Contracts, Information Persistence, and Renegotiation*

Bruno Strulovici

November 7, 2011

Abstract

This paper studies how renegotiation and information persistence shape long-term contracts in principal-agent relationships. Truthful contracts that are renegotiation-proof, according to a concept tailored to account for persistence in the agent's type, are characterized by their sensitivity to the reports of the agent. The sensitivity of the *optimal* renegotiation-proof contract is increasing in information persistence and in the discount rate of the agent, and causes immiserization. Renegotiation-proof contracts are self-correcting off the equilibrium path. These results still hold when the agent is also subject to moral hazard. In that case, a lower cost of effort of the agent can reduce the payoff of the principal by increasing the severity of the agency problem.

1 Introduction

Financial bailouts, mortgage regulations, monetary and fiscal policies, and corporate pay cuts during recessions are instances in which regulators, firms, and other economic agents have the possibility to renegotiate implicit or explicit formal contracts in reaction to their economic environment. In these environments, the types of the agents (income, skills, output) are often privately observed and correlated over time.

*Northwestern University. I am particularly grateful to Alessandro Pavan, Hector Chade, Roozbeh Hosseini, Jean Tirole, and Daron Acemoglu, Philippe Aghion, Bruno Biais, Mehmet Ekmekci, Daniel Garrett, Brett Green, Christian Hellwig, Johannes Hörner, Parag Pathak, Christine Parlour, Francisco Ruiz-Aliseda, Larry Samuelson, Ilya Segal, Itai Sher, Juuso Toikka, Mike Whinston, Alex Wolitzky 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, the Harvard/MIT theory seminar, U.C. Berkeley (Haas School of Business), and Boston University (School of Management) for questions and comments.

Despite its importance in economics, renegotiation has resisted numerous formalization attempts by theorists into a single, universally accepted concept. Even in the simpler context of repeated games, where no exogenous state exists, no consensus has emerged about the right concept of renegotiation. For example, *internal consistency* compares only the continuation payoffs of a given equilibrium, providing a conceptual lower bound for what renegotiation should, intuitively, entail.[1] At the extreme opposite, concepts such as *strong renegotiation proofness* (Farrell and Maskin, 1989) require comparisons across all possible equilibria, and are so strong that they may eliminate all equilibria. Other concepts amount to a fixed point problem determining the set of equilibria that are renegotiation proof (for example the "social norms" by Asheim (1991)). Such concepts lead to problems of existence and multiplicity for the set of renegotiation-proof contracts.

These problems would seem to become even more severe in environments, pervasive in economics, where agents are subject to persistent and privately observed shocks. In that case, however, a strengthening of internal consistency can prove powerful enough to drastically reduce the set of renegotiation-proof contracts and yield clear predictions about their properties. The idea is to compare continuation payoffs of a given equilibrium not only across histories leading to the same state, but also across histories leading to *distinct* states, thereby obtaining a concept of *state consistency*. If the state affects the set of feasible payoffs, as in the case of Markovian cash flows generated by an agent, one cannot directly compare continuation equilibria following distinct states. However, one may be able to *transform* a continuation equilibrium starting from a given state into a continuation equilibrium starting from a different state, and compare the continuation payoffs of transformations leading to the same state. State consistency is weaker than strong renegotiation proofness, and provides a natural way to construct social norms for stochastic games.

This paper has two related objectives. The first one is to study the properties of renegotiation-proof contracts when the agent has private persistent information. The second one is to explore the concept of state consistency and illustrate it in a contracting environment.

The agent generates mean-reverting cash-flows whose evolution depends on the agent's effort and on exogenous shocks. Both efforts and cash flows are privately observed by the agent. The agent continually reports and transfers cash flows to the principal, possibly lying and diverting some of the cash flows. When the agent has exponential utility and effort cost, a natural class of contract transformations emerges, which creates a rich and well structured class of challengers to the continuation contract, at any time. Given any incentive compatible contract, the challengers are entirely constructed from continuations of the initial contract and transformations thereof. Renegotiation-

[1] The concept is due to Bernheim and Ray (1989). A closely related concept is that of "weak renegotiation proofness," introduced by Farrell and Maskin (1989). Despite its apparent innocuity, the concept has been criticized by Asheim (1991) and Abreu and Pearce (1991) for the time invariance that it implicitly assumes.

proofness means that the continuation contracts do weakly better than their challengers. The concept proves so powerful in this context that it reduces the set of renegotiation-proof contracts to a single-parameter family, which is easily analyzed.

The paper characterizes the set of all truthful, 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 the agent's incentive to truthfully report his cash flows and his incentive for effort. For any such contract, all contractual variables have exact formulas as a function of the sensitivity parameter and the continuation utility of the agent. The sensitivity parameter describes how the agent's continuation utility varies with his reports, and can take any value between 0 and the coefficient of absolute risk aversion of the agent.

For those contracts, the agent wants to report cash-flows truthfully not only on the equilibrium path, but also *after any possible deviation.* Even if, at time zero, the principal has a mistaken belief about the initial cash-flow, it is strictly optimal for the agent to immediately correct this mistake.[2]

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*: almost surely, it becomes arbitrarily negative as time elapses. This extends results obtained for the case of full commitment when the agent has a strictly concave utility function.[3]

The sensitivity of the optimal contract is increasing in the persistence of the agent's information. One intuition for this result is that persistence increases the importance of the current report, as it makes that report more relevant for the future, and thus the contract should be more sensitive to this information.

The optimal sensitivity is also increasing with the discount *rate* of the agent: the less patient the agent is, and the more sensitive the contract is to the reports. Intuitively, a more impatient agent is more tempted by short-term deviations, and needs stronger incentives to take his future utility into account. Since effort is also determined by the sensitivity of the contract, this leads to a counter-intuitive result: the higher the discount rate of the agent, and the more effort he puts in equilibrium.

Furthermore, the sensitivity of the optimal contract is decreasing in the magnitude of exogenous

[2]Such 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.

[3]See Thomas and Worrall (1990).

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 the 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 the agent's marginal value of lying, which, in turn, can only be obtained by an providing arbitrarily large utility flow.[4] 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): the agent may mix between reporting his true type and another type. 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

[4]See Section 5.2 and, in particular, Equation (14).

