

Renegotiation and Contracts with Moral Hazard and Persistent Private Information*

Bruno Strulovici

September 27, 2011

Abstract

How does renegotiation shape contracts between a principal and an agent subject to persistent private information and moral hazard? This paper introduces a concept of renegotiation-proofness, which adapts to stochastic games the concepts of weak renegotiation-proofness and internal consistency by exploiting natural comparisons across states. When the agent has exponential utility and cost of effort, each separating renegotiation-proof contract is characterized by a single “sensitivity” parameter. The optimal contract among those always causes immiserization. Reducing the agent’s cost of effort can harm the principal by increasing the tension between moral hazard and reporting problems. For the constructed contracts, truthful reporting is optimal on but also off the equilibrium path.

*Northwestern University. I am particularly grateful to Alessandro Pavan, Hector Chade, Roozbeh Hosseini, Jean Tirole, and Mehmet Ekmekci, Daniel Garrett, Christian Hellwig, Johannes Hörner, Parag Pathak, Francisco Ruiz-Aliseda, Larry Samuelson, Ilya Segal, Itai Sher, Juuso Toikka, Mike Whinston, and seminar audiences at Arizona State University, the University of Minnesota, Northwestern University, the University of Tokyo, Yale University, the Ecole Polytechnique and the Ecole HEC, Stanford University, the Northwestern CSIO-Toulouse IDEI 2011 Conference, Ohio State University, the SAET 2011 Conference, the ESSET 2011 Conference, the Hong Kong/UBC Summer Microeconomics 2011 Conference, and the Harvard/MIT theory seminar, for questions and comments.

1 Introduction

Financial bailouts, mortgage regulations, monetary and fiscal policies, and corporate pay cuts during recessions are examples where regulators, firms, and other economic agents have incentives to renegotiate implicit or formal contracts between them in reaction to their economic environment. Such renegotiation takes place as aggregate and idiosyncratic shocks unfold, and in the midst of potentially severe agency problems.

Despite its importance in economics, renegotiation has resisted numerous formalization attempts by theorists into a satisfactory concept. Indeed, even in the simpler context of repeated games, where no exogenous state exists, no consensus has emerged about the “right” concept of renegotiation. The problem would appear to become even more severe in environments, pervasive in economic applications, where agents are subject to persistent and privately observed shocks. For example, consider a firm manager who must report and transfer cash-flows to his shareholders and is thus prone to diverting cash-flows. It is natural to assume that cash-flows are serially correlated and, therefore, that past reports affect the interpretation or credibility of present reports. How does renegotiation take place in such environment? What sort of contracts survive renegotiation?

This paper argues that renegotiation can actually *simplify* the economic analysis of environments with persistent private information, making the concept both realistic and useful for the analysis. Indeed, the possibility of renegotiation prevents complicated forms of history dependence. For example, punishment phases become unavailable if it is commonly known that such phases will be renegotiated upon. In a contracting environment, this simplification also applies to the history of reports that the agent has made.

To exploit this intuition, the present paper introduces a new concept of renegotiation, which extends to stochastic games the weakest form of renegotiation in repeated games, and explores its consequences for a principal-agent model where the agent is subject to persistent private information. As most concepts of renegotiation in the literature, the definition is *set theoretic*: a contract is renegotiation-proof if it survives, after each possible history, the comparison with some set of *challengers*. For example, in repeated games, internal consistency¹ requires that the continuation payoffs never be Pareto ranked. Otherwise, players could move from a Pareto dominated continuation play to the continuation play that dominates it.

¹See Bernheim and Ray (1989).

Applied to the contracting environment of this paper, internal consistency already imposes a drastic simplification on continuation payoffs, turning the problem into a Markovian one.²

In a stochastic game, internal consistency has limited bite however, as it only applies to comparisons of continuation payoffs starting from the *same* state. A natural question then arises: can agents also make comparisons *across* states? And how do such comparisons affect the set of renegotiation-proof contracts? This paper proposes a concept of renegotiation-proofness based on such comparisons. In an abstract sense, comparisons arise when contracts starting from a given state can be transformed into contracts starting from another state. For example, a contract that provides a certain continuation utility to an agent can easily be translated or scaled so as to provide a different continuation utility to the agent. Which transformation is relevant largely depends on the context, and in any case must preserve incentive compatibility of the initial contract.

Building on and extending the setting of Williams (2011), this paper provides a thorough examination of a specific example, where the agent generates privately observed cash-flows that depend on his effort and exogenous shocks, and where the cash-flows are, other things equal, *mean reverting*. In that example, assuming exponential utility and effort cost for the agent, a natural class of contract transformation emerges, which creates a rich and well structured class of challengers. Given any incentive compatible contract, the challengers are entirely constructed from continuations of the initial contract and transformations thereof. Renegotiation proves so powerful in this context that it reduces the set of renegotiation-proof contracts to a single-parameter family, which is very easily analyzed.

Before discussing properties of these contracts, it is useful to explain further the concept of renegotiation. Suppose that there are two states of the world: “high” and “low,” corresponding to some aggregate demand for a good. Consider a contract underwritten between a firm and its supplier when the state of the world is high. If the state of the world becomes bad, the standard concept of internal consistency, or weak renegotiation proofness,³ is completely silent about how the continuation contract in that state should look like, compared to the initial contract. However, it may be possible for the parties to consider how the initial contract should look like if it had been written in the low state, and how that contract compares to the continuation contract that are currently facing, after that history. Such link

²While payoffs are easily shown to be Markovian, proving that actions are also Markovian is less straightforward. See Section 4.1.

³See Farrell and Maskin (1989).

establishes a relation between continuation payoffs across states.

In this paper, the principal has the sole initiative to propose new contracts.⁴ The viewpoint adopted here is that if the principal can come up with natural comparisons across states, then these comparisons should be taken into account as they generate sensible challengers for the continuation contracts after any state. Internal consistency presumes a minimal form of ability: the principal recognizes only histories that lead to same state, and is able to compare his continuation payoffs across any such histories. Consistency across states, the concept studied in this paper, brings an additional bite if the principal also has the ability to recognize similarities across histories leading to distinct states.

Characterization of renegotiation-proof contracts in terms of their challengers may be seen as an axiomatization of their properties. For example, starting from a given incentive compatible contract, it is often possible to achieve a different level of promised utility by simply translating or scaling (depending on the utility function) the consumption stream of the initial contract, without affecting incentive compatibility. Such transformation suggests a particular functional relation for how consumption depends on promised utility, and that relation can be justified by considering “translated” or “scaled” contracts as possible challengers of the initial contract after any history leading to the relevant level of promised utility.⁵

The paper characterizes the set of all separating, state-consistent contracts, and the optimal contract among those, assuming that the agent has exponential utility and effort cost functions. Any renegotiation-proof contract is characterized by a single “sensitivity” parameter, which determines both the agent’s incentive to truthfully report cash flows and his incentive for effort. For any such contract, all contractual variables have exact formulas as a function of the sensitivity parameter. The range of these sensitivity parameters depends only on a single parameter of the model, the risk-aversion coefficient of the agent. In particular, it is unaffected by the presence of moral hazard, compared to a pure cash-flow diversion problem.

All renegotiation-proof contracts are such that the agent wants to report cash-flows truthfully, *not only on the equilibrium path, but also after any possible deviation*. Thus, even if, at time zero, the principal has a wrong belief about the initial cash-flow, a renegotiation-

⁴In particular, the relation between renegotiation and bargaining emphasized by Abreu and Pearce (1991) and Abreu et al. (1993) does not apply here.

⁵This viewpoint is somewhat similar to axiomatizing representations of preferences in decision theory. Here, it is the contractual form that is axiomatized.

proof contract of the type studied here will induce the agent to immediately reveal the true cash-flow level.⁶

The optimal contract is obtained by maximizing a closed-form objective function with respect to the sensitivity parameter, which makes it easy to study its properties.

As a result of risk-aversion, the agent's continuation utility under the optimal contract exhibits immiserization:⁷ almost surely, it becomes arbitrarily negative as time elapses. Immiserization arises for all parameter values of the model. Intuitively, it is less costly to provide the right incentives to the agent when his continuation utility is low, because he is more sensitive to small changes in the consumption flow. Other things equal, this makes the principal prefer to provide more utility flow today and let continuation utility drift downwards, compared to the constant promised utility and utility flow arising in the first-best contract.

Furthermore, the sensitivity of the optimal contract is decreasing in the magnitude of exogenous shocks: risk-aversion makes it more costly for the principal to provide a given level of promised utility, when fluctuations are wider. A lower sensitivity parameter, in turn, reduces the principal's ability to elicit high effort from the agent and increases his incentive to underreport cash flows, other things equal, which always results in a lower payoff for the principal.

Because the agent's information is persistent, the combined presence of moral hazard and adverse selection cannot be reduced to a pure moral hazard or to a pure adverse selection problem (in contrast to what earlier literature has pointed out for the i.i.d. case). In particular, *reducing the agent's cost of effort can arbitrarily reduce the principal's second-best payoff*, even though it always improves the principal's first-best payoff. The intuition may be explained as follows. With lower effort cost, the agent arbitrarily increases his promised utility, for any fixed positive sensitivity parameter, by putting more effort. Although such effort increases the principal's cash flows, the increase in promised utility is more costly to the principal, due to concavity of the agent's utility. To offset this problem, the principal optimally reduces the sensitivity of the contract. This, however, reduces his ability to induce truthful reporting from the agent. The only way the principal can do that is by also reducing

⁶Such feature is interesting for example if the agent is a newly-arrived CEO who discovers, upon taking the job, that the financial situation of his firm is worse than what outsiders think. In such case, the contracts studied here give the agent the incentive to correctly book a nonrecurring loss on the firm's accounts.

⁷See Thomas and Worrall (1990).

the agent's marginal value of lying, which, in turn, can only be obtained by an providing arbitrarily large utility flow.⁸

As a result, a principal who could invest initially to reduce the effort cost of the agent may optimally choose to forgo this option: Even if productivity improvements can be made at little or no cost, such improvements can reduce the principal's second-best payoff.

Absent full commitment, it is well known that the revelation principle need not hold (see in particular Bester and Strausz (2001)). This paper characterizes truthful (or separating) renegotiation-proof contracts, and the optimal such contract. *A priori*, the optimal renegotiation-proof contract may involve pooling: i.e., the principal may prefer a contract where he does not learn how the cash-flows of the agent evolve. However, as discussed in Section 7, the principal has an incentive to elicit information from the agent, and cannot commit not to do so. In contrast to the setting analyzed by Laffont and Tirole (1988), the principal can repeatedly propose new contracts, and potentially learn, gradually, the type of the agent. Renegotiation harms the principal not only because it affects the agent's ex ante incentives but also, potentially, by reducing the principal's ability *not* to learn about the agent's type.

Reporting incentives are linear for any arbitrary contract, which implies that the agent is either indifferent between telling the truth and lying, or wishes to lie at maximal (infinite) rate, either upwards or downwards. This makes it necessary to model jumps in the agent's reports. A contract must specify how such jumps affect promised utility. For the renegotiation-proof contracts studied here, this contractual relation is naturally pinned down by the sensitivity parameter characterizing each of these contracts. The agent's incentives are characterized by a Hamilton-Jacobi-Bellman equation with an impulse response component, which provides a new (in the contracting literature, to the author's knowledge) and simple way to deal with the possibility of unbounded drift of the reporting process. Using this technique, it is possible to derive the agent's value function not only on the equilibrium path, but also after any possible deviation. This value is a very simple function of the agent's promised utility, of the current gap between reported and actual cash flows, and of the sensitivity parameter.

The paper contributes to the literature on renegotiation⁹ and more particularly contracting

⁸See Section 5.1 and, in particular, Equation (22).

