Renegotiation-Proof Contracts with Persistent States*

Bruno Strulovici

May 31, 2022

Abstract

This paper introduces a new approach to model renegotiation in contractual relationships and applies it to study how renegotiation shapes long-term contracts in principal-agent relationships with persistent states. The structure of players' payoff and state dynamics generates an algebraic structure over contractual equilibria and determines the set of alternatives considered through renegotiation. Using recent advances in functional stochastic differential equations, the paper derives an Observability Theorem and a Revelation Principle to address asymmetric information over persistent variables in diffusion models. Truthful renegotiation-proof contracts are characterized by a single number—their sensitivity to the agent's report—and are self-correcting off the equilibrium path. The sensitivity of the optimal contract is increasing in information persistence and *decreasing* in players' patience.

1 Introduction

Financial bailouts, fiscal policies, and corporate pay cuts during recessions are instances in which regulators, firms, and other economic agents have the possibility to *renegotiate explicit or implicit contracts* in reaction to shocks in their environment. In these instances, moreover, state variables such as agents' revenue or productivity are often privately observed and correlated over time.

Despite its relevance for economics, renegotiation has proved challenging to model and analyze. Even for repeated games, there is no universally accepted concept of "renegotiation-proof" equilibrium. One concept, *internal consistency*, is often perceived as a minimal requirement, and defined as follows: an equilibrium is internally consistent if there do not exist two histories such that the continuation equilibrium following the first history is Pareto dominated by a continuation equilibrium following the second history. Presumably, players facing the first continuation equilibrium

*Northwestern University. This paper subsumes the 2011 working paper entitled "Renegotiation-Proof Contracts with Moral Hazard and Persistent Private Information." I am particularly grateful to Alex Bloedel, Hector Chade, Roozbeh Hosseini, R. Vijay Krishna, Nikolai Krylov, Kyohei Okumura, Alessandro Pavan, as well as to Bruno Biais, Mehmet Ekmekci, Daniel Garrett, Brett Green, Christian Hellwig, Johannes Hörner, Francisco Ruiz-Aliseda, Larry Samuelson, Ilya Segal, Itai Sher, Juuso Toikka, and various seminar audiences for questions and comments. Part of this research was conducted while I was visiting Kyoto University, who hospitality is gratefully acknowledged.

could think of moving to the second, Pareto superior continuation equilibrium, which would destroy the initial equilibrium.[1] Internal consistency is weak because it does not impose any comparison with any other equilibrium. For example, indefinitely playing a bad Nash equilibrium is internally consistent.

At the opposite extreme, *strongly renegotiation-proof* equilibria (Farrell and Maskin (1989)) must sustain comparisons with all internally consistent equilibria and may fail to exist. Other concepts involve a fixed point problem determining the set of equilibria that are renegotiation proof (see Asheim (1991)) and are subject to existence and multiplicity issues.

State Consistency: These problems would seem even more severe in environments, pervasive in economics, with persistent shocks. In these environments, however, a strengthening of internal consistency can prove powerful enough to drastically reduce the set of renegotiation-proof equilibria and yield sharp 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, thus creating a concept of *state consistency.*

For example, if a state variable describes the current productivity of a firm, one cannot directly compare continuation equilibria starting from different productivity levels (or "states") since the production frontier and, more generally, the physical environment is different across states. However, one may be able to *transform* a continuation equilibrium starting from a given state into another continuation equilibrium starting from a different state, for example by using some homotheticity argument. One may then compare the continuation payoffs of various transformations that generate equilibria starting from the same state, and use Pareto dominance to determine whether an equilibrium is vulnerable to renegotiation. State consistency is weaker than strongly renegotiation-proofness because it does not require that players envision radically different equilibria, but only that they think by analogy with the equilibrium that they are playing.