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 proposes a way to analyze renegotiation in stochastic games and contracting environments with persistent types.[5] It builds on concepts of renegotiation-proofness introduced for repeated games by Bernheim and Ray (1989) and Farrell and Maskin (1989).[6] Gromb (1994) studies a binary-state dynamic model of debt contracts, and compares continuation payoffs across the two states, in a way that is similar to the approach taken in this paper. Like all the papers mentioned so far, the concept of renegotiation in this paper is set theoretic: agents move a new equilibrium (or contract) when it is Pareto optimal to do so. Recent work in the context of repeated games has explored another approach where the bargaining game between agents at each period is explicitly modeled.[7]

The paper also builds on the literature on dynamic contracting with persistent private information initiated by Fernandes and Phelan (2000).[8] In addition to the employer-employee interpretation, the paper can be interpreted as an insurance model or an optimal taxation model where the agent reports income and gets subsidized or taxed according to his reports.[9] Renegotiation may take the form of a change in regulation that the social planner cannot commit not to perform, in the tradition of Kydland and Prescott (1977). The presence of moral hazard and persistent information

[5]Laffont and Tirole (1988,1990) Dewatripont (1989), and Fudenberg and Tirole (1990) study renegotiation in finite period models. 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 conceptual problems arising in repeated games and stochastic games with an infinite horizon. While conceptually simpler, renegotiation in finite horizon and discrete time also raises difficulty. For example, Benoît and Krishna (1993) show that the shape of renegotiation-proof contracts can largely depend on the number of periods.

[6]See also Pearce (1987), Asheim (1991), Abreu and Pearce (1991), Abreu et al. (1993), and the discussion of Section 9.

[7]Recent work includes Kranz and Ohlendorf (2010) and Miller and Watson (2011) and is closer to the relational contract literature.

[8]The 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.

[9]See Golosov et al. (2003), Golosov and Tsivinsky (2006), and Farhi and Werning (2010).

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 managerial compensation contracts with full commitment, when the type of the manager is persistent. Many features of the model in this paper are based on Williams (2011), who focuses on full commitment and a pure reporting problem.[10] 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. A similar problem arises in DeMarzo and Sannikov (2011), where both the principal and the agent can learn the agent's skill, but the agent privately observes his effort. Sannikov (2011) considers a moral hazard problem where the agent's action have long-term consequences, as in the present paper. None of these papers considers renegotiation.

The paper is organized as follows. Section 2 presents the contractual setting of the paper. Section 3 introduces *state consistency*, a new concept of renegotiation-proofness for stochastic games, which is applied here to a contracting environment. Section 4 characterizes state-consistent contracts in the principal-agent setting. Section 5 establishes several comparative statics: i) The sensitivity of the optimal contract is increasing with persistence and the discount *rate*, and decreases with the noise volatility; ii) The optimal renegotiation-proof contract always exhibits immiserization: the agent's promised utility drifts to minus infinity almost surely; iii) Increasing (at not cost) the productivity of the agent can reduce the principal's payoff. Section 6 shows that all state-consistent contracts are incentive compatible, even off the equilibrium path (i.e., the agent has a strict 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 other concepts of renegotiation and extensions of the model. Proofs omitted from the main text

[10] 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 follows a first-order approach. Fukushima and Waki (2009) propose a numerical analysis for a setting with persistent private information. Doepke and Townsend (2006) introduce a numerical method to analyze the optimal contract with moral hazard and adverse selection. In 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. All these papers assume full commitment.

are in the Appendix.

2 Setting

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

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

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, 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: $\phi(a) = \bar{\phi}\exp(\chi a)$. If χ is infinite, the agent chooses $A_t = 0$ for any incentive level, because the cost of effort is flat on $(-\infty, 0)$ and infinite on $(0, +\infty)$. The problem thus reduces to a pure cash flow diversion model. Much of the analysis and the intuition can be obtained from just looking at that case, and the reader may wish to set A_t and $\phi(A_t)$ identically to 0 and focus on the pure cash flow diversion model for a first reading of the paper.[11]

The Brownian Motion B is the exogenous source uncertainty in the model. There is a fixed a probability space $(\Omega, \mathcal{F}, P)$ satisfying the usual conditions, such that each outcome ω is identified with a path realization for B.[12]

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

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

L_t is the rate at which the agent "lies" about the true increment dX_t of the cash-flow.[13]

[11] Beyond obvious realism, there are several reasons to consider moral hazard in addition to the reporting problem. First, the interaction between the two agency problems yields some surprising results, as explored in Section 5. Second, moral hazard takes a particularly relevant form: the agent's current actions have a long-term impact, which is closely related to persistence of agent's type. Third, all the results obtained in the pure cash-flow diversion case still hold when moral hazard is added, which shows their robustness.

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

[13] Such 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. One may put constraints on the lies of the agents, for example by requiring that $Y_t \leq X_t$. The contracts derived here obviously remain incentive compatible in the presence of such constraints, and their form is unaffected by that.

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

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

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.[14]

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, as described by Equation (3).

The principal observes only the reports $\{Y_t\}$, but the initial cash-flow is publicly known:[15] $Y_0 = X_0$.

The agent's strategy consists of a lying process L and an effort process A adapted to the agent's information X.[16]

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

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

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, and does not affect the results.)

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

[15] As mentioned earlier, this will turn out not to matter for renegotiation-proof contracts, because the agent finds strictly optimal to correct any previous reporting mistake.

[16] X determines Y, given the agent's reporting strategy.

The principal's expected payoff is

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

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.[17] Similarly, the principal is at all times fully committed to providing the agent with 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,[18] the promised utility of the agent satisfies

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

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.,

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

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.[19]

The agent's problem is to solve

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

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

[18] See, e.g., Karatzas and Shreve (1991)

[19] This 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.

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 reports, while 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 Concepts of Renegotiation

In the context of repeated games, an equilibrium is said to be *weakly renegotiation-proof* (WRP) (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.[20]

Applied to stochastic games, internal consistency has a natural extension: say that an equilibrium is weakly renegotiation-proof, or internally consistent, if there are no two histories *leading to the same state*, such that the continuation payoffs following the first history Pareto dominate those following the second history. This extension does not make comparisons across states, and has therefore limited power. In particular, it implies that the players' payoffs are Markovian, but not necessarily that the equilibrium itself is Markovian.[21]

In the present setting, the "state" is the current cash-flow generated by the agent. Since we are analyzing a contracting rather than an equilibrium problem, we will take the state to also include the continuation utility of the agent, and a contract will be said to be 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.[22,23]

[20]Although 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).

[21]To guarantee that any equilibrium is Markovian, stronger assumptions are needed, as discussed in Section 4.