⁹See Laffont and Tirole (1988,1990) Dewatripont (1989), and Fudenberg and Tirole (1990). Battaglini (2007) studies renegotiation with a persistent type, and finds that separating contracts are optimal for a large class of parameter values. His model has two periods and is therefore not subject to the

and renegotiation with an infinite horizon.¹⁰ Several definitions of renegotiation-proofness have been proposed for repeated games, in particular by Bernheim and Ray (1989) and Farrell and Maskin (1989), none of which is fully satisfactory.¹¹ These notions are even more problematic for stochastic games. Indeed, the weak notions (weak renegotiation proofness and internal consistency) do not allow comparisons across states.¹² An exception is Gromb (1994), who studies a binary-state dynamic model of debt contracts, and compares continuation payoffs across the two states.

The paper also contributes to the literature on dynamic contracting with persistent private information, as initiated by Fernandes and Phelan (2000).¹³ In addition to the employer-employee interpretation, the paper can be interpreted as an insurance model or an optimal taxation model¹⁴ where the agent reports income and gets subsidized or taxed according to his reports. Renegotiation here amounts to a change in regulation that the social planner cannot commit not to perform. The presence of moral hazard and persistent information combines ingredients of the literature on career concerns with contracts, as studied by Gibbons and Murphy (1990). In the model studied here, putting a higher effort today raises all future cash flows (as if the agent were investing in skills), but reduces the transfers payments to the agent, other things equal. Garrett and Pavan (2010a, 2010b) study full-commitment contracts in managerial compensation settings for which the manager's state is persistent.

Many features of the underlying model are based on Williams (2011), who focuses on full commitment and a pure reporting problem. Zhang (2008) also proposes a continuous-time model with persistent information and a binary type. Discrete-time models of contracting with persistent information include Tchistyi (2006), who considers a pure cash-flow diversion problem with binary cash flows, and Kapicka (2006) who considers a first-order approach. Fukushima and Waki (2009) propose a numerical analysis for a setting with persistent pri-

conceptual problems arising in repeated games and stochastic games with an infinite horizon.

¹⁰For limits of renegotiation-proof contracting with finite horizon, see Benoît and Krishna (1993).

¹¹See Pearce (1987), Asheim (1991), Abreu and Pearce (1991), Abreu et al. (1993), and the discussion in Section 9.

¹²These weak notions do imply Markov payoffs in principal-agent environments, as used by Bergemann and Hege (2005).

¹³The recursive approach, which plays a major role in the renegotiation concept studied here, dates back to Green (1987), Spear and Srivastava (1987), and Thomas and Worrall (1990), who introduce the use of promised utility as a state variable.

¹⁴For recent models of each type, see respectively Golosov and Tsvinsky (2006), who study a particular kind of information persistence, and Golosov et al. (2003).

vate information. Doepke and Townsend (2006) introduce a numerical method to analyze the optimal contract with moral hazard and adverse selection.¹⁵ All these papers assume full commitment. Another kind of contracting with persistent private information concerns delegated experimentation as analyzed by Bergemann and Hege (2005), Hörner and Samuelson (2010), and Garfagnini (2010). In these papers, the agent’s effort intensity to learn about the value of an action is privately observed, and may result in the principal and the agent having different beliefs about that value, hence the private information of the agent.

Section 2 introduces the principal-agent setting with moral hazard and persistent private information, and Section 3 introduces a new concept of renegotiation-proofness for stochastic games, called *state consistency*. Section 4 characterizes state-consistent contracts in the principal-agent setting. Section 5 establishes several comparative statics: i) the optimal renegotiation-proof contract always exhibits immiserization: the agent’s promised utility drifts to minus infinity almost surely, ii) increasing productivity (at not cost) can reduce the principal’s payoff, iii) increasing noise of the technology always reduces the principal’s payoff and its ability to induce effort. Section 6 shows that all state-consistent contracts are incentive compatible, even off the equilibrium path (i.e., the agent has an incentive to correct any misreport that occurred earlier). Section 7 considers the possibility issue of pooling contracts. Section 8 considers the case where the agent has a permanent outside option, and shows how the analysis of renegotiation-proofness is affected by this new constraint. Section 9 discusses extensions of the model. Proofs omitted from the main text are in the Appendix.

2 Setting

An agent generates some cash flow X_t at time t , governed by the dynamic equation

$$dX_t = [(\xi - \lambda X_t) + A_t] dt + \sigma dB_t, \quad (1)$$

where $A_t \in \mathbb{R}$ is the agent’s effort at time t and B is the standard Brownian motion. The cash flow has a mean reversion component, with speed λ and long run average ξ/λ . A low (high) mean-reversion speed λ results in high (low) persistence of the cash flows and, hence,

¹⁵In Doepke and Townsend (2006), income at any given period only depends on the agent’s action at the previous period, and hence does not exhibit the type of direct persistence studied in this and other cited papers.

of the agent's private information. Precisely, λ is the rate at which the impact of current cash-flow level on future cash flows decays over time.

The agent incurs a cost $\phi(a)$ to produce effort a , where ϕ is increasing and convex. In the computations to follow, ϕ will be exponential: $\phi(a) = \bar{\phi} \exp(\chi a)$.

The Brownian Motion B is for now the only source of exogenous uncertainty. There is a fixed a probability space (Ω, \mathcal{F}, P) satisfying the usual conditions, such that each outcome ω is identified with a path realization for B .¹⁶

The agent reports and transfers to the principal a cash flow Y_t such that

$$dY_t = dX_t + L_t dt = [(\xi - \lambda X_t) + A_t + L_t] dt + \sigma dB_t. \quad (2)$$

L_t is the rate at which the agent "lies" about the true increment dX_t of the cash-flow.¹⁷

The *gap* $G_t = Y_t - X_t$ between the reported and actual incomes Y_t and X_t satisfies

$$G_t = \int_0^t L_s ds.$$

The case of a pure reporting problem with persistent private information and no moral hazard ($a = 0$ everywhere) obtains at the limit for the exponential cost function as χ becomes arbitrarily large.¹⁸

The principal provides a consumption process C_t to the agent. A *contract* is a consumption process C adapted to the filtration generated by the report process Y .¹⁹

Whenever a recursive formulation of the contract is used with the agent's promised utility as one component of the state, a contract must also stipulate an effort process \bar{A} adapted to Y . That effort is used to compute i) the drift of the agent's continuation utility, and ii) the "innovation" (or "surprise") in the agent's reported cash-flow increment, compared to the expected increment. (See (3).)

¹⁶Section 4.1 introduces randomization across contracts. In that case, the probability space must be enlarged to account for this other source of uncertainty.

¹⁷Such lie may be unbounded. Section 6 allows the agent to report jumps in his cash flows, and propose a natural extension of the contract to this case.

¹⁸In that case, the cost function becomes flat on $(-\infty, 0]$ and arbitrarily large for $a > 0$, and the agent's optimal action converges to zero as $\chi \rightarrow \infty$.

¹⁹Section 7 allows the principal to use public randomization, in addition to the reports, to determine consumption.

The principal observes only the reports $\{Y_t\}$, but the initial cash-flow is publicly known, so that $Y_0 = X_0$.

The agent's strategy is a lying process L and an effort process A adapted to the agent's information X .²⁰

Given a contract, the agent maximizes his expected utility. The resulting value function is

$$V_0 = \sup_{L,A} \left\{ E \left[\int_0^\infty e^{-rt} (u(C_t + X_t - Y_t) - \phi(A_t)) dt \right] \right\},$$

where u is a strictly concave utility function. Computations to follow focus on the case where u is exponential: $u(c) = -\exp(\theta c)$ for some risk-aversion coefficient θ . The set of promised utility is then $\mathcal{W} = (-\infty, 0)$.

A contract (C, \bar{A}) is *incentive compatible given (w, y)* if it is optimal for the agent to report and transfer truthfully the real cash-flow process X and to implement the stipulated action \bar{A} , given that the initial cash flow is y , and if the resulting expected lifetime utility for the agent is w . For simplicity, the effort process will sometimes be dropped from the definition: *a contract C is incentive compatible given (w, y)* if it is optimal for the agent to report and transfer truthfully the real-cash flow process X , given that the initial cash flow is equal to y , and that this, along with the action process optimally chosen by the agent, yields an expected lifetime utility of w to the agent.

The agent immediately consumes the sum of the consumption C_t provided by the principal and of the difference $X_t - Y_t$ between real and transferred cash flows. (the possibility of private savings is discussed in Section 9).

The principal's expected payoff is

$$E \left[\int_0^\infty e^{-rt} (Y_t - C_t) dt \right].$$

The contract must initially provide the agent with some minimal expected promised utility w :

$$V_0 \geq w.$$

As is often assumed in the literature on dynamic contracting, the agent loses his outside option after time zero, and is fully committed to the contract.²¹ Similarly, the principal is

²⁰This is without loss of generality, since X determines Y through L .

²¹The case where the agent keeps his outside option throughout the lifetime of the contract is discussed in detail in Section 8.

at all times fully committed to providing the agent his promised utility, although he may propose at any time new contracts that preserve the agent's promised utility.

Persistence of private information creates a complex strategic environment. For example, if the agent has lied even for a short period before time t , he has affected the report history $Y^t = \{Y_s\}_{s \leq t}$. He has therefore affected his future consumption flow C and his future incentives to report the truth. Hidden actions add to this complexity: current effort affects immediate cash-flows but also future ones, owing to their persistence, and hence the entire consumption process and the distribution of all future states.

From the Martingale Representation Theorem,²² the promised utility of the agent satisfies

$$dW_t = (rW_t - u(C_t + Y_t - X_t) + \phi(\bar{A}_t))dt + \sigma S_t d\tilde{B}_t, \quad (3)$$

for some process S_t adapted to the filtration of the principal, and where \tilde{B}_t is a Brownian motion under the probability measure where the agent reports truthfully and chooses the prescribed effort \bar{A}_t , i.e.,

$$d\tilde{B}_t = \frac{dY_t - (\xi - \lambda Y_t)dt - \bar{A}_t dt}{\sigma} = \frac{(\xi - \lambda X_t) - (\xi - \lambda Y_t) + A_t - \bar{A}_t + L_t dt + \sigma dB_t}{\sigma}.$$

The sensitivity S_t describes how promised utility varies with reports from the agent, and is chosen by the principal in the recursive formulation of the problem.²³

The agent's problem is to solve

$$V_0 = \sup_{L,A} E \left[\int_0^\infty e^{-rt} (u(C_t + Y_t - X_t) - \phi(A_t)) dt \right],$$

subject to (1), (2), and (3).

If the agent misreports his cash-flow increment dX_t at time t , he affects two things: the change of his promised utility, which is sensitive to the report dY_t by a factor S_t , and the change of the principal's future consumption C . Indeed, for a given promised utility, the principal must provide higher payments to the agent, other things equal, if he thinks that the cash-flow is lower. The former channel gives an incentive for the agent to make high

²²See, e.g., Karatzas and Shreve (1991)

²³This representation may be compared to the discrete time, where the principal would have to specify changes in promised utility for each possible report of the agent. Continuous time linearizes the problem, and the derivative S_t , a single number, completely describes how reported cash-flow increments affect the agent's promised utility.

reports, which the second channel gives him an incentive to report a lower cash-flow. For the contract to be incentive compatible, these two incentives must balance each other.

3 Renegotiation-Proofness: Concepts and State Consistency

In repeated games, an equilibrium is said to be weakly renegotiation-proof (Farrell and Maskin, 1989), or internally consistent (Bernheim and Ray, 1989), if there are no two histories such that the continuation payoffs after the first history Pareto dominate those following the second history.²⁴

With stochastic games, internal consistency is too weak, because it does not allow comparison across states. In the present setting, the state consists of the current cash-flow and of the utility promised to the agent. A contract is internally consistent if there are no two histories leading to the *same* underlying state (cash flow and promised utility) such that the principal gets a strictly higher continuation payoff after the first history than after the second one.^{25,26}

To see how internal consistency may be strengthened for stochastic games, it is useful to think about its rationale. Internal consistency presumes that, after observing the second history, *the principal is able to recognize* that he could use the continuation contract following the first history and achieve a higher payoff.²⁷ This *cognitive ability should extend to other natural comparisons*, as described next.