State consistency imposes an algebraic structure on renegotiation-proof equilibria. By transforming an equilibrium starting from one state into equilibria starting from other states, one generates an *orbit* within the set of equilibria, in a group-theoretic sense that the paper makes precise. This orbit has a structure that reflects the nature of the transformation. State consistency imposes a comparison between continuations of an equilibrium and the orbit that it generates. As a result, a state-consistent equilibrium inherits many properties of its orbit. For example, in the simple model of Section 2, the structure of the principal's payoff inherits many properties from the structure of the

[1] Internal consistency is due to Bernheim and Ray (1989) and closely related to Farrell and Maskin's (1989) weakly renegotiation-proof equilibria. Although internal consistency is a concept mild, Abreu and Pearce (1991) and Asheim (1991) point out that it does assume time invariance.

output dynamics and of the agent's utility and effort cost functions, and the class of renegotiation-proof contracts is reduced to a one-parameter family.

Private Information and Renegotiation: An additional complication arises if some persistent state variables are observed only by some party. Even if parties can communicate, they may misreport private information, and parties could have to renegotiate under asymmetric information.

One approach is to restrict attention to equilibria in which parties truthfully reveal their information. In a truthful equilibrium, all parties know the state variables at all times and state consistency can be used to study truthful, renegotiation-proof equilibria.

With renegotiation, however, focusing on truthful equilibria may *a priori* be restrictive since the absence of commitment precludes the use of the standard Revelation Principle. However, renegotiation also incites parties to try and erase any inefficiency in ongoing agreements, which requires some exchange of information. When parties can communicate frequently, parties can renegotiate away any inefficiency stemming from information asymmetries and cannot commit not to do so. If efficiency requires full information disclosure, focusing on truthful equilibrium may thus be without loss of generality when parties can continually communicate.[2]

Observability and Revelation Principle in a Diffusion Model: This paper takes a different approach to address private information. It studies a continuous-time model, in which communication continually takes place. In the model, which is due to Williams (2011), an agent privately observes a stochastic cash flow process, which he reports to a principal. The agent can lie: if the true cash flow process is $\{X_t\}_{t\geq 0}$, the agent reports a cash flow process $\{Y_t\}_{t\geq 0}$ where Y_t is a function of $\{X_s\}_{s\leq t}$ and of the agent's reporting strategy L. Mathematically, the lying process L is *adapted*[3] to the filtration generated by X and is assumed to enter the reporting process Y through its drift: $dY_t = dX_t + L_t dt$ until Section 9, which considers reporting jumps.

The agent's equilibrium strategy, L, is known to the principal. Therefore, given a reporting history

[2]The main theorem of Strulovici (2017) formalizes this intuition: it shows in a principal-agent model with asymmetric information that when (i) allocation efficiency requires full information disclosure and (ii) parties can communicate with arbitrarily high frequency before the allocation takes place, information is fully revealed and the transaction is efficient. Maestri (2017) shows a similar result when parties transact in each period: information is revealed arbitrarily quickly relative to the parties' discount rate as parties become arbitrarily patient. These papers consider an explicit game of renegotiation—with offer and acceptance decisions—whereas renegotiation in the present paper is modeled using a set-theoretic approach.

[3]The requirement that the agent's strategy is adapted to what the agent observes is standard in continuous-time models, including in Williams (2011). This assumption often plays a crucial role in the analysis, e.g., to apply the Martingale Representation Theorem (see, e.g., Sannikov (2008) and a large subsequent literature), but it is a priori *with* loss of generality. See Section **??**.

$\{Y_s\}_{s\le t}$ the principal's belief about X_t need not coincide with Y_t. For example, if the agent often underreports his cash flow in equilibrium, the principal's belief about X_t will be higher than the report Y_t. The question, then, is to determine how much information the principal can back out about X from observing the process Y and knowing the reporting strategy L. Since L can depend arbitrarily on history, the question would seem a priori hard to answer.

Using recent advances in the analysis of functional stochastic differential equations,[4] this paper provides the following Observability Theorem: (i) Under very mild[5] regularity conditions on L, the principal can infer the true process X, (ii) Even without such conditions, there always exists a *belief equation* based on the principal's information, which, whenever it has a well-defined solution, reveals the true cash flow X.

The Observability Theorem implies that the principal knows the agent's cash flow and continuation utility at all times and that parties have symmetric information at all times. This addresses the difficulty raised by private information and validates the state consistency approach.