[22]If the agent is indifferent between an old and a new contract, he is assumed to accept the new contract. The principal can always give an infinitesimal share of the gain to the agent to convince him to switch.

[23]This 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, and need not *a priori* be true. See Section 4.1.

To see how internal consistency may be strengthened to analyze 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.[24] This cognitive ability should extend to other natural comparisons, as illustrated next.

Suppose that we are given an incentive-compatible contract $C = \{C_t\}_{t\geq 0}$ yielding a promised utility w_1 to the agent and a payoff π_1 to the principal and, for simplicity, that the agent has linear utility: $u(c) = c$. The contract obtained by translating consumption from the initial contract to $\tilde{C}_t = C_t + r(w_2 - w_1)$ provides utility w_2 to the agent. Let π_2 denote the payoff of the new contract for the principal. Now suppose that, starting from the contract C, the agent reaches some history at which his promised utility is $W_t = w_2$ and the continuation for the principal is equal to $\Pi_t < \pi_2$. In such scenario, the principal can reason that, by replacing the continuation contract by $\tilde{C}$, reset from time 0, he will achieve the strictly higher payoff π_2, and provide the same utility w_2 to the agent. Such comparison enlarges the class of challengers to a given continuation contract, by adding transformations of the initial contract that make it compatible with another state (here, a different level of continuation utility).

Such comparison can be motivated in different ways. Firstly, provided that the transformation is simple enough, it is realistic to assume that the principal can indeed come up with such comparison.

Secondly, the kind of challengers obtained through such transformations is a *subset* of the class of challengers obtained by stronger existing concepts of renegotiation. Consider, in particular, Farrell and Maskin's concept of strong renegotiation proofness, adapted to the present contracting environment. A WRP is strongly renegotiation-proof (SRP) if there is no history after which the principal's continuation payoff can be strictly improved by the continuation payoff of another WRP contract. Clearly, an SRP contract must withstand the sort of comparisons described here.

Thirdly, the type of comparisons established in the present paper can be seen as a way to generate a *social norm* for renegotiation-proof contracts, as studied by Asheim (1991).[25] A social norm $\mathcal{N}$ is, by definition, a set of equilibria. $\mathcal{N}$ is said to be internally stable if any two equilibria belonging to the social norm are not Pareto ranked, and to be externally stable if any equilibrium that is not part of the social norm and that dominates (for some continuation) an equilibrium of the social norm is dominated (for some continuation) by an equilibrium in the social norm. This paper provides a way to construct "social norms" (in a contractual environment, but the idea could be applied to more general stochastic games) by including in the social norm all transformations of any contract

[24] Of 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.

[25] A related concept is studied by Abreu et al. (1993)

that already belongs to the norm.

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 held, 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 should have the same payoff as the initial contract C. We will say that $\mathcal{G}$ is a *transformation group* operating on the set of contracts.[26]

The transformation group constructed in this paper involves two layers of challengers. First, we will show that drift linearity of the cash-flow implies that if a contract is IC given (w, y), then the same contract, seen as a functional of the underlying Brownian Motion (describing exogenous shocks) inferred from the reports is also IC given (w, y') for any y'. Renegotiation-proofness with respect to such transformation is closely related to the following property: the consumption provided to the agent is, at each time, only a function of his promised utility, not of the cash flow that he generates (reported cash flow affects the promised utility of the agent, and through this channel, his consumption). Second, exponential utility 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).

The transformation group thus has a further interpretation. Because each additional layer of challengers generated by the transformation group narrows the set of renegotiation-proof contracts and, therefore, increases the set of properties that they have, the present analysis may be seen as an *axiomatization* of the properties of a contract, based on renegotiation-proofness: linearity of the cash flow drift implies cash-flow independence of consumption, while exponential utility implies

[26] Although not necessary for the results, the transformations studied in the paper are also compatible in the following sense: $(\mathcal{G}_{w_3,w_2} \circ \mathcal{G}_{w_2,w_1})(C) = \mathcal{G}_{w_3,w_1}(C)$. $\mathcal{G}$ is usually not a group in the algebraic sense, because the composition $\mathcal{G}_{w_3,w_2} \circ \mathcal{G}_{w_1,w_0}$ only needs to be defined when $w_2 = w_1$. However, for the group constructed in the next section, the composition is easily defined when $w_2 \neq w_1$.

that flow utility is proportional to promised utility.

For now, we return to the definition of renegotiation-proofness based on general transformation groups. For any incentive-compatible contract C, let $\Pi(C)$ denote the principal's payoff under that contract.

Definition 1 *A contract C that is incentive compatible given (w, y) is* state-consistent *relative to $\mathcal{G}$ (or, simply,* renegotiation-proof*) if, after any history leading up to any state $(\tilde{w}, \tilde{y})$ and continuation contract $\tilde{C}$,*

$$\Pi(\tilde{C}) \geq \Pi(\mathcal{G}_{(\tilde{w},\tilde{y}),(w,y)}(C)),$$

and, reciprocally,

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

This definition requires not only that continuation payoffs must sustain the comparison with transformations of the initial contract, but also the reverse: the initial contract must sustain the comparison with transformations of the continuation contracts consistent with the initial state.

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 transformation 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, 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

In this section, state consistency is exploited, first, to characterize the principal's payoff as a function of the state, for any renegotiation-proof contract. This payoff characterization is then used to derive properties of contractual variables, for any such contract.

Focusing on renegotiation-proof equilibria can simplify the analysis of environments with persistent private information. Intuitively, the possibility of renegotiating equilibria prevents complicated forms of history dependence on the equilibrium path. For example, punishment phases become unavailable if it is commonly known that such phases will be renegotiated upon. This observation, commonly made in the context of repeated games, takes a particular strength for stochastic games with incomplete information, because it destroys the credibility of "threat keeping constraints" or "marginal promised utility."[27]

The relation between payoffs and contractual variables is more subtle than it would appear *a priori*. In particular, while weak renegotiation proofness trivially implies that the principal's continuation payoff can only depend on the current states (agent's promised utility and cash flow), it does not immediately imply that the contractual variables themselves (consumption, sensitivity, effort) are also Markovian. Indeed, there could *a priori* exist multiple contracts that yield the same Markovian payoff to the principal, and the principal may credibly exploit these various continuation contracts to affect the agent's ex ante incentives. Section 4.2 shows how to address this concern, by exploiting the payoff formula derived in Section 4.1.