To illustrate, suppose that some incentive compatible contract $\{C_t\}_{t \geq 0}$ results in some promised utility w_1 for the agent and some payoff π_1 to the principal. Now suppose that the principal must provide utility $w_2 = w_1 + \Delta w$ to the agent and that the agent has linear utility (just for the purpose of this illustration). Then, the translated contract with consumption

²⁴Although apparently benign and well accepted in the literature, these concepts have been criticized by Asheim (1991), Abreu and Pearce (1991) and Abreu et al. (1993).

²⁵It is assumed that if the agent is indifferent, he agrees to switch to this new contract. The principal can always give an infinitesimal share of the gain to the agent to convince him to switch.

²⁶This extension of internal consistency, which was also made by Bergemann and Hege (2005), immediately implies that the principal's payoff, after any history, only depends on the current state (w, y) . Showing that the contractual variables themselves are Markov is more difficult, need not be true. See Section 4.1.

²⁷Of course, such change could affect the agent's incentives to report truthfully ex ante. Eventually, the goal is to characterize contracts that are both renegotiation-proof and incentive compatible.

process $\tilde{C}_t = C_t + r\Delta w$ provides utility w_2 to the agent and does not change the agent's incentives.²⁸ Let π_2 denote the resulting payoff for the principal. The principal can make the following comparison: suppose that, starting from state w_1 , the contract reaches some time at which the promised utility of the agent is $W_t = w_2$, and the continuation for the principal is Π_t . If $\Pi_t < \pi_2$, the principal could reason that, by restarting the contract exactly as it did from time zero, but uniformly halving the consumption process, he could achieve the higher payoff π_2 . This comparison creates a comparison across states.

This suggests the following definition. Suppose that, starting from any contract C that is incentive compatible given (w_1, y) , there is an operation \mathcal{G} that “transforms” this contract into another contract $\tilde{C} = \mathcal{G}_{w_2, w_1}(C)$ that is incentive compatible given (w_2, y) (in the previous example, the transformation was to translate the consumption process). The contract C is *consistent from w_1 to w_2* if the payoff $\tilde{\pi}_2$ achieved by \tilde{C} is not strictly greater than the continuation payoff achieved under C after any history such that $W_t = w_2$ and $Y_t = y$. If that condition did not hold, the principal could obtain a higher payoff by proposing contract \tilde{C} , after the relevant history, instead of the continuation contract initially specified.

The operation should be reversible in the following sense: if, starting from the transformed contract $\tilde{C} = \mathcal{G}_{w_2, w_1}(C)$, one applies a similar operation $\mathcal{G}_{w_1, w_2}(\tilde{C})$ to get a contract that is incentive compatible given (w_1, y) , the resulting contract is the initial contract C . Although not necessary for the results, the transformations considered in the paper will be consistent in the following sense: $(\mathcal{G}_{w_3, w_2} \circ \mathcal{G}_{w_2, w_1})(C) = \mathcal{G}_{w_3, w_1}(C)$. We will say that \mathcal{G} is a *transformation group* operating on the set of contracts.²⁹ Choosing $w_3 = w_1$, the previous conditions imply that the operation \mathcal{G}_{w_1, w_1} is the identity for each w_1 : it leaves any contract that is IC given (w_1, y) unchanged.

The transformation group constructed in Section 4.3.1 has a natural motivation, given the context of this paper. First, we will show that drift linearity of the cash-flow dynamic equation implies that if a contract is IC given (w, y) , the same contract, seen as a functional of the exogenous shocks (Brownian motion) inferred from the reports, is also IC given (w, y') for any y' . A transformation group built on this property implies that for any renegotiation-proof contract, consumption is independent of current income. Second, exponential utility

²⁸This natural conjecture is taken as granted for now, and checked in Section 4.3.2.

²⁹ \mathcal{G} not a group strictly speaking, because we need to consider the composition $\mathcal{G}_{w_3, w_2} \circ \mathcal{G}_{w_1, w_2}$ only if $w_2 = w_1$, even though such composition may be well-defined even when $w_2 \neq w_1$, as for the group constructed in later sections.

implies that if a contract is IC given (w, y) , translating the consumption process by an appropriately chosen constant results in a contract that is IC given (w', y) for any given w' . Incorporating such transformation in the group will imply a functional relation between consumption and promised utility, for any renegotiation-proof contract. Taken together, these properties define a transformation group \mathcal{G} whose elements $\mathcal{G}_{(w_2, y_2), (w_1, y_1)}$ are indexed by pairs of initial state (w_1, y_1) and final state (w_2, y_2) .

As we shall see, the transformation group defined in the next section can be seen as an axiomatization based on renegotiation-proofness of contracts whose consumption are linear in some function of the promised utility and are independent of current income.

For now, we return to the definition of renegotiation-proofness based on general transformation groups.

DEFINITION 1 *A contract C that is incentive compatible given (w, y) is state-consistent (relative to \mathcal{G}) or renegotiation-proof if, after any history leading up to any state (\tilde{w}, \tilde{y}) and continuation contract \tilde{C} , the continuation payoff $\Pi(\tilde{C})$ for the principal is weakly greater than $\Pi(\mathcal{G}_{(\tilde{w}, \tilde{y}), (w, y)}(C))$, and reciprocally, the initial payoff $\Pi(C)$ is weakly greater than $\Pi(\mathcal{G}_{(w, y), (\tilde{w}, \tilde{y})}(\tilde{C}))$.*

Thus, not only the continuation payoffs must sustain comparison with transformations of the initial contract, but the reverse is also true: the initial contract must sustain the comparison with transformations of the continuation contracts consistent with the initial state.

Remark: by definition, a renegotiation-proof contract is incentive compatible for some state (w, y) .

The transformation group \mathcal{G} is *monotone* if, for any two contracts C, C' that are incentive compatible given (w, y) , and yield principal payoffs $\Pi(C) \leq (<) \Pi(C')$, and any other state (\tilde{w}, \tilde{y}) , the payoffs of the transformed contracts $\tilde{C} = \mathcal{G}_{(\tilde{w}, \tilde{y}), (w, y)}(C)$ and $\tilde{C}' = \mathcal{G}_{(\tilde{w}, \tilde{y}), (w, y)}(C')$ satisfy $\Pi(\tilde{C}) \leq (<) \Pi(\tilde{C}')$.

PROPOSITION 1 *Suppose that \mathcal{G} is monotone and let C denote any contract that is incentive compatible for some state (w, y) and state-consistent with respect to \mathcal{G} . Then, after any finite history ending with state (\tilde{w}, \tilde{y}) , the continuation payoff for the principal is equal to his initial payoff under the contract $\mathcal{G}_{(\tilde{w}, \tilde{y}), (w, y)}(C)$.*

In principle, there may be several transformations groups to consider. The next result shows that, for a contract to be renegotiation-proof with respect to several monotone transforma-

tion groups, these groups must, taken individually, yield the same concept of renegotiation-proofness. Consider two monotone transformation groups \mathcal{G} and $\tilde{\mathcal{G}}$. Given a contract C that is IC given (w, y) , suppose that there is a state (w', y') for which the groups yield different payoffs: $\Pi(C') = \Pi(\mathcal{G}_{(w', y'), (w, y)}(C)) > \Pi(\tilde{\mathcal{G}}_{(w', y'), (w, y)}(C)) = \Pi(\tilde{C}')$. Applying the transformations again from (w', y') to (w, y) and using monotonicity, we get

$$\Pi(\tilde{\mathcal{G}}_{(w, y), (w', y')} \circ \mathcal{G}_{(w', y'), (w, y)}(C)) > \Pi(C).$$

Therefore, that C is dominated by a simple composition of elements in the two groups. In such case, we will say that C is *unstable* with respect to $(\mathcal{G}, \tilde{\mathcal{G}})$, otherwise, C is called *stable*. Intuitively, instability implies that the ability to consider transformations that yield different payoffs after a given history prevents the existence of a renegotiation-proof contract.

PROPOSITION 2 (CONCEPT EQUIVALENCE FOR STABLE CONTRACTS) *Let \mathcal{G} and $\tilde{\mathcal{G}}$ be two monotone transformation groups, and C be a contract that is incentive compatible given some state (w, y) and stable with respect to \mathcal{G} and $\tilde{\mathcal{G}}$. Then, C is renegotiation-proof with respect to \mathcal{G} if and only if it is renegotiation-proof with respect to $\tilde{\mathcal{G}}$. Moreover, the principal's payoffs for the \mathcal{G} and $\tilde{\mathcal{G}}$ transformations of C to any state (w', y') are identical.*

4 Characterization of State-Consistent Contracts

4.1 Contractual Variables are Markov

As observed in the previous section, the principal's continuation payoff for any internally consistent contract only depends, at any time, on the current state (w, y) . Let $\Pi(w, y)$ denote the principal's payoff.

More can be obtained: the *contractual variables* themselves only depend on (w, y) , under some additional conditions. The principal can choose two variables at each time: the consumption rate c provided to the agent, and the sensitivity s of the promised utility to reported cash-flow. For a given internally consistent contract C and states (w, y) , let $\mathcal{K}(w, y) = \{(C_t(\omega), S_t(\omega)) : (t, \omega) \in \mathbb{R}_+ \times \Omega, W_t(\omega) = w, Y_t(\omega) = y\}$ denote the set of consumption and sensitivity levels that may arise, under contract C , after some history leading to state (w, y) . The goal is to show that these sets are in fact singletons, so that the consumption and sensitivity chosen by the principal indeed only depend on w and on y . Allowing

the principal to randomize across two continuation contracts convexifies the set $\mathcal{K}(w, y)$.^{30,31} Weak renegotiation proofness implies that the principal is free to choose his current actions optimally within $\mathcal{K}(w, y)$, which is captured by the Hamilton-Jacobi-Bellman equation³²

$$0 = \sup_{(c,s) \in \mathcal{K}(w,y)} \left\{ y - c - r\Pi(w, y) + \Pi_w(w, y) (rw - u(c) + \phi(a(s))) + \Pi_y((\xi - \lambda y) + a(s)) + \frac{1}{2}\Pi_{ww}s^2 + \Pi_{wy}s\sigma + \frac{1}{2}\Pi_{yy}\sigma^2 \right\}, \quad (4)$$

where $a(s)$ is the effort level optimally chosen by the agent, and is given by the first-order condition $\phi'(a(s)) = s$.³³ The objective is strictly concave in c and in s , provided that i) Π_w and Π_{ww} are negative (i.e., the principal's payoff is decreasing in the agent's promised utility, other things equal, and strictly concave in the agent's promised utility), Π_y is positive (i.e., the principal's payoff is increasing in the current cash-flow, keeping promised utility constant), and ii) u is strictly concave, $\phi(a(s))$ is weakly concave in s , and $a(s)$ is weakly convex in s . The latter set of conditions is satisfied if ϕ is exponential and Section 4.3.3 independently shows that the first set of conditions is always satisfied for state-consistent contracts with exponential utility and cost functions.

Strict concavity of the objective function and convexity of the domain $S(w, y)$ imply that the maximizing pair $c(w, y), s(w, y)$ is unique, which shows that the contractual variables are Markov (and, therefore, that the set $\mathcal{K}(w, y)$ had to be a singleton).

As a result, the agent faces a standard optimal control problem where the variables are the public, contractual state variables w and y , and the actual cash-flow x that is privately observed by the agent.

³⁰The agent chooses his report and action, at each instant t , before observing the outcome of the randomization.

³¹The randomization adds a new source of potential uncertainty, only for the present argument: The randomization does not arise on the equilibrium path.

³²Bergemann and Hege (2005, Theorem 3) also exploit the Bellman equation to show that any weakly renegotiation-proof contract must be Markov. That paper does not address the possibility of multiple maximizers of the Bellman equation, for any given state.

³³The argument for pinning down the effort level given the sensitivity parameter s is similar to Sannikov (2008, Proposition 2), which does not assume a Markovian structure. In the present setting, effort also has an indirect impact of reported cash flows and rewards, but this impact is identical to the impact of a lie, and must vanish for any incentive compatible (i.e., truthful) contract, as illustrated by Equations (7) and (8) for the Markovian case.

