. We set $A = T = 1$, $a = .1$, $\sigma = 2$ and $\gamma_0 = 5$.




Figure 5: An impulse response function to variable pay (β_t) given a shock at time 0 to beliefs. We plot $E^{\mathbb{P}}[\beta_t|\delta_0]$ for $\delta_0 = -1, 0, 1, 2$. We set $A = T = 1$, $a = .1$, $\sigma = 2$ and $\gamma_0 = 5$.




Figure 6: The principal's value to contracting is the function $-\frac{1}{A} \ln(-V(0, Y_0 - C_0 = 0, \delta_0))$. We plot this as a function of δ_0 for two values of γ_0 . The plot sets $T = a = A = 1$ and $\sigma = .5$.




Figure 7: β_t^* as a function of time for $\delta_t = 0$ and $\delta_t = 1.5$ in the first- and second-best. We set $A = a = T = \sigma = 1$ and $\gamma_0 = 5$. For references, the first-best value of β_t^* under these parameter with homogeneous beliefs is 0.5, while the second-best value is 0.67.




Figure 8: The principal's value to contracting is the function $-\frac{1}{A} \ln \left(-\hat{V}(0, Y_0 - C_0 = 0, \delta_0) \right)$. We plot this as a function of δ_0 for two values of γ_0 for both the first- and second-best. The plot sets $T = a = A = 1$ and $\sigma = .5$.




Figure 9: An impulse response function to variable pay (β_t) given a shock at time 0 to beliefs. We plot both $E^A[\beta_t|\delta_0]$ and $E^P[\beta_t|\delta_0]$ for $\delta_0 = -1, 0, 1, 2$. We set $A = T = 1$, $a = .1$, $\sigma = 2$ and $\gamma_0 = 5$.