Building on this result, the paper provides the following Revelation Principle: any state-consistent equilibrium is outcome equivalent to a truthful state-consistent equilibrium, as long as the transformation group used to define state consistency satisfies a simple monotonicity condition.

Renegotiation-Proof Contracts with Persistent Private Information: The main application studied in this paper revisits a central question in macroeconomics:[6] the agent who generates the cash flow is risk averse and the principal is risk neutral. The agent reports and transfers his cash flow to the principal and receives in exchange a transfer from the principal (equivalently, the agent is subsidized or taxed depending on his cash flow report). The principal proposes a *contract* (a transfer process adapted to the agent's report process) to the agent, and an equilibrium consists of a contract and a reporting strategy for the agent. The principal wishes to give the agent some expected lifetime utility at the smallest possible cost. The basic tension is that, in order to properly insure the agent, the principal must know the agent's cash flow, but the agent may benefit from underreporting cash flow to get a higher subsidy (or lower tax) or, conversely, overreport his cash flow to get rewarded by a higher continuation utility.

From the principal's perspective there are two state variables: the current cash flow and the agent's continuation utility. When the agent has exponential utility, a natural class of equilibrium trans-

[4] A functional SDE is an SDE whose drift and volatility at time t depend on the path of the solution until t.

[5] The theorem provides two independent sufficient conditions: a local Lipschitz condition and an "arbitrarily small delay" condition, which means that the agent's lying strategy cannot depend on what he observed in the last ε units of time, where ε is arbitrarily small. This condition is particularly mild because the process X is continuous. The theorem also requires a local boundedness condition.

[6] See Thomas and Worrall (1990), Williams (2011), Bloedel, Krishna, and Leukhina (2020) and references therein.

formations from one state to another emerges, which creates a rich and well structured class of challengers for the equilibrium and its continuations.

Together with an additional comparison (to convex combinations of equilibria), state consistency pins down the structure of all renegotiation-proof equilibria and reduces the family of renegotiation-proof equilibria to a simple, low-dimensional family, which may be viewed as the quotient of the set of renegotiation-proof equilibria with respect to their orbits.

Any truthful renegotiation-proof contract (i.e., a contract with which truth-telling forms an equilibrium) is characterized by a single "sensitivity" parameter, which determines the agent's incentive to truthfully report his cash flows. For such a 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.

The class of renegotiation-proof contracts contains the contract studied by Williams (2011) and is equivalent to the class of stationary contracts studied by Bloedel, Krishna, and Strulovici (2020), who provide a self-insurance implementation of these contracts. The optimal renegotiation-proof contract is obtained by maximizing a closed-form objective with respect to the sensitivity parameter, which makes it easy to derive comparative statics of the optimal sensitivity parameter with respect to information persistence, discounting, and risk. The agent's flow utility optimally has a negative drift for all parameters of the model, but the occurrence of immiserisation depends on the strength of cash flow persistence and volatility (Bloedel, Krishna, and Strulovici).

Reporting Incentives and Reporting Jumps: Reporting incentives are linear for any arbitrary contract, which implies that the agent either is indifferent between telling the truth and lying, or wishes to lie at maximal (infinite) rate, either upwards or downwards. To account for this, the agent's strategy space is enlarged to allow jumps in the agent's reports. For the contracts characterized in this paper, there is a natural way to specify how such jumps affect the agent's continuation utility. 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 of dealing with the possibility of unbounded drift of the reporting process. This technique is used to derive in closed form the agent's value function not only on the equilibrium path, but also after any possible deviation.[7]

With truthful renegotiation-proof contracts, the agent wants to report cash flows truthfully not

[7]Another approach, which does not use jumps is to allow the agent's optimal strategy to be ill defined off the equilibrium path. In this case, the agent's value function is only a viscosity supersolution of the agent's HJB equation. This approach is explained in Bloedel, Krishna, and Strulovici (2020).

only on the equilibrium path, but also *after any possible deviation.* If, there was any mistake in the report, it is strictly optimal for the agent to immediately correct this mistake.[8]

Literature Review: This paper proposes a way to analyze renegotiation in contractual relationships with