4.2 Necessary Conditions for Incentive Compatibility

The agent's Hamilton-Jacobi-Bellman (HJB) equation is

$$0 = \sup_{l,a} \{ [u(c(w,y) + x - y) - \phi(a)] - rv(w,y,x) \\ + v_w [rw - u(c(w,y)) + \phi(\bar{a}(w,y)) + s(w,y)(l + (a - \bar{a}(w,y)) + \lambda(y - x))] \\ + v_y [(\xi - \lambda x) + l + a] + v_x [(\xi - \lambda x) + a] + q(w,y,x) \}, \quad (5)$$

where $q(w,y,x) = \sigma^2 \left(\frac{1}{2}(v_{ww}s^2 + v_{xx} + v_{yy}) + v_{wx}s + v_{w,y}s + v_{xy} \right)$ is independent from l and from a .

Incentive compatibility and optimality imply several regularity properties of v .³⁴ First, if the contract is incentive compatible, $v(w,x,x) = w$ for all w,x , which implies that

$$v_w(w,x,x) = 1 \quad (6)$$

for all w and x . Second, as will be seen shortly, the Envelope Theorem implies that $v_x(w,y,x)$ exists and is continuous for all (w,y,x) . Taken together, these observations imply that $v_y(w,y,x)$ exists whenever $y = x$, and that

$$v_y(w,x,x) = -v_x(w,x,x)$$

for all w,x .

PROPOSITION 3 *If the contract $(w,x) \mapsto p(w,x)$ is incentive compatible, then there exists a viscosity solution to (5) such that*

$$v(w,x,x) = w,$$

such that

$$v_x(w,x,x) = -v_y(w,x,x),$$

and where

$$s(w,x) + v_y(w,x,x) = 0 \quad (7)$$

for all w and x . Moreover, the agent's optimal effort $a(w,y,x)$ satisfies

$$\phi'(a(w,x,x)) = v_x(w,x,x) = s(w,x). \quad (8)$$

³⁴In general, the value function v is only a viscosity solution to the equation (see Fleming and Soner, 2006, Chapter 5).

Reciprocally, if there exists a solution \bar{v} to (5) such that $\bar{v}(w, x, x) = w$ for all w, x , then \bar{v} is the value function of the agent, the contract is incentive compatible, and (8) holds.

Incentive compatibility is characterized by finding a *global* solution v to (5), i.e., a solution that holds for all possible values of the state (w, y, x) , and then check that (7) holds (or, equivalently, that $v(w, x, x) = w$) for that solution.

The relevant derivatives of v can now be computed. Recall that

$$v(w, y, x) = \sup_{L, A} E \left[\int_0^\infty e^{-rt} (u(\mathcal{C}_t(Y_s : s \leq t) + X_t - Y_t) - \phi(A_t)) dt \right],$$

where $\mathcal{C}_t(\cdot)$ is, for each t , a functional that determines the consumption provided to the agent at time t given past reports $Y_s : s \leq t$. If the initial cash-flow is increased by ε , this affects the distribution of future incomes and, keeping the lying process fixed, of future reports. However, by a change of variable, one can control the path of the report process Y_t , and make it independent from the initial cash-flow change. Recall that

$$dY_t = [(\xi - \lambda X_t) + A_t + L_t]dt + \sigma dB_t.$$

Making the change of variable $\bar{L}_t = L_t + (\xi - \lambda X_t) - (\xi - \lambda Y_t)$, one gets

$$dY_t = [(\xi - \lambda Y_t) + A_t + \bar{L}_t]dt + \sigma dB_t. \quad (9)$$

The agent's strategy can be restated as choosing \bar{L} , rather than L :

$$v(w, y, x) = \sup_{\bar{L}, A} E \left[\int_0^\infty e^{-rt} (u(\mathcal{C}_t(Y_s : s \leq t) + X_t - Y_t) - \phi(A_t)) dt \right].$$

subject to $Y_0 = y$, $X_0 = x$, (9), and

$$dX_t = (\xi - \lambda X_t)dt + A_t + \sigma dB_t.$$

$$dW_t = (rW_t - u(\mathcal{C}_t(Y_s : s \leq t)) + \phi(\bar{A}_t))dt + S_t (dY_t - ((\xi - \lambda Y_t) + \bar{A}_t)dt).$$

If the contract is incentive compatible, it is optimal to set $\bar{L}_t = 0$ whenever initial conditions are such that $y = x$. By the Envelope Theorem (Milgrom and Segal, 2002), this implies that $v_x(w, x, x)$ can be computed by evaluating the objective function at $\bar{L}_t \equiv 0$ or, equivalently, under the report process Y_t starting from $y_0 = x_0$. Under this approach, W , Y , \mathcal{C} , and \bar{A} are independent from the initial condition x for the actual report. One has

$$v_x(w, x, x) = \int_0^\infty e^{-rt} \frac{d}{dx} E[u(X_t - Y_t + \mathcal{C}(Y_s : s \leq t))]dt.$$

Since the distribution of $\{Y_s\}_{s \leq t}$ is independent from the initial condition x , the inner derivative simply equals

$$E \left[u'(X_t - Y_t + C_t(Y_s : s \leq t)) \frac{dX_t}{dx} \right].$$

The process X , as defined by the dynamic equation (1), is a generalization of Ornstein-Uhlenbeck processes, and can be explicitly integrated:

$$X_t = e^{-\lambda t} x + \int_0^t e^{\lambda(s-t)} (A_s + \xi) ds + \int_0^t e^{\lambda(s-t)} \sigma dB_s \quad (10)$$

This implies that $\frac{dX_t}{dx} = e^{-\lambda t}$ and, therefore, that

$$v_y(w, x, x) = - \int_0^\infty e^{-(r+\lambda)t} u'(C_t + X_t - Y_t) dt. \quad (11)$$

4.3 State Consistency

4.3.1 Comparing Contracts Across Cash-Flows Levels

Starting with initial conditions w, y, x , the agent's value for a given strategy (L, A) is

$$V(w, y, x, L, A) = E \left[\int_0^\infty e^{-rt} (u(c(W_t, Y_t) - G_t) - \phi(A_t)) dt \right] \quad (12)$$

subject to (1), (2),

$$\begin{aligned} dW_t &= (rW_t - u(C_t) + \phi(\bar{A}_t))dt + s(W_t, X_t + G_t)(L_t dt + A_t dt + \lambda G_t + dB_t), \\ dG_t &= L_t dt \end{aligned}$$

and the initial conditions $W_0 = w$ and $Y_0 = y, X_0 = x, G_0 = y - x$.

The principal's expected payoff is

$$\pi(w, y) = E_{w,y} \int_0^\infty e^{-rt} (Y_t - C_t) dt.$$

Suppose that the contract C is incentive compatible given (w, y) , with $C_t = \mathcal{C}(Y_s : s \leq t)$ for some functional \mathcal{C} .

Starting from a different cash-flow level \hat{y} , and given a report process \hat{Y}_t , suppose that the principal pays the consumption process $\hat{C}_t = \mathcal{C}(\hat{Y}_s : s \leq t)$, where \hat{Y} is constructed as follows: $\hat{Y}_0 = y$, and

$$d\tilde{Y}_t = d\hat{Y}_t - (\xi - \lambda \hat{Y}_t) dt + (\xi - \lambda \tilde{Y}_t) dt.$$

The intuition for this construction is as follows. First, the principal reconstructs the reports \tilde{Y} that the agent would have made, had he started from y instead of \hat{y} , under the same realization of the Brownian path that generated report history \hat{Y} , assuming that the agent is truthful and follows the action process A . Second, the principal provides the consumption that he would have provided under the contract C , had the agent started from y instead of \hat{y} and made the report \tilde{Y} .

In the Appendix, it is shown that this construction yields an incentive compatible contract given (w, \hat{y}) . If L, A was an optimal strategy for the agent, starting from y , it must also be optimal given the new contract.

PROPOSITION 4 *Suppose that the contract process $C_t = \mathcal{C}(Y_s : s \leq t)$ along with the prescribed effort process A_t is incentive compatible given (w, y) . Then, the reconstructed process $\hat{C}_t = \mathcal{C}(\tilde{Y}_s : s \leq t)$ along with the same prescribed effort process A is incentive compatible given (w, \hat{y}) .*

The contract \hat{C} is called the (w, \hat{y}) -version of C .

The processes C and \hat{C} have the same distribution. Therefore, the principal has the same expected consumption cost under these two contracts. The only difference for the principal, then, is the expectation Υ of the discounted cash-flow stream transferred to him by the agent. Thus,

$$\Upsilon(w, y) = E \int_0^\infty e^{-rt} Y_t dt.$$

If the contract is truthful, $Y_t = X_t$ as given by (10). Therefore,

$$\Upsilon(w, y) = \int_0^\infty e^{-rt} \left(e^{-\lambda t} y + E \left[\int_0^t e^{\lambda(s-t)} (A_s + \xi) ds \right] \right) dt.$$

After simplification,

$$\Upsilon(w, y) = \frac{y}{r + \lambda} + \frac{\xi}{r(r + \lambda)} + \int_0^\infty e^{-rt} \frac{\alpha_t}{r + \lambda} dt, \quad (13)$$

where

$$\alpha_t = E[A_t].$$

4.3.2 Comparing Contracts Across Promised-Utility Levels

From now on, the agent is assumed to have the utility function $u(c) = -\exp(-\theta c)$ and the cost function $\phi(a) = \bar{\phi} \exp(\chi a)$. In particular, utility is always negative.

Suppose that, starting from initial conditions (w_0, y) , the contract C is incentive compatible and induces the effort process A . That is, letting $v(L, A|C)$ denote the agent's expected utility when he follows strategy L, A and given contract C ,

$$v(0, A|C) = w_0 \geq v(L', A'|C)$$

for all (L', A') . Now consider another promised utility level $w_1 = \beta w_0$ for $\beta \in (0, \infty)$ and the initial state (w_1, x) . Define a new contract (\hat{C}, \hat{A}) as follows

$$\begin{aligned}\hat{C}_t &= C_t - \frac{\log(\beta)}{\theta} \\ \hat{A}_t &= A_t + \frac{\log(\beta)}{\chi}\end{aligned}$$

PROPOSITION 5 *(C, A) is incentive compatible and provides expected utility w_0 if and only if (\hat{C}, \hat{A}) is incentive compatible and provides expected utility w_1 .*

Proof. For any \hat{L}', \hat{A}' , let $L' = \hat{L}'$ and $A' = \hat{A}' - \log(\beta)/\chi$. Then,

$$v(\hat{L}', \hat{A}'|\hat{C}) = \beta v(L', A'|C) \leq \beta v(0, A|C) = \beta w_0 = w_1,$$

with the inequality being tight if $\hat{L}' = 0$ and $\hat{A}' = \hat{A}$. ■

The contract \hat{C} is called the (w_1, y) -version of C .

Let $\Pi(C, A)$ denote the expected payoff for the principal when the agent receives the consumption process C and follows effort process A , and let $\Pi(w, y)$ denote the value function of the principal starting from state (w, y) . The previous analysis shows that³⁵

$$\Pi(w_1, y) = \frac{\log(\beta)}{r} \left(\frac{1}{\theta} + \frac{1}{\chi(r + \lambda)} \right) + \Pi(w_0, y).$$

The previous analysis yields the following result.

PROPOSITION 6 *To any contract C that is incentive compatible given (w, y) corresponds another contract C' that is incentive compatible give (w', y') , called the (w', y') -version of C .*

The principal's payoffs across versions satisfy the following relation:

$$\Pi(w', y') = \frac{(y' - y)}{r + \lambda} + \frac{\log(w'/w)}{r} \left(\frac{1}{\theta} + \frac{1}{\chi(r + \lambda)} \right) + \Pi(w, y). \quad (14)$$

³⁵The effort taken at any time has a decaying effect on all future incomes, with decaying rate λ , discounted at rate r . This explains the factor $(r + \lambda)$ in the denominator.

4.3.3 Form of State-Consistent Contracts

Combined together, Sections 4.3.1 and 4.3.2 define a class of transformations $\mathcal{G}_{(\tilde{w}, \tilde{y}), (w, y)}$ which, to any contract C that is incentive compatible given (w, y) , associates a contract $\tilde{C} = \mathcal{G}_{(\tilde{w}, \tilde{y}), (w, y)}(C)$ that is incentive compatible given (\tilde{w}, \tilde{y}) .