4.1 Payoffs

As observed in the previous section, the principal's continuation payoff for any internally consistent (and, *a fortiori*, state consistent) contract only depends, at any time, on the current state (w, y).

[27] See Fernandes and Phelan (2000) and Williams (2011).

Let $\Pi(w, y)$ denote the principal's payoff.

4.1.1 Comparing Contracts Across Cash-Flows Levels

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

$$V(w, y, x, L, A) = E\left[\int_0^\infty e^{-rt}\left(u(C_t - G_t) - \phi(A_t)\right) dt\right] \tag{4}$$

subject to (1), (2),

$$dW_t = (rW_t - u(C_t) + \phi(\bar{A}_t))dt + S_t(L_t dt + A_t dt + \lambda G_t + \sigma dB_t),$$

$$dG_t = L_t dt$$

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

The principal's expected payoff is

$$\Pi(w, y) = E_{w,y}\left[\int_0^\infty e^{-rt}(Y_t - C_t)dt\right].$$

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 $\hat{C}_t = \mathcal{C}(\tilde{Y}_s : s \leq t)$, where $\tilde{Y}$ is constructed as follows: $\tilde{Y}_0 = y$, 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 $\tilde{Y}$.

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

Proposition 3 *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 $(w, \hat{y})$-*transform* of C and denoted $\mathcal{G}_{(w,\hat{y}),(w,y)}(C)$.

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 (26). 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,

$$\Upsilon(w, y) = \frac{y}{r + \lambda} + \frac{\xi}{r(r + \lambda)} + \int_0^\infty e^{-rt} \frac{\alpha_t}{r + \lambda} dt, \tag{5}$$

where

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

4.1.2 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

$$\hat{C}_t = C_t - \frac{\log(\beta)}{\theta}$$

$$\hat{A}_t = A_t + \frac{\log(\beta)}{\chi}$$

Proposition 4 *(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 that[28]

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

The previous analysis yields the following result.

Proposition 5 *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:

$$\Pi(w', y') = \frac{(y' - y)}{r + \lambda} + \frac{\log(w'/w)}{r} \left(\frac{1}{\theta} + \frac{1}{\chi(r+\lambda)} \right) + \Pi(w, y). \tag{6}$$

4.1.3 Transformation Group and Payoffs

Combined together, Sections 4.1.1 and 4.1.2 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. Therefore, state consistency implies the following relationships for the principal's payoff function, as a consequence of Proposition 1.[29] The relation is expressed with respect to some reference state, chosen to be $(-1, 0)$.

Proposition 6 *The group $\mathcal{G}$ constructed in Sections 4.1.1 and 4.1.2 preserves incentive compatibility and is monotone. For any state-consistent contract, and any time t,*

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

[28] The effort taken at any time has a decaying effect on future cash flows, with decaying rate λ, discounted at rate r. This explains the factor $(r + \lambda)$ in the denominator.

[29] As mentioned in Section 4.1, Proposition 6 has been established without assuming Markov contractual variables, and is thus independent from the analysis of Section 4.1.

4.2 Contractual Variables

To pin down the contractual variables c, a, and s, 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.

4.2.1 Contractual Variables are Markov

First, one can show that the *contractual variables* themselves ar Markov: they only depend on (w, y). 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 equation[32]

$$0 = \sup_{(c,s)\in\mathcal{K}(w,y)} \Big\{ y - c - r\Pi(w,y) + \Pi_w(w,y)\,(rw - u(c) + \phi(a(s))) + \Pi_y((\xi - \lambda y) + a(s)) \\ + \frac{1}{2}\Pi_{ww}s^2\sigma^2 + \Pi_{wy}s\sigma + \frac{1}{2}\Pi_{yy}\sigma^2 \Big\}, \quad (7)$$

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

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

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

[32] Bergemann 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.

[33] The 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 on reported cash flows and rewards, but this impact is identical to the impact of a lie, and must vanish for any truthful contract, as illustrated by Equation 16 for the Markovian case.

satisfied if ϕ is exponential and Proposition 6 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.

4.2.2 Functional Form of Contractual Variables

Consider some state-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.

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.

$$0 = \sup_{(c,s) \in \mathcal{I}(w,y)} \Big\{ y - c - r\Pi(w, y) + \Pi_w(w, y)\,(rw - u(c) + \phi(a(s))) + \Pi_y((\xi - \lambda y) + a(s)) + \frac{1}{2}\Pi_{ww} s^2 \sigma^2 + \Pi_{wy} s\sigma + \frac{1}{2}\Pi_{yy}\sigma^2 \Big\}, \quad (8)$$

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 6 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 $\mathcal{I}(\tilde{w}, \tilde{y})$, and vice versa. Given the functional form of Π, this implies that $(\bar{c}, \bar{s})$ solves (8) 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 7 *Any state-consistent contract has the form*

$$C_t = c_1 - \frac{\log(-W_t)}{\theta}$$

$$A_t = a_1 + \frac{\log(-W_t)}{\chi}$$

$$S_t = -W_t\bar{s}$$

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 problem reduces to an optimization with respect to variables c_1, a_1, and $\bar{s}$. Incentive compatibility constraints imposes a relationship between these three variables, which we express in terms of utility and cost of effort. Let $u_1 = u(c_1)$ and $\phi_1 = \phi(a_1)$. The following result is shown in Section 6.

Proposition 8 *A renegotiation-proof contract is incentive compatible if and only if*

$$\phi_1(\bar{s}) = \frac{\bar{s}}{\chi}, \tag{9}$$

and

$$u_1(\bar{s}) = -\frac{\bar{s}(\chi\lambda + \bar{s})}{\chi(\theta - \bar{s})}, \tag{10}$$

It remains for the principal to optimize his payoff with respect to $\bar{s}$.

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 payoff*[34]

$$\frac{1}{\theta r}\left[\kappa\log(\phi_1(\bar{s})) + \log(-u_1(\bar{s})) + \frac{\kappa+1}{r}\left(r + u_1(\bar{s}) - \phi_1(\bar{s}) - \frac{1}{2}\sigma^2\bar{s}^2\right) + A(\kappa, w, \bar{\phi}, y, \xi)\right], \tag{11}$$

where

$$A(\kappa, w, \bar{\phi}, y, \xi) = (\kappa+1)\log(-w) - \kappa\log(\bar{\phi}) + (\kappa\chi)(ry + \xi).$$

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

5 Comparative Statics

5.1 Persistence

The first result is to show that the sensitivity of the contract, $\bar{s}$ is increasing in the persistence of the agent's type.