The transformations are clearly monotone, so that state consistency implies the following relationships for the principal's payoff function, as a consequence of Proposition 1.³⁶

PROPOSITION 7 *For any state-consistent contract, and any time t ,*

$$\Pi(W_t, Y_t) = \frac{Y_t}{r + \lambda} + \frac{\log(-W_t)}{r} \left(\frac{1}{\theta} + \frac{1}{\chi(r + \lambda)} \right) + \Pi(-1, 0).$$

To pin down the contractual variables, one proceeds as in Section 4.1. State consistency extends the set of feasible challengers that the principal may consider after any possible history, which may be exploited to show that the contractual variables must take a very specific form.

Consider some internally-consistent contract C that is incentive compatible given some arbitrary conditions (w_0, y_0) . For any outcome ω and time t , let $C^t(\omega)$ denote the continuation contract of C at time t as outcome ω unfolds. For any (\tilde{w}, \tilde{y}) , let

$$\Gamma^C(\tilde{w}, \tilde{y}) = \{C^t(\omega) : (t, \omega) \in [0, \infty) \times \Omega \text{ s.t. } (W_t, Y_t) = (\tilde{w}, \tilde{y})\}$$

denote the set of all continuation contracts generated by C after any history leading up to state (\tilde{w}, \tilde{y}) . Let also

$$\mathcal{C}(w, y) = \cup_{(\tilde{w}, \tilde{y}) \in \mathcal{W} \times \mathbb{R}} \left\{ \mathcal{G}_{(w, y), (\tilde{w}, \tilde{y})}(\tilde{C}) : \tilde{C} \in \Gamma^C(\tilde{w}, \tilde{y}) \right\}$$

denote the set of all contracts that are (w, y) -versions of continuation contracts of C . The set $\mathcal{C}(w, y)$ defines the class of all challengers that the principal can consider, after any history leading up to state (w, y) , to replace the current continuation contract. That set is larger than the one corresponding to internal consistency, where only the continuation contracts starting from the same state can be compared.

Finally, let $\mathcal{I}(w, y) = \{(\tilde{C}_0, \tilde{S}_0) : \tilde{C} \in \mathcal{C}(w, y)\}$. $\mathcal{I}(w, y)$ is the set of initial consumption-sensitivity pairs for all contracts that are incentive compatible given (w, y) and generated from some continuation contract of C .

³⁶As mentioned in Section 4.1, Proposition 7 has been established without assuming Markov contractual variables, and is thus independent from the analysis of Section 4.1.

After any such history the principal can choose, among the pairs in $\mathcal{I}(w, y)$, one that maximizes his payoff, as captured by the principal's HJB equation.

$$0 = \sup_{(c,s) \in \mathcal{I}(w,y)} \left\{ y - c - r\Pi(w, y) + \Pi_w(w, y) (rw - u(c) + \phi(a(s))) + \Pi_y((\xi - \lambda y) + a(s)) + \frac{1}{2}\Pi_{ww}s^2 + \Pi_{wy}s\sigma + \frac{1}{2}\Pi_{yy}\sigma^2 \right\}, \quad (15)$$

where $\phi'(a(s)) = s$. Let (\bar{c}, \bar{y}) denote the optimum for (w, y) , which is unique, since the functional form of Proposition 7 and the fact that $\phi(a) = \bar{\phi} \exp(\chi a)$ imply strict concavity of the objective. It is easy to see that for any (\tilde{w}, \tilde{y}) , the pair $(\tilde{c}, \tilde{s}) = (\bar{c} - \log(-\beta)/\theta, \bar{s}\beta)$, where $\beta = \tilde{w}/w$, belongs to $\mathcal{I}(\tilde{w}, \tilde{w})$, and vice versa. Given the functional form of Π , this implies that (\bar{c}, \bar{s}) solves (15) at (w, y) if and only if (\tilde{c}, \tilde{s}) solves it at (\tilde{w}, \tilde{y}) , as is easily checked.

This establishes the following result.

PROPOSITION 8 *A state-consistent contract has the form*

$$\begin{aligned} C_t &= c_1 - \frac{\log(-W_t)}{\theta} \\ A_t &= a_1 + \frac{\log(-W_t)}{\chi} \\ S_t &= -W_t \bar{s} \end{aligned}$$

for all t , where c_1 , a_1 and \bar{s} are the consumption, effort, and sensitivity provided at time 0 by the version of the contract starting with promised utility -1 and any cash-flow level.

Therefore, the principal's optimal contracting problem reduces to optimization with respect to variables c_1 , a_1 and \bar{s} , subject to incentive compatibility constraints. More conveniently, consider the variables

$$u_1 = u(c_1)$$

and

$$\phi_1 = \phi(a_1).$$

Equivalently, $c_1 = -\log(-u_1)/\theta$ and $a_1 = \log(\phi_1/\bar{\phi})/\chi$.

The principal's problem is to maximize $\Pi(w, x)$ with respect to u_1 and ϕ_1 . From (13), this is equivalent to

$$\max_{u_1, \phi_1} E \int_0^\infty e^{-rt} \left(\frac{\alpha_t}{r + \lambda} - E[C_t] \right),$$

where

$$\alpha_t = a_1 + \frac{E[\log(-W_t)]}{\chi}$$

and

$$E[C_t] = c_1 - \frac{E[\log(-W_t)]}{\theta}.$$

Therefore, one needs to compute $E[\log(-W_t)]$. Recall that

$$dW_t = [rW_t - u(C_t) + \phi(A_t)]dt + S_t\sigma dB_t.$$

With the exponential specification, $u(C_t) = -W_t u_1$ and $\phi(A_t) = -W_t \phi_1$. This implies that, letting $\omega(t) = E[W_t]$ (so that $\omega(0) = w$)

$$\omega'(t) = (r + u_1 - \phi_1)\omega(t),$$

and hence that

$$E[W_t] = e^{(r+u_1-\phi_1)t}w. \tag{16}$$

Moreover, (11) and (7), combined with $u'(C_t) = -\theta u(C_t)$, imply that

$$s(w) = E \int_0^\infty e^{-(r+\lambda)t} \theta u(C_t) dt = \theta u_1 \int_0^\infty e^{-(r+\lambda)t} E[W_t] dt.$$

Combining this with (16) yields

$$s(w) = -w\bar{s},$$

where

$$\bar{s} = \frac{\theta(-u_1)}{\lambda - u_1 + \phi_1} > 0. \tag{17}$$

Notice that $\bar{s} < \theta$ for all $u_1 < 0$, $\lambda > 0$ and $\phi_1 \geq 0$. Intuitively, if \bar{s} were higher than θ , the agent would want to exaggerate the cash-flow in order to artificially increase his promised utility. The cost of earning actually less than what is reported to the principal affect the utility by a rate θ , which would be dominated by the increase in promised utility as measured by the sensitivity parameter \bar{s} . Incentive compatible rules this case out.

Incentive compatibility imposes an additional relation between ϕ_1 and \bar{s} . Precisely, if the contract is incentive compatible,

$$\phi_1(\bar{s}) = \frac{\bar{s}}{\chi},$$

as implied by equations (7) and (8).

This, along with (17), implies that

$$u_1(\bar{s}) = -\frac{\bar{s}(\chi\lambda + \bar{s})}{\chi(\theta - \bar{s})}. \quad (18)$$

Letting $Z_t = \log(-W_t)$, Itô's formula implies that

$$dZ_t = (r + u_1 - \phi_1)dt - \frac{1}{2}\sigma^2\bar{s}^2 dt - \bar{s}dB_t.$$

Therefore,

$$E \log(-W_t) = \log(-w) + (r + u_1 - \phi_1)t - \frac{1}{2}\sigma^2\bar{s}^2 t.$$

The principal's objective is to maximize

$$\frac{a_1}{r(r + \lambda)} - \frac{c_1}{r} + \left(\frac{1}{\chi(r + \lambda)} + \frac{1}{\theta} \right) \int_0^\infty e^{-rt} E \log(-W_t) dt.$$

After further simplifications and multiplication of the objective by θr , the following result obtains. Let $\kappa = \frac{\theta}{\chi(r + \lambda)} > 0$.

PROPOSITION 9 *The optimal contract is determined by choosing \bar{s} so as to maximize the principal's expected payoff³⁷*

$$\kappa \log(\phi_1(\bar{s})) + \log(-u_1(\bar{s})) + \frac{\kappa + 1}{r} \left(r + u_1(\bar{s}) - \phi_1(\bar{s}) - \frac{1}{2}\sigma^2\bar{s}^2 \right) + (\kappa + 1) \log(-w) - \kappa \log(\bar{\phi}). \quad (19)$$

5 Comparative Statics

5.1 Immiserization

From Proposition 9, one may easily show that the drift of promised utility is negative, for the case of pure reporting. In that case, Equation (19) reduces to

$$\max_{\bar{s}} \log(-u_1(\bar{s})) + \frac{1}{r} \left(r + u_1(\bar{s}) - \frac{1}{2}\sigma^2\bar{s}^2 \right) + (\kappa + 1) \log(-w), \quad (20)$$

³⁷The last two terms of the objective are independent from \bar{s} and, therefore, have no impact on the maximization.

or equivalently,

$$\max_{u_1} \log(-u_1) + \frac{1}{r} \left(r + u_1 - \frac{1}{2} \sigma^2 \bar{s}(u_1)^2 \right) + (\kappa + 1) \log(-w), \quad (21)$$

where

$$\bar{s}(u_1) = \frac{\theta(-u_1)}{\lambda - u_1}.$$

If the volatility σ were equal to zero (no private information), or if λ were infinite (no persistence), the final, quadratic term in (21) would vanish, and the optimal u_1 would equal $-r$, implying that the drift of W_t , $r + u_1$ is equal to zero. This would amount to pure consumption smoothing: given concavity of the agent's utility function, the cheapest way to give him a promised utility of w is through a constant consumption flow of rw , which keeps W_t constant (or, where σ is nonzero, implies that W_t is a martingale). In general however, the principal also needs to mitigate the agent's incentive to misreport the cash-flow. That incentive is captured by the term

$$v_y = E \left[\int_0^\infty e^{-(r+\lambda)t} u'(C_t) dt \right]. \quad (22)$$

Therefore, this incentive to misreport is lower, other things equal, if u' is lower. Since u is concave, this means that providing more consumption has the additional benefit, other things equal, of reducing marginal utility and, therefore, the agent's incentive to misreport. Providing more consumption today, compared to pure consumption smoothing, results in a negative drift for promised utility and, therefore, in immiserization.

Mathematically, the first-order condition of (21) includes the term $\bar{s}(u_1) \bar{s}'(u_1)$. The sensitivity $\bar{s}(u_1)$ is decreasing in u_1 , as may easily be checked. This implies that the optimal u_1 is strictly greater than $-r$. Therefore, the drift of W_t , which equals $W_t(r + u_1)$, is *negative*, since W_t is negative.

The result also holds in the presence of moral hazard, as shown in the Appendix.

PROPOSITION 10 (IMMISERIZATION) *For all parameters $(r, \lambda, \xi, \theta, \chi, \bar{\phi})$, the optimal contract implies a negative drift for W_t .*

5.2 Changes in Effort Cost

This section shows that an arbitrarily flat cost function for the effort of the agent may hurt the principal.

Recall the first-order condition for effort

$$\phi'(a(w)) = \bar{s}(-w) = \frac{\phi(a(w))}{\chi}.$$

When the effort cost parameter χ goes to zero, it becomes arbitrarily cheap for the agent to undertake any effort level.³⁸ If \bar{s} is strictly positive, this means that as χ goes to zero, the agent makes an arbitrarily large effort, which is very costly to the principal, as it results in arbitrarily large log utility (i.e., W_t gets arbitrarily close to zero). To avoid this situation, the principal has to reduce the sensitivity \bar{s} to a level arbitrarily close to zero. However, this reduces his ability to ensure truth-telling, the other channel through which the agent can deviate.

To keep inducing truth-telling, the principal has to reduce the magnitude of, v_y , the marginal benefit from lying. By an argument similar to the one used for the immiserization result (see Equation (22)), this can only be done by providing more immediate utility to the agent. Owing to agent's decreasing marginal utility, this gets arbitrarily costly to the principal. As a result, promised utility dives at a rate arbitrarily close to the discount rate r , and there is no consumption smoothing.

The principal also receives arbitrarily large cash flows from the agent's effort, which may offset the amount of consumption that he must provide to the agent. However, if the agent's initial promised utility is high enough, the cost exceeds the benefits, and the principal's payoff gets arbitrarily negative.³⁹

Let $\Pi(\chi)$ denote the principal's expected payoff under the optimal state-consistent contract. The following result is proved in the Appendix.

PROPOSITION 11 *For all parameters $(r, \lambda, \xi, \theta, \bar{\phi})$, if $w > -1/r$, $\Pi(\chi)$ diverges to $-\infty$ as χ goes to zero. If $w < -1/r$, $\Pi(\chi)$ diverges to $+\infty$ as χ goes to zero.*

The comparative statics with respect to the scaling parameter $\bar{\phi}$ are straightforward. From (19), a higher scaling parameter for the cost function does not affect the optimal sensitivity \bar{s} , and reduces the principal's objective only through the last term. This suggests that what matters most Proposition 11 is the curvature of the cost function rather than its scale: a

³⁸More precisely, the cost function becomes flat: the marginal cost of effort converges everywhere to zero.

³⁹As χ goes to zero, the immiserization effect becomes muted, as shown at the end of the proof of Proposition 11.

flatter cost functions makes all actions more similar from the agent's viewpoint, whereas a homogeneous increase or decrease of the cost functions does not affect the agent's preferences across actions (indeed, the agent's optimal effort cost $\phi(\bar{s}) = \bar{s}/\chi$ is independent from $\bar{\phi}$).

5.3 Impact of Noise

In contrast to the previous result, it is always in the principal's interest to reduce the noise, or volatility, of the agent's output, and the optimal sensitivity coefficient is decreasing in volatility.

PROPOSITION 12 *The optimal sensitivity s is decreasing in σ .*

Proof. The objective 19 is submodular in σ and s . The result follows from Topkis (1978).

6 Verification of Incentive Compatibility

One must verify that for any $\bar{s} < \theta$, the contract constructed in Section 4 is incentive compatible. Under the contract, the promised utility evolves as

$$dW_t = (r + u_1 - \phi_1)W_t dt + s(W_t)(dG_t + A_t dt + \lambda G_t dt + \sigma dB_t), \quad (23)$$

and the agent consumes $c(W_t) - G_t$, where $G_t = Y_t - X_t$. Therefore, the agent only cares about X_t and Y_t through their difference G_t . The agent's optimization is reduced to

$$v(w, g) = \sup_{L, A} E \left[\int_0^\infty e^{-rt} (u(c(W_t) - G_t) - \phi(A_t)) dt \right]$$

subject to $dG_t = L_t dt$, $G_0 = g$, and (23). The HJB equation for this problem is

$$0 = \sup_{l, a} \left\{ u(c(w) - g) - \phi(a) - rv(w, g) + v_w (rw - u(c(w)) + \phi(a(w)) + s(w)(a - a(w) + l + \lambda g)) + v_g l + \frac{1}{2}(s(w))^2 \sigma^2 v_{ww} \right\}, \quad (24)$$

where $u(\cdot)$, $\phi(\cdot)$, $a(\cdot)$, $c(\cdot)$ and $s(\cdot)$ have the forms given in Section 4.

One must show that there exists a solution to (24) such that $v(w, 0) = w$ for all w . This will establish that, under the proposed contract, the best the agent can achieve is his promised

utility, whenever the current report is correct. It is natural to conjecture a solution of the form $v(w, g) = wf(g)$ for some function f to determine. Incentive compatibility will be established if one finds a solution f such that $f(0) = 1$, meaning that when the gap is zero, the promised utility is exactly w . With this form, the first-order condition⁴⁰ with respect to a yields $\phi(a) = \frac{f(g)s(w)}{\chi}$, so the Bellman equation becomes, after simplification and dividing throughout by $(-w)$,

$$0 = \sup_{\ell} \left\{ u_1 \exp(\theta g) + f(g)(-u_1 + \bar{s}(\ell + \frac{1}{\chi}(\log f(g)) + \lambda g)) - f'(g)\ell \right\}. \quad (25)$$

The objective is linear in ℓ , which has unbounded domain. If the contract is not truthful, the agent therefore wants to lie at an infinite rate. To accommodate for this, the agent is now allowed to report jumps in the cash flows. This expands the reporting domain of the agent.

Let $W_{t+}(\Delta L)$ denote the promised utility of the agent after he reports a jump ΔL in the cash flow at time t . For contracts with a fixed sensitivity parameter, as considered here, a natural closure of the contract is to stipulate that

$$W_{t+}(\Delta L) = \exp(-\bar{s}\Delta L)W_t. \quad (26)$$

To see this, notice that if the agent lies at an arbitrarily large rate K between times t and $t + \varepsilon$, his promised utility satisfies, ignoring second-order effects, the dynamic equation

$$dW_t = \bar{s}(-W_t)K dt.$$

This yields $W_{t+\varepsilon} = \exp(-K\bar{s}\varepsilon)W_t$, and results in a gap change $G_{t+\varepsilon} = G_t + K\varepsilon$. Combining the last two equations yields $W_{t+\varepsilon} = \exp(-\bar{s}(G_{t+\varepsilon} - G_t))W_t$, which explains (26).

Report jumps amount to impulse controls on the part of the agent (see for example, Øksendal and Sulem (2004)). The HJB equation (25) becomes

$$0 = \max \left\{ \sup_{\ell \in \mathbb{R}} \left\{ u_1 \exp(\theta g) + f(g)(-u_1 + \bar{s}(\ell + \frac{1}{\chi}(\log f(g)) + \lambda g)) - f'(g)\ell \right\}, \right. \\ \left. \sup_{\Delta L \in \mathbb{R} \setminus \{0\}} \left\{ \exp(-\bar{s}\Delta L)f(g + \Delta L) - f(g) \right\} \right\}. \quad (27)$$

⁴⁰Note that the objective is strictly concave in a and the domain of a is open, so that the first-order condition pins down the unique optimum.

The function $f(g) = \exp(\bar{s}g)$ solves the equation. Indeed, with that value for f , the second term of the equation is always equal to zero. Therefore, it suffices to show that

$$\sup_{\ell} \left\{ u_1 \exp(\theta g) + \exp(\bar{s}g) \left(-u_1 + \bar{s} \left(\ell + \frac{1}{\chi}(\bar{s}g) + \lambda g \right) \right) - \bar{s} \exp(\bar{s}g)\ell \right\} \leq 0 \quad (28)$$

for all g . That term is independent of ℓ , and reduces to

$$u_1 \exp(\theta g) + \exp(\bar{s}g) \left(-u_1 + \bar{s} \left(\frac{1}{\chi}(\bar{s}g) + \lambda g \right) \right)$$

Convexity of the exponential implies that, for all g ,

$$\exp(\theta g) > \exp(\bar{s}g) + \exp(\bar{s}g)(\theta - \bar{s})g.$$

Since $u_1 < 0$, (28) will be satisfied if

$$\exp(\bar{s}g) \left(u_1(\theta - \bar{s})g + \bar{s} \left(\frac{1}{\chi}(\bar{s}g) + \lambda g \right) \right) \leq 0.$$

The second factor is zero, from (18), which concludes the proof.

An optimal control associated with the Bellman equation is to set $\Delta L = -g$ if $g \neq 0$ and $\Delta L = 0$ otherwise, and ℓ always equal to zero. It is optimal for the agent to i) always report truthfully if he has been truthful in the past, and ii) immediately correct any existing gap between real and reported cash flows. It means, in particular, that if the principal did not know the initial cash flow, the contract is still incentive compatible. This optimal control is essentially unique. Indeed, between two impulse controls, the first term of the Bellman equation must equal zero, which only holds for $g = 0$.

Ponzi Schemes

We also need to check that cannot gain by over-reporting forever or under-reporting forever. We argue here informally (to be completed). Over-reporting by a large amount and over a long period can never benefit the agent. Indeed, high reports get rewarded by a promised utility increase at rate \bar{s} , but reduce the net transfers to the agent which reduce his utility by a rate $\theta > \bar{s}$. Thus, neglecting second-order effects, large and durable over-reports cost too much consumption to the agent relative to the reward they generate.

Underreporting does not work either, due to the mean reversion of the cash flow process. If the agent makes very negative reports, the principal believes that the current cash flow is very low, and hence that the drift is very high, due to mean reversion. Therefore, further

negative news become more and more “surprising” by the principal, and are accordingly punished by an increasingly large drop in the continuation utility. Simple computations show that, in such case, the promised utility of the agent diverges to $-\infty$ in finite time: negative Ponzi schemes are not sustainable.

7 Renegotiation and Separating Contracts

Previous sections have focused on separating contracts. It is well known that, with renegotiation, the revelation principle need not apply.

To understand the impact of renegotiation on contracting, it is useful to consider the striking, extreme case where the cost of effort of the agent goes to zero. In that case, Section 5.2 has shown that the payoff of the principal, for any separating contract, becomes arbitrarily negative, provided that initial promised utility is high enough.

With commitment, the principal could easily avoid this problem. For example, consider a contract that proposes $C_t = Y_t + b$, for some constant b . With that specification, the agent’s payoff is independent of his reports: he always gets a total consumption $(Y_t + b) + (X_t - Y_t) = X_t + b$. The cost to the principal is simply b/r , which is finite. By choosing b appropriately, the principal can always achieved any given promised utility to the agent.⁴¹

Such contract is not renegotiation-proof. To illustrate, suppose that, after some time, the cash flow X_t becomes very high and, just for now, that the principal knows it. Then, the principal could propose the agent a low payment $b_1 < b$ in the short term, and a high payment $b_2 > b$ in the future. Owing to mean reversion, the agent expect his cash flow to go down in the future, and given the concavity of his utility function, may prefer this new contract. The principal can strictly improve his payoff with this contract.

Now suppose that the principal cannot observe the cash flow. He could still propose the above contract to the agent. If the cash flow is low, the agent will reject this contract, while if the cash-flow is high, he will accept it. Thus, not only will the contract be renegotiated in some cases, but the principal will in any case learn more about the agent’s type. Because the principal cannot commit not to renegotiate, he cannot commit not to learn more about

⁴¹Even with a very low cost, the payoff of the agent is always bounded, because of mean reversion: the higher the cash flow, and the more negative the cash-flow drift, for given effort. Therefore, the cash-flows cannot grow arbitrarily large.

the agent's type through such renegotiation proposal.

One might think that the previous argument is limited in scope, i.e., that the principal may learn something about the agent, but not the precise cash flow. Indeed, this limitation arises in the two-period model of Laffont and Tirole (1988), where the principal does not get the chance to propose yet another contract after partially learning the type of the agent. However, the structure of the present model suggests otherwise. With continuous time, the principal can propose arbitrarily many contracts in any small time interval, which potentially allows him to learn arbitrarily precisely the cash flow.⁴² Moreover, the argument made above is “self similar” in the sense that no matter how small the uncertainty is about the current cash flow, the principal could always propose a contract of the form above but more precisely targeted to exploit a small difference in cash flow levels, that tells him a bit more about the cash flow.

The previous discussion suggests that the principal cannot avoid learning about the agent's type. In a simplified version of the setting with finite horizon, discrete time, and a pure reporting problem, Strulovici (2011) shows that, the combined assumptions of cash-flow mean reversion and concave utility imply that the optimal renegotiation-proof contract is always separating.⁴³

Another potential issue concerns the possibility of renegotiation off the equilibrium path. If deviations were detectable, the principal could have an incentive to renegotiate the contract. In the setting of this paper, however, the all report histories may arise on the equilibrium path, since the drift of an Itô process only affects the probability distribution of its possible sample paths, not the set of paths that can occur.⁴⁴ Therefore, the principal never realizes that the agent has gone off the equilibrium path, and thus does not renegotiate the contract. Given this, it is optimal for the agent to be truthful, since the contract is incentive compatible.