PROPOSITION 10 *The optimal sensitivity $\bar{s}$ is decreasing in λ.*

One possible intuition is that with more persistence, current report have more bearings about future cash flows. When we shut down the moral hazard component of the model (setting A_t identically equal to zero or, equivalently, letting χ go to infinity), the proof is straightforward. The objective of the principal in that case is to maximize, removing terms and factors that do not influence the optimal $\bar{s}$,

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

where

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

The function u_1 is submodular in $(\bar{s}, \lambda)$, as is easily checked by computing the cross derivative. Therefore, the second term is submodular. The logarithmic term breaks up into separate functions of $\bar{s}$ and λ, and is therefore modular in $(\bar{s}, \lambda)$. Submodularity of the objective implies that $\bar{s}$ is decreasing in λ (see Topkis 1978). The intuition here can be back traced to the incentive compatibility condition (18), specialized to the case where $\phi_1 = 0$:

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

A lower λ (higher persistence) implies a higher sensitivity $\bar{s}$.

The intuition and proof for the general case with moral hazard is more involved.[35]

To isolate the impact of persistence, one may want to change ξ so as to keep the cash flows's long term average, ξ/λ, constant. However, changes ξ has no impact on the optimal level of $\bar{s}$, as may be seen from 11.

[35] Formally, the objective function is of the form $f(s,\lambda) + (\kappa(\lambda)+1)g(s,\lambda)$, where f and g are submodular in (s,λ), f is increasing in s, g is decreasing in s, and κ is decreasing in λ. Increasing λ yields a lower marginal value of s for f and g. However, it also decreases the weight of g relative to f which makes higher values of s more attractive, since f is increasing in s. Showing that the overall effect is to decreasing the optimal value of s is left for the appendix.

5.2 Immiserization

From Proposition 9, one may easily show that the drift of promised utility is negative, for the case of pure reporting (effort is set to 0, shutting down the moral hazard component). In that case, Equation (11) reduces to

$$\max_{\bar{s}} \log(-u_1(\bar{s})) + \frac{1}{r}\left(r + u_1(\bar{s}) - \frac{1}{2}\sigma^2\bar{s}^2\right) + (\kappa+1)\log(-w), \tag{12}$$

or equivalently,

$$\max_{u_1} \log(-u_1) + \frac{1}{r}\left(r + u_1 - \frac{1}{2}\sigma^2\bar{s}(u_1)^2\right) + (\kappa+1)\log(-w), \tag{13}$$

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 (13) would vanish, and u_1 would, at the optimum, be equal $-r$, implying that the drift of W_t 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 by captured by the term

$$v_y = E\left[\int_0^\infty e^{-(r+\lambda)t}u'(C_t)dt\right]. \tag{14}$$

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 (13) 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 (and does not vanish), since W_t is negative.

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

Proposition 11 (Immiserization) *For all parameters $(r, \lambda, \xi, \theta, \chi, \bar{\phi})$, the continuation utility W_t of the agent at the optimal renegotiation-proof contract has a strictly negative drift.*

5.3 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.[36] 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 truthtelling, the other channel through which the agent can deviate.

To keep inducing truthtelling, 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 (14)), 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.[37]

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

Proposition 12 *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 (11), 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 12 is the curvature of the cost function rather than its scale: a flatter cost functions

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

[37] As χ goes to zero, the immiserization effect becomes muted, as shown at the end of the proof of Proposition 12.

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 $\bar{\phi}$).

5.4 Impact of Noise and Discount Rate

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 13 *The optimal sensitivity $\bar{s}$ is decreasing in σ.*

Proof. The objective 11 is submodular in σ and $\bar{s}$. The result follows from Topkis (1978).

Perhaps more surprisingly, the sensitivity is decreasing in the patience of the players, as shown in the Appendix.

PROPOSITION 14 *The optimal sensitivity $\bar{s}$ is increasing in r.*

One possible intuition for this result is that reports have more importance, other things equal if the players care more about the future: because the persistence of the agent's type is perceived as more important in that case.

6 Necessary and Sufficient Conditions for Incentive Compatibility

Beyond its obvious importance for the present analysis, the proof for verifying incentive compatibility of renegotiation-proof contracts is interesting for several reasons. First, it illustrates a strategy, potentially useful in other problems, to deal with an unbounded reporting domain by expending the strategy space of the agent. Second, this extension provides a way to think about private savings (see Section 9). Third, it clarifies why Ponzi schemes (of two kinds: always overreporting or always underreporting) are not sustainable in the present setting.

For the renegotiation-proof contracts constructed in Section 4, the promised utility of the agent evolves as

$$dW_t = (r + u_1 - \phi_1)W_t dt + s(W_t)(dG_t + A_t dt + \lambda G_t dt + \sigma dB_t), \tag{15}$$

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 problem is therefore reduced to

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

subject to $dG_t = L_t dt$, $G_0 = g$, and (15). The HJB equation for this problem is

$$0 = \sup_{l,a} \{ u(c(w) - g) - \phi(a) - rv(w,g) \\ + v_w \left(rw - u(c(w)) + \phi(a(w)) + s(w)(a - a(w) + l + \lambda g) \right) + v_g l + \frac{1}{2}(s(w))^2 \sigma^2 v_{ww} \Big\}, \quad (16)$$

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

It is easy to show by applying the same controls starting from different values of w that the value function has the form $v(w,g) = wf(g)$ for some function f to determine. Incentive compatibility obtains 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 condition[38] with respect to a yields $\phi(a) = \frac{f(g)s(w)}{\chi}$, so the Bellman equation becomes, after simplification and dividing throughout by $(-w)$,

$$0 = \sup_{\ell} \left\{ u_1 \exp(\theta g) + f(g)(-u_1 + \bar{s}(\ell + \frac{1}{\chi}(\log f(g)) + \lambda g)) - f'(g)\ell \right\}. \quad (17)$$