⁴²That feature may also arise in a model of “dialogue” where the principal and the agent alternate contract offers and acceptance/rejection decisions. Such protocol need not lead to full type revelation, but is likely to result in more revelation than in a two-period model, where the principal can make only two contract proposals, in effect committing not to react to the last piece of information he gets from the agent.

⁴³There, renegotiation-proofness is the (standard) finite-horizon concept, based on backward induction.

⁴⁴This observation is, essentially, a qualitative version of Girsanov's theorem. See, e.g., Karatzas and Shreve (1991).

8 Permanent Outside Option

Suppose that the agent is allowed to leave the contract at any time and get a continuation utility $\underline{w} < 0$. This imposes the individual-rationality constraint $W_t \geq \underline{w}$ at all times. How does this new constraint affect renegotiation? The gist of the previous analysis is unchanged. First, any internally-consistent contract has continuation payoffs that only depend on the current state (w, y) . Second, comparing across cash-flow levels, the argument of Section 4.3.1 goes through, so that contractual variables should depend only on promised utility, not on the cash-flow level. Moreover, the principal's payoff function should vary across cash-flow levels according to Equation (13).

However, the constraint raises a difficulty for comparing contracts across initial promised utility. Indeed, starting for some contract C that is individually rational and incentive compatible given (w_1, y) for $w_1 > \underline{w}$, there is no guarantee that, for $w_2 \in (\underline{w}, w_1)$, the (w_2, y) -version of C will also be individually rational. Indeed, that version scales the continuation utility process W_t of the agent by a factor w_2/w_1 , compared to contract C , and may violate the individual rationality constraint. Therefore, the individual rationality constraint reduces the set of challengers to any given continuation contract, which prevents us from exploiting earlier comparisons to derive the closed-form formulas obtained in Section 4.3.2.

Conceptually however, the problem is very similar to the unconstrained case. At one extreme, for w far above \underline{w} , the optimal state-consistent contract should be very similar to optimal contract of the unconstrained case, and the payoff and contractual functions should be well approximated by the closed-form functions derived for that case. At the other extreme, if $w = \underline{w}$, the principal has very few options to keep the agent in the relationship. Indeed, the only contracts that guarantee that the constraint is not violated are those for which i) the sensitivity parameter of the promised utility is exactly zero (for otherwise the promised utility of the agent might drop below \underline{w}), and ii) the drift is positive (to push W_t higher away from \underline{w}), for example by providing a low utility flow. Such extreme contractual characteristics are clearly not required for w high above \underline{w} .

The cross-state comparison provides, even in the constrained case, valuable information about the principal's continuation payoff. Precisely, one direction of the unconstrained analysis carries over to the constrained case, providing a whole family of inequalities comparison for the principal payoffs. Suppose that w_2 is *higher* than w_1 . In that case, the (w_2, y) -version

of C does satisfy individual rationality if C did. This observation implies the following: for any individually-rational and state-consistent contract C , let $\Pi(w, y)$ denote the principal's payoff under any continuation contract of C following a history ending up with state (w, y) .⁴⁵

PROPOSITION 13 *For any states $(w, y), (w', y')$ such that $w' \geq w$,*

$$\Pi(w', y') \geq \frac{(y' - y)}{r + \lambda} + \frac{\log(w'/w)}{r} \left(\frac{1}{\theta} + \frac{1}{\chi(r + \lambda)} \right) + \Pi(w, y). \quad (29)$$

This is intuitive: the farther away one gets from the constraint w , and the more flexibility one has to choose consumption/effort processes achieving the given promised utility and, therefore, the higher the payoff of the principal can get relative to versions of more constrained contracts starting with lower promised utility.

In line with the previous argument, one may further conjecture that, for the optimal individually-rational renegotiation-proof contracts, the sensitivity factor is increasing in w (rather than constant for the unconstrained case), going from 0 for $w = \underline{w}$ to the unconstrained optimum \bar{s} (i.e., the maximizer of (19)), as w gets arbitrary large.

9 Discussion

9.1 Summary: Concepts of Renegotiation

This paper has introduced a new concept of renegotiation, which extends to stochastic games the notion of internal consistency, by exploiting natural comparisons across states. With stochastic games, some underlying state affects the physical environment of the players. However, players may be able to recognize that there are clear relations between the sets of feasible continuation games across different states, in the same way that they recognize, in a repeated-game setting, the relation between continuation payoffs across different histories. Here, state consistency was applied to the particular setting of a principal-agent relationship where the principal can make arbitrary transfer to the agents and has alone the initiative to propose a new contract and where the agent has specific utility and cost functions. The concept is portable, and can clearly be applied to other environments.⁴⁶

⁴⁵By internal consistency, the principal's payoff depends only on (w, y) .

⁴⁶For example, Gromb (1994) considers renegotiation of debt contracts with a binary state (whether investment occurred in the last period), and compares continuation contracts across these two states.

Strong Renegotiation-Proofness and State Consistency State consistency is a weak concept of renegotiation-proofness.⁴⁷ There are different ways to strengthen the notion. For example, in the context of repeated games, a contract is said to be strongly renegotiation proof if it is weakly renegotiation-proof, and there is no weakly renegotiation-proof contract that Pareto dominates it after any history. The notion is easily extended to stochastic games, requiring that for no value of the underlying state the continuation payoffs of the contract are Pareto dominated by those of a weakly renegotiation proof contract. Of course, there need not in general exist such contract. In the present setting, one may show that if a strongly renegotiation proof contract exists, then it has to be state consistent.⁴⁸ This result shows how continuation payoffs vary with the underlying state, as in Section 4.

9.2 Robustness and Extensions

Choice of a Production Technology One consequence of Section 5.2 is that, should the principal choose the production technology as captured by the effort cost function $\phi(\cdot)$, he may choose a technology with lower productivity, even if a high-productivity technology has a comparable or even lower cost. In contrast, Proposition 12 and Equation (19) imply that the principal's payoff is decreasing with the noise in the production technology, and that increased noise reduces the principal's ability to reward the agent's effort.

Reporting Constraints The contracts constructed in this paper continue to be incentive compatible if the agent has constraints on cash-flow reports and transfers. For example, the agent could be prevented from over-reporting his cash flows. Indeed, such constraints only restrict the agent's strategy space and, hence, the set of possible deviation. Moreover, the entire analysis goes through for that case: the transformation group is unchanged, and the renegotiation-proof contracts have the same functional form.

Private Savings Private savings do not offer the agent an enlarged reporting space, since he is already able to make arbitrary reports. Furthermore, using the extended reporting space described in Section 6, where the agent can make jumps in his reports, the agent can

⁴⁷Indeed, for the case of repeated games, the concept reduces to internal consistency (or weak renegotiation-proofness).

⁴⁸The argument can be sketched easily: for any two states, say (w, x) and (\tilde{w}, \tilde{x}) , a strongly renegotiation proof contract C must not result in a lower payoff, at state (\tilde{w}, \tilde{x}) , than its transformation $\mathcal{G}_{(\tilde{w}, \tilde{x}), (w, x)}(C)$, which is weakly renegotiation proof if C is.

replicate whatever change of consumption the agent could obtain through private savings. To illustrate, suppose that the agent borrows an amount $b_t = b > 0$ for $t \in [0, T]$, and repays it by some constant annuity $b_t = -c$ for $t \in [T, \infty]$ such that $\int_0^\infty e^{-rt} b_t = 0$, achieving a real consumption $\bar{C}_t = C_t + b_t$, if the contract provides consumption C_t . One way to replicate this consumption stream without saving is for the agent to report a lump sum drop in his income at time zero. For a report $-\Delta L$, the agent's promised utility jumps downwards to $W_+ = \exp(s\Delta L)W_0$. The jump in consumption provided by the principal is $\Delta C = -s\Delta L/\theta$, from Proposition 8. The real consumption change of the agent, however, is $\Delta C + \Delta L = (1 - \frac{s}{\theta})\Delta L$, which is strictly positive, remembering that $s < \theta$ (see (17) and the comment later). Choosing ΔL appropriately yields the jump b of consumption that the agent could obtain by privately borrowing. By adjusting his lie continually, the agent can make sure to obtain the consumption stream that he would have obtained, should he have borrowed. The repayment can be obtained similarly.

Additional Signal A natural extension would allow the principal to receive a secondary signal about the agent's action. The promised utility would then depend on both the agent's report and on that signal, allowing another the principal to use an additional instrument, the sensitivity to that other signal. This would enlarge the set of incentive compatible contracts to a two-dimensional set and mitigate the impact of a flatter cost function of the agent on the principal's payoff.

10 Appendix

Proof of Proposition 1

State consistency already implies that after any history leading up to any state (\tilde{w}, \tilde{y}) and continuation contract \tilde{C} , $\Pi(\tilde{C}) \geq \Pi(\mathcal{G}_{(\tilde{w}, \tilde{y}), (w, y)}(C))$, for otherwise the principal could use the transformation $\mathcal{G}_{(\tilde{w}, \tilde{y}), (w, y)}(C)$ instead of \tilde{C} . Now suppose that the inequality is strict. Monotonicity of $\mathcal{G}_{(w, y), (\tilde{w}, \tilde{y})}$ implies that $\Pi(\mathcal{G}_{(w, y), (\tilde{w}, \tilde{y})}(\tilde{C})) > \Pi(\mathcal{G}_{(w, y), (\tilde{w}, \tilde{y})} \circ \mathcal{G}_{(\tilde{w}, \tilde{y}), (w, y)}(C)) = \Pi(C)$, which contradicts state consistency of C : the principal would do better by starting from $\mathcal{G}_{(w, y), (\tilde{w}, \tilde{y})}(\tilde{C})$ than starting from C , given (w, y) . ■

Proof of Proposition 4

We prove more generally that if a strategy (L, A) adapted to the filtration generated by the

underlying Brownian motion dominates another adapted strategy (L', A') under the initial contract, then it also dominates it under the new contract. By construction,

$$\begin{aligned} E \left[\int_0^\infty e^{-rt} \left(u \left(\mathcal{C}(Y_s : s \leq t) - \int_0^t L_s ds \right) - \phi(A_t) \right) dt \right] \\ \geq E \left[\int_0^\infty e^{-rt} \left(u \left(\mathcal{C}(Y'_s : s \leq t) - \int_0^t L'_s ds \right) - \phi(A'_t) \right) dt \right] \end{aligned} \quad (30)$$

where

$$dY_t = \left[L_t + A_t + \left(\xi - \lambda \left(Y_t - \int_0^t L_s ds \right) \right) \right] dt + \sigma dB_t$$

and

$$dY'_t = \left[L'_t + A'_t + \left(\xi - \lambda \left(Y'_t - \int_0^t L'_s ds \right) \right) \right] dt + \sigma dB_t$$

and the initial conditions $Y_0 = Y'_0 = y$.

Now consider the initial condition \hat{y} . The reporting processes under strategies (L, A) and (L', A') are respectively

$$d\hat{Y}_t = \left[L_t + A_t + \left(\xi - \lambda \left(\hat{Y}_t - \int_0^t L_s ds \right) \right) \right] dt + \sigma dB_t$$

and

$$d\hat{Y}'_t = \left[L'_t + A'_t + \left(\xi - \lambda \left(\hat{Y}'_t - \int_0^t L'_s ds \right) \right) \right] dt + \sigma dB_t$$

and subject to the initial condition $\hat{Y}_0 = \hat{Y}'_0 = \hat{y}$.