A priori, the function f defining the value function need not be everywhere differentiable, but it is a viscosity solution of (17). Moreover, $v(w,g)$ is clearly decreasing in the gap g, which implies that f is increasing. This implies that f is left and right differentiable at 0. Incentive compatibility implies that, for $g = 0$, it is optimal for the agent not to increase the gap above zero, so that

$$\bar{s} - f_r(0) \leq 0$$

and it is optimal not to decrease it below zero, so that

$$\bar{s} - f_l(0) \geq 0$$

Therefore, incentive compatibility implies that $f_l(0) \leq \bar{s} \leq f_r(0)$. Now suppose that one of these inequalities is strict, for example, $\bar{s} < f_r(0)$. This implies that there exists a right neighborhood of 0, such that for $g \in (0, \eta)$ it is strictly optimal for the agent to reduce the gap to zero, at an infinitely negative rate. 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 $G_{t+\varepsilon} = G_t - K\varepsilon$. Therefore, the change in the value function must satisfy, for $g \in [0, \eta)$, $f(g) = \exp(\bar{s}g)f(0)$. Since $f(0) = 1$ for an incentive compatible contract, this implies that

$$f_r(0) = \bar{s},$$

[38]Since the objective is strictly concave in a and the domain of a is open, the first-order condition pins down the unique optimum.

which contradicts our assumption that $\bar{s} < f_r(0)$. Thus, necessarily, $f_r(0) = \bar{s}$. By a similar argument, $f_l(0) = \bar{s}$. In conclusion, f is differentiable at 0, and thus the value function $v(w, x, y)$ of the agent is differentiable with respect to x and to y whenever $x = y$. In the Appendix (Lemma 1), this differentiability is used in an envelope argument to show that, necessarily,

$$\bar{s} = f'(0) = v_x(w, x, x) = -v_y(w, x, x) = \int_0^\infty e^{-(r+\lambda)t} u'(C_t + X_t - Y_t) dt.$$

Simple computations, also in the appendix (Lemma 2), imply that

$$\bar{s} = \frac{\theta(-u_1)}{\lambda - u_1 + \phi_1} \in (0, \theta). \tag{18}$$

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 compatibility rules this case out.

Therefore, we have established that (10) and (9) are necessary for incentive compatibility. From now one, we assume that these relations hold, and establish sufficiency.

The objective (17) 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

$$W_{t_+}(\Delta L) = \exp(-\bar{s}\Delta L) W_t. \tag{19}$$

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 (19).

Report jumps amount to impulse controls on the part of the agent (see for example, Øksendal and Sulem (2004)). The HJB equation (17) becomes

$$0 = \max\left\{\sup_{\ell\in\mathbb{R}}\left\{u_1\exp(\theta g) + f(g)(-u_1 + \bar{s}(\ell + \frac{1}{\chi}(\log f(g)) + \lambda g)) - f'(g)\ell\right\},\right.$$
$$\left.\sup_{\Delta L\in\mathbb{R}\setminus\{0\}}\{\exp(-\bar{s}\Delta L)f(g+\Delta L) - f(g)\}\right\}. \quad (20)$$

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

$$\sup_{\ell}\left\{u_1\exp(\theta g) + \exp(\bar{s}g)\left(-u_1 + \bar{s}\left(\ell + \frac{1}{\chi}(\bar{s}g) + \lambda g\right)\right) - \bar{s}\exp(\bar{s}g)\ell\right\} \leq 0 \quad (21)$$

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 function implies that, for all $g \neq 0$,

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

Since $u_1 < 0$, (21) 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 (10), 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 5.3 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.[39]

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 the agent's type through

[39] Even 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.

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.[40] 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.[41]

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, 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.[42] 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.

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).

[40] That 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.

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

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

Second, comparing across cash-flow levels, the argument of Section 4.1.1 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 (5).

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 4.1.2.

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).[43]

PROPOSITION 15 *For any states (w, y),(w', y') such that $w' \geq w$,*

$$\Pi(w', y') \geq \frac{(y' - y)}{r + \lambda} + \frac{\log(w'/w)}{r} \left(\frac{1}{\theta} + \frac{1}{\chi(r + \lambda)} \right) + \Pi(w, y). \tag{22}$$

This is intuitive: the farther away one gets from the constraint $\underline{w}$, and the more flexibility one has to choose consumption/effort processes achieving the given promised utility and, therefore, the

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

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 (11)), 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.[44] With stochastic games, some underlying state affects the physical environment of the players. When it is possible to construct an equilibrium starting from a given state from another equilibrium starting from a different state, such construction can be used to narrow down the set of renegotiation-proof equilibria, by comparing continuation equilibria at a given state with transformations to that state of continuation equilibria starting from other states.

This procedure can be justified in several ways.

First, the equilibria obtained through such constructions are necessarily challengers in some mainstream notion of strong renegotiation-proofness. For example, the concept of strongly renegotiation proof equilibria (Farrell and Maskin (1989) imposes comparisons between any continuation of the equilibrium under consideration with the continuations of *all* weakly renegotiation proof equilibria, and those necessarily include the constructions previously mentioned.

Second, these comparisons also give a particular meaning, for stochastic games, to the concept of "social norms" studied by Asheim (1991): the comparisons across states are a natural candidate for generating the set of equilibria that constitute a social norm, across states.

Finally, there is a cognitive interpretation of these comparisons: players may be able to recognize that there are clear relations between the sets of feasible continuation equilibria across different states, in the same way that they recognize, in a repeated game, whether continuation payoffs across different histories are Pareto ranked.

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

More generally, the set of renegotiation-proof equilibria becomes smaller as the set of comparisons becomes larger. In the example studied here, as we add layers of challengers, we impose more properties on the set of renegotiation proof contracts, each layer implying particular properties of the renegotiation-proof contracts: comparison across cash flow levels implies that consumption and other contractual variables are independent of the current cash flows level. Similarly, comparison across promised utilities, combined with the exponential utility and cost specifications, implies that a translation property of the consumption process with respect to the promised utility.

It would be useful to compare the optimal state-consistent contract with the optimal contract with full commitment. For the case of pure reporting, the optimal full commitment contract has been studied by Williams (2011), who provides a general analysis of necessary and sufficient conditions for optimality. For the case of exponential utility, Williams shows that if the optimal contract is Markovian (with the states being the promised utility and the current cash flow) then it creates no immiserization. Combined with the results of the present paper, which imply that the optimal Markovian contract causes immiserization, Williams's observation implies that the optimal full-commitment contract must also depend on the marginal promised utility.[45]

It would also be interesting to characterize the optimal weakly renegotiation-proof contract. Here again, the solution is unknown. As pointed out in Section 4.1, the payoff of the principal is Markovian in that case. However, this alone does not guarantee that the contractual variables themselves are Markovian, since the principal could in principle benefit from mixing across different contracts that provide him with the same payoff. Even when contractual variables are Markovian, it is unclear whether the sensitivity function $(w, y) \mapsto s(w, y)$ takes the form $s(w, y) = \bar{s}(-w)$ for some constant $\bar{s}$, as in the contracts that are studied in this paper. While the analysis of optimal contracts for such settings requires a separate study, one may conjecture that some of the intuition provided in the present paper, concerning the impact of persistence and discounting, in particular, is likely to carry over to such environments.