The reconstructed processes are, by definition, such that

$$d\tilde{Y}_t = d\hat{Y}_t - (\xi - \lambda\hat{Y}_t)dt + (\xi - \lambda\tilde{Y}_t)dt$$

and

$$d\tilde{Y}'_t = d\hat{Y}'_t - (\xi - \lambda\hat{Y}'_t)dt + (\xi - \lambda\tilde{Y}'_t)dt$$

subject to $\tilde{Y}_0 = \tilde{Y}'_0 = y$. Combining the previous equations yields

$$d\tilde{Y}_t = \left[L_t + A_t + \left(\xi - \lambda \left(\tilde{Y}_t - \int_0^t L_s ds \right) \right) \right] dt + \sigma dB_t$$

$$d\tilde{Y}'_t = \left[L'_t + A'_t + \left(\xi - \lambda \left(\tilde{Y}'_t - \int_0^t L'_s ds \right) \right) \right] dt + \sigma dB_t$$

subject to the conditions $\tilde{Y}_0 = \tilde{Y}'_0 = y$. The comparison between objective functions is identical to (30), and subject to the same dynamic equations and constraints, which establishes that the strategy (L, A) dominates (L', A') under the new contract as well. \blacksquare

Proof of Proposition 10

Let $x = \phi_1 - u_1$. It suffices to show that, for \bar{s} maximizing (19), $x < r$. From the expressions of $\phi_1(\bar{s})$ and $u_1(\bar{s})$,

$$\phi(x) = \frac{\bar{\theta}x}{\bar{\theta} + \lambda + x}$$

where $\bar{\theta} = \theta/\chi$, and

$$u(x) = -\frac{x(\lambda + x)}{\bar{\theta} + \lambda + x}$$

The objective can therefore be expressed in terms of x . Its derivative with respect to x is

$$(\kappa + 1)\frac{1}{x} - (\kappa + 1)\frac{1}{\bar{\theta} + \lambda + x} + \frac{1}{\lambda + x} - \frac{\kappa + 1}{r} (1 + \sigma^2\chi^2\phi(x)\phi'(x)).$$

Suppose that $x \geq r$. Then, using that $\kappa = \bar{\theta}/(r + \lambda)$ it is easy to show that

$$-(\kappa + 1)\frac{1}{\bar{\theta} + \lambda + x} + \frac{1}{\lambda + x} \leq 0.$$

Since also $\phi'(x) > 0$, the derivative is negative for $x \geq r$, showing that the optimum is achieved for $x < r$. ■

Proof of Proposition 11

Recall from Proposition 9 that

$$\Pi(\chi) = \kappa \log(\phi_1(\bar{s})) + \log(-u_1(\bar{s})) + \frac{\kappa + 1}{r} \left(r + u_1(\bar{s}) - \phi_1(\bar{s}) - \frac{1}{2}\sigma^2\bar{s}^2 \right) + (\kappa + 1) \log(-w) - \kappa \log(\bar{\phi}), \quad (31)$$

where $\phi_1(\bar{s}) = \bar{s}/\chi$, $\kappa(\chi) = \frac{\theta}{\chi(r+\lambda)}$, $u_1(\bar{s}) = -\bar{s}(\chi\lambda + \bar{s})/(\chi(\theta - \bar{s}))$, and $\bar{s}(\chi) \in [0, \theta]$.

As χ goes to zero, κ is of order $1/\chi$ and, after neglecting second-order terms and terms independent from \bar{s} , the objective equals

$$\kappa \log \bar{s} + \frac{\kappa}{r} \left(\frac{-\bar{s}\theta}{\chi(\theta - \bar{s})} - \frac{1}{2}\sigma^2\bar{s}^2 \right).$$

The maximum can only be attained for \bar{s} arbitrarily small, otherwise the second term would be of order $(1/\chi^2)$ (taking into account the factor κ), arbitrarily negative, and dominate all other terms. Precisely, \bar{s} must be at most of order χ . Let $\bar{s} = \alpha\chi + o(\chi)$, for some $\alpha \geq 0$ to be chosen by the principal. The objective, after dropping second-order terms and terms independent from α , becomes

$$\kappa \log \alpha + \frac{\kappa}{r}(-\alpha).$$

Therefore, the optimum sensitivity is equal to $\bar{s} = r\chi + o(\chi)$. This implies that $\phi_1(\chi) \sim r$ and $u_1(\chi) \sim -r^2\chi/\theta$. The objective is equal to

$$\kappa(\chi) \log(r(-w)) + o\left(\frac{1}{\chi}\right).$$

This shows that $\Pi(\chi)$ diverges to $+\infty$ if $w < -1/r$ and to $-\infty$ if $w > -1/r$.⁴⁹

The drift of the promised utility is equal to $r + u_1 - \phi_1$. Since, $\phi_1 \sim r$ while $u_1(\chi) \sim -r^2\chi/\theta$, one concludes that

$$\lim_{\chi \rightarrow 0} r + u_1(\chi) - \phi_1(\chi) = 0.$$

This shows that immiserization becomes arbitrarily slow as χ goes to zero. ■

References

- ABREU D., PEARCE, D. (1991) “A Perspective on Renegotiation in Repeated Games,” in *Game Equilibrium Models II*, (Reinhard Selten Ed.), Springer-Verlag, Berlin.
- ABREU D., PEARCE, D., AND E. STACCHETTI (1993) “Renegotiation and Symmetry in Repeated Games,” *Journal of Economic Theory*, Vol. 60, pp. 217–240.
- ASHEIM, G. (1991) “Extending Renegotiation-Proofness to Infinite Horizon Games,” *Games and Economic Behavior*, Vol. 3, pp. 278–294.
- BATTAGLINI, M. (2007) “Optimality and Renegotiation in Dynamic Contracting,” *Games and Economic Behavior*, Vol. 60, pp. 213–246.
- BENOÎT, J.-P., KRISHNA, V. (1993) “Renegotiation in Finitely Repeated Games,” *Econometrica*, Vol. 61, pp. 303–323.
- BERGEMANN, D., HEGE, U. (2005) “The Financing of Innovation: Learning and Stopping,” *Rand Journal of Economics*, Vol. 36, pp. 719–752.
- BERNHEIM, B.D., RAY, D. (1989) “Collective Dynamic Consistency in Repeated Games,” *Games and Economic Behavior*, Vol. 1, pp. 295–326.

⁴⁹That result is illustrated numerically for parameter values $r = \lambda = 5\%$, $\theta = 1$, $\sigma = 0.2$ and $\chi = 0.01$. In that case, $\Pi \sim 686$ for $w = -40$, while $\Pi \sim -2310$ for $w = -2$. For these values of w , the payoffs respectively get arbitrarily positive and arbitrarily negative as χ gets closer to zero.

- BESTER, H., STRAUZ, R. (2001) “Contracting with Imperfect Commitment and the Revelation Principle: The Single Agent Case,” *Econometrica*, Vol. 69, pp. 1077–1098.
- DEMARZO, P., FISHMAN, M. (2007) “Optimal Long-Term Financial Contracting,” *Review of Financial Studies*, Vol. 20, pp. 2079–2128.
- DEMARZO, P., SANNIKOV, Y. (2006) “Optimal Security Design and Dynamic Capital Structure in a Continuous-Time Agency Model,” *Journal of Finance*, Vol. 61, pp. 2681–2724.
- DEWATRIPONT, M. (1989) “Renegotiation and Information Revelation over Time: The Case of Optimal Labor Contracts,” *Quarterly Journal of Economics*, Vol. 104, pp. 589–619.
- DOEPKE, M., TOWNSEND, R. (2006) “Dynamic Mechanism Design with Hidden Income and Hidden Actions,” *Journal of Economic Theory*, Vol. 126, No. 1, pp. 235–285.
- FARRELL, J., MASKIN, E. (1989) “Renegotiation in Repeated Games,” *Games and Economic Behavior*, Vol. 1, pp. 327–360.
- FERNANDES, A., PHELAN, C. (2000) “A Recursive Formulation for Repeated Agency with History Dependence,” *Journal of Economic Theory*, Vol. 91, pp. 223–247.
- FLEMING, W., SONER, H. M. (2006) *Controlled Markov Processes and Viscosity Solutions*, Second Edition, Springer, New York, N.Y.
- FUDENBERG, D., TIROLE, J. (1990) “Moral Hazard and Renegotiation in Agency Contracts,” *Econometrica*, Vol. 58, No. 6, pp. 1279–1319.
- FUKUSHIMA, K., WAKI, Y. (2009) “Computing Dynamic Optimal Mechanisms When Hidden Types Are Markov,” Working Paper, University of Minnesota.
- GARFAGNINI, U. (2010) “Delegated Experimentation,” Working Paper, Northwestern University.
- GARRETT, D., PAVAN, A. (2010a) “Dynamic Managerial Compensation: On the Optimality of Seniority-based Schemes,” Working Paper, Northwestern University.
- GARRETT, D., PAVAN, A. (2010b) “Managerial Compensation in a Changing World,” Working Paper, Northwestern University.

- GIBBONS, R., MURPHY, K. (1992) “Optimal Incentive Contracts in the Presence of Career Concerns: Theory and Evidence,” *Journal of Political Economy*, Vol. 100, No. 3, pp. 468–505.
- GOLOSOV, M., KOCHERLAKOTA, N., AND A. TSYVINSKI (2003) “Optimal Indirect and Capital Taxation,” *Review of Economic Studies* Vol. 70, pp. 569–587.
- GOLOSOV, M. TSYVINSKI, A. (2006) “Designing Optimal Disability Insurance: A Case for Asset Testing,” *Journal of Political Economy*, Vol. 114, pp. 257–279.
- GREEN, E. J. (1987) “Lending and the Smoothing of Uninsurable Income,” in *Contractual Arrangements for Intertemporal Trade*, E.C. Prescott and N. Wallace (Editors), Minneapolis, University of Minnesota Press.
- GROMB, D. (1994) “Renegotiation in Debt Contracts,” Working Paper, London Business School.
- HÖRNER, J., SAMUELSON, L. (2010) “Incentives for Experimenting Agents,” Discussion Paper No. 1726, Cowles Foundation.
- KAPICKA, M. (2007) “Efficient Allocations in Dynamic Private Information Economies with Persistent Shocks: A First-Order Approach.” Working paper, University of California, Santa Barbara.
- KARATZAS, I., SHREVE, S. (1991) *Brownian Motion and Stochastic Calculus*, Springer, New York, N.Y.
- LAFFONT, J.-J., TIROLE, J. (1988) “The Dynamics of Incentive Contracts,” *Econometrica*, Vol. 56, No. 5, pp. 1153–1175.
- LAFFONT, J.-J., TIROLE, J. (1990) “Adverse Selection and Renegotiation in Procurement,” *Review of Economic Studies*, Vol. 57, No. 4, pp. 597–625.
- MILGROM, P., SEGAL, I. (2002) “Envelope Theorems for Arbitrary Choice Sets,” *Econometrica*, Vol. 70, No. 2, pp. 583–601.
- ØKSENDAL, B., SULEM, A. (2004) *Applied Stochastic Control of Jump Diffusions*, Springer Verlag, Berlin, Germany.

- PEARCE, D. (1987) “Renegotiation-Proof Equilibria: Collective Rationality and Intertemporal Cooperation,” *Cowles Foundation Discussion Paper*, No. 855, Yale University.
- SANNIKOV, Y. (2008) “A Continuous-Time Version of the Principal-Agent Problem,” *Review of Economic Studies*, Vol. 75, No. 3, pp. 957-984.
- SPEAR, S., SRIVASTAVA, S. (1987) “On Repeated Moral Hazard with Discounting,” *Review of Economic Studies*, Vol. 53, pp. 599–617.
- STRULOVICI, B. (2011) “Contracting and Renegotiation: When are Optimal Contracts Separating?,” Working Paper, Northwestern University.
- TCHISTYI, A. (2006) “Security Design with Correlated Hidden Cash Flows: The Optimality of Performance Pricing,” Working Paper, Stern School of Business.
- THOMAS, J., WORRALL, T. (1990) “Income Fluctuation and Asymmetric Information: An Example of a Repeated Principal-Agent Problem,” *Journal of Economic Theory*, Vol. 51, pp. 367–390.
- TOPKIS, D. M. (1978) “Minimizing a Submodular Function on a Lattice,” *Operations Research*, Vol. 26, pp. 305–321.
- WILLIAMS, N. (2009) “Persistent Private Information,” forthcoming in *Econometrica*.
- ZHANG, Y. (2009) “Dynamic Contracting with Persistent Shocks,” *Journal of Economic Theory*, Vol. 144, pp. 635–675.