9.2 Robustness and Extensions

Choice of a Production Technology One consequence of Section 5.3 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 13 and Equation (11) imply that the principal's payoff is decreasing with the noise in the production technology, and that increased noise reduces the

[45] The contract derived by Williams, based on a numerical observation, belongs to the class of renegotiation-proof contracts studied here, with $u_1 = -r$, and is strictly dominated by the optimal renegotiation-proof contract.

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 per se an enlarged reporting space to the agent, since he is already able to make arbitrary reports. Furthermore, using the extended reporting space and contracts described in Section 6, where the agent can make jumps in his reports and is rewarded or punished at rate $\bar{s}$, the agent can replicate whatever change of consumption he 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 cash flow 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 7. 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 (18) 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 3

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,

$$E\left[\int_0^\infty e^{-rt}\left(u\left(\mathcal{C}(Y_s : s \le t) - \int_0^t L_s ds\right) - \phi(A_t)\right) dt\right] \geq E\left[\int_0^\infty e^{-rt}\left(u\left(\mathcal{C}(Y'_s : s \le t) - \int_0^t L'_s ds\right) - \phi(A'_t)\right) dt\right] \quad (23)$$

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 follow the equations

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

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

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 (23), 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. ■

Necessary Conditions for Incentive Compatibility

LEMMA 1 *For any renegotiation-proof contract* $(c, \bar{a}, s)$,

$$v_y(w, x, x) = -v_x(w, x, x) = -\int_0^\infty e^{-(r+\lambda)t} u'(C_t + X_t - Y_t) dt. \tag{24}$$

Proof of Lemma 1

The value function of the agent satisfies the optimization problem

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

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 cash flows 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

$$dY_t = [(\xi - \lambda Y_t) + A_t + \bar{L}_t]dt + \sigma dB_t. \tag{25}$$

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

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

subject to $Y_0 = y$, $X_0 = x$, (25), and

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

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

If the contract is incentive compatible, it is optimal to set $\bar{L}_t = 0$ whenever initial conditions are such that $y = x$. Section 6 established that the value function is differentiable with respect to x and y whenever $x = y$. An application of the Envelope Theorem (Milgrom and Segal, 2002, Theorem 1) then 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, and

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

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

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

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

$$X_t = e^{-\lambda t} x + \int_0^t e^{\lambda(s-t)}(A_s + \xi) ds + \int_0^t e^{\lambda(s-t)} \sigma dB_s \tag{26}$$

This implies that $\frac{dX_t}{dx} = e^{-\lambda t}$ and yields the formula of Lemma 1. ■

LEMMA 2 *The following equality holds for any truthful contract:*

$$E \int_0^\infty e^{-(r+\lambda)t} u'(C_t) dt = (-w) \frac{(-u_1)\theta}{\lambda - u_1 + \phi_1}.$$

Since $u'(C_t) = -\theta u(C_t)$,

$$E \int_0^\infty e^{-(r+\lambda)t} u'(C_t) dt = E \int_0^\infty e^{-(r+\lambda)t} \theta u(C_t) dt = \theta u_1 \int_0^\infty e^{-(r+\lambda)t} E[W_t] dt.$$

Moreover,

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

With exponential utility and cost functions, $u(C_t) = -W_t u_1$ and $\phi(A_t) = -W_t \phi_1$. Letting $\vartheta(t) = E[W_t]$ ($\vartheta(0) = w$), this implies that

$$\frac{d\vartheta}{dt}(t) = (r + u_1 - \phi_1)\vartheta(t),$$

and hence that

$$E[W_t] = e^{(r+u_1-\phi_1)t} w. \tag{27}$$

Integrating this expression over time yields the result.

Proof of Proposition 9

The principal's objective is to maximize $\Pi(w, x)$ with respect to u_1 and ϕ_1. From (5), and neglecting for now terms that are not affected by the contract choice, this is equivalent to solving

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

where

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

and

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

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

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

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

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

Replacing a_1, c_1 and $E\log(-W_t)$ by their formulas in terms of ϕ_1, u_1, t, and the parameters of the model, and integrating the last term proves the proposition. ■

Proof of Proposition 10

Dividing (11) by $\kappa + 1$ (which does not affect the optimal value of $\bar{s}$) and ignoring terms that are independent of $\bar{s}$, it suffices to show that the function

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

is submodular in (s, λ). The cross-derivative with respect to $\bar{s}$ and λ is equal to

$$\frac{\theta\chi}{(\chi(r+\lambda)+\theta)^2}\left[\frac{1}{\chi\lambda + \bar{s}} + \frac{1}{\theta - s}\right] - \frac{\chi}{(\chi\lambda+\bar{s})^2}\frac{\chi(r+\lambda)}{\theta + \chi(r+\lambda)} - \frac{1}{r}\frac{\theta}{(\theta - s)^2}.$$

Multiplying this expression throughout by the denominators, we need to show that, for all nonnegative parameters, $\theta, \lambda, r, s, \chi$ such that $s < \theta$ the expression

$$A = r(\theta - \bar{s})^2\chi^2(r+\lambda)(\theta + \chi(r+\lambda)) + \theta(\chi(r+\lambda)+\theta)^2(\chi\lambda+\bar{s})^2 - r(\theta + \chi\lambda)\theta\chi(\chi\lambda + \bar{s})(\theta - \bar{s})$$

is nonnegative. Developing this expression with Maple yields

$$A = r^3\chi^3\theta^2 + r^2\chi^2\theta^3 + 2\theta\chi^4 r\lambda^3 + 2\theta\chi^3\lambda^3\bar{s} + \theta\chi^2\lambda^2\bar{s}^2 + 2r^2\chi^3\theta^2\lambda + 2r\chi^3\theta^2\lambda^2 + 2r^2\chi^2\bar{s}^2\theta + 2r^2\chi^3\bar{s}^2\lambda + \theta\chi^4\lambda^4 + \theta^3\chi^2\lambda^2 + r^3\chi^3\bar{s}^2 + \theta^3\bar{s}^2 + 2\theta^2\chi^3\lambda^3 + 2r\chi^2\theta^2\bar{s}\lambda + 3r\chi^3\theta\bar{s}\lambda^2 + 4r\chi^2\bar{s}^2\lambda\theta + r\chi^3\bar{s}^2\lambda^2 + 2\theta^3\chi\lambda\bar{s} + 3\theta^2\chi r\bar{s}^2 + 4\theta^2\chi^2\lambda^2\bar{s} + 2\theta^2\chi\lambda\bar{s}^2 + \theta\chi^4 r^2\lambda^2 - \theta^3\chi r\bar{s} - 2r^2\chi^2\theta^2\bar{s} - 2r^3\chi^3\theta\bar{s} - 2r^2\chi^3\theta\bar{s}\lambda. \quad (28)$$

Therefore, it suffices to show that the last four terms, which are negative, are dominated by the positive terms. The last negative term is dominated by $2r^2\chi^3\bar{s}^2\lambda + 2r^2\chi^3\theta^2\lambda$: the difference is equal to $2r^2\chi^3\lambda(\theta^2 + \bar{s}^2 - \theta\bar{s})$, which is positive. We now focus on all the terms not containing λ. Let, for the rest of this proof only, $x = \bar{s}/(r\chi)$ and $y = \theta/(r\chi)$, so that $0 \le x \le y$. The sum of all the terms that do not contain λ can be reexpressed, after division by $(r\chi)^5$, as

$$3x^2y^2 + y^3x^2 + x^2 + 2x^2y + y^2 + y^3 - xy^3 - 2y^2x - 2xy, \tag{29}$$

or

$$y^3(1-x)^2 + xy(1-y)^2 + (y-x)^2 + xy(3yx + 2x - 1).$$

The last (and only problematic) term is positive if $3x^2 + 2x - 1 \ge 0$ or, equivalently, if $x \ge 1/3$. It is also positive if $y \ge 1/(3x)$. There remains to examine the case in which $x < 1/3$ and $y < 1$. We can reexpress (29) as

$$y^3(1-x)^2 + (y-x)(y-x-2xy) + xy(y^2 + 3xy).$$

The second term and only problematic term, call it B, is positive if $y - 2xy > x$. Since $x < 1/3$, we have $y - 2xy > \frac{1}{3}y$. Therefore, B is positive if $y \ge 3x$. Suppose instead that $y < 3x$. We can rewrite (29) as

$$(x-y)^2 + y(x-y)^2 + x^2y + y^3x^2 + 3x^2y^2 - xy^3. \tag{30}$$

Since $y < 1$ and $y < 3x$, we have $xy^3 < xy^2 < 3yx^2$. The expression $3x^2y^2 + x^2y - xy^3$ is therefore positive if $3x^2y^2 \ge 2x^2y$, or if $y \ge 2/3$. Suppose instead that $y < 2/3$. In that case, we have $xy^3 < \frac{2}{3}xy^2 < 2yx^2$. The expression $3x^2y^2 + x^2y - xy^3$ is thus positive if $3x^2y^2 \ge x^2y$, or if $y \ge 1/3$. This leads us to the final case where $y < 1/3$. We now have $xy^3 < \frac{1}{3}xy^2 < yx^2$. The expression $x^2y - xy^3$ is therefore positive. This shows that, for all parameter values, the modified objective function is submodular in $(\bar{s}, \lambda)$, and concludes the proof of the proposition. ■

Proof of Proposition 11

Let $x = \phi_1 - u_1$. It suffices to show that, for $\bar{s}$ maximizing (11), $x < r$. 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

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

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 $x < r$.[46] ■

Proof of Proposition 12

Recall from Proposition 9 that

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

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 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

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

This shows that $\Pi(\chi)$ diverges to $+\infty$ if $w < -1/r$ and to $-\infty$ if $w > -1/r$.[47]

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

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

[46] A simpler proof is to observe that i) $x(\bar{s}) = r$ when $\sigma = 0$ (since perfect consumption smoothing is optimal), ii) $x(\bar{s}) \propto \bar{s}/(\theta - \bar{s})$ is an increasing function of $\bar{s}$ (from (9) and (10)), and iii) $\bar{s}$ is decreasing in σ, by submodularity of the objective function (see Proposition 13).

[47] That 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.

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

Proof of Proposition 14

The optimal value of $\bar{s}$ is unchanged if one multiplies (11) by r and gets rid of the term r in the factor $(r + u_1 - \phi_1 - 1/2\sigma^2\bar{s}^2)$. The resulting objective function is equal to

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

The function $r \mapsto \kappa(r)r$ is increasing in r, while $\phi_1(\bar{s})$ and $u_1(-\bar{s})$ are increasing in $\bar{s}$. Therefore, the first two terms of the objective are supermodular. The last term is also supermodular, because $\kappa(r)+1$ is decreasing in r, while $u_1 - \phi_1 - \frac{1}{2}\sigma^2\bar{s}^2$ is decreasing in $\bar{s}$. The result then obtains from an application of the comparative statics theorem for supermodular functions of Topkis (1978). ■

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.

Bester, H., Strausz, R. (2001) "Contracting with Imperfect Commitment and the Revelation Principle: The Single Agent Case," *Econometrica*, Vol. 69, pp. 1077–1098.

DeMarzo, P., Sannikov, Y. (2011) "Learning, Termination, and Payout Policy in Dynamic Incentive Contracts," Working Paper, Princeton University and Stanford University.

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.

FARHI, E., WERNING, I. (2010) “Insurance and Taxation over the Life Cycle,” Working Paper, Harvard University and MIT.

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, Northwerstern University.

GARRETT, D., PAVAN, A. (2010b) “Managerial Compensation in a Changing World,” Working Paper, Northwerstern 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.

KRANZ, S., OHLENDORF, S. (2010) *Renegotiation-Proof Relational Contracts with Side Payments*, Working Paper, Mannheim University.

KYDLAND,F., PRESCOTT, E. (1977) "Rules Rather than Discretion: The Inconsistency of Optimal Plans," *Journal of Political Economy*, Vol. 85, No. 3, pp. 473–492.

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.

MILLER, D., WATSON, J. (2011) "A Theory of Disagreement in Repeated Games with Bargaining," Working Paper, UC San Diego.

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 PrincipalAgent Problem," *Review of Economic Studies*, Vol. 75, No. 3, pp. 957-984.

SANNIKOV, Y. (2011) "Moral Hazard and Long-Run Incentives," Working Paper, Princeton University.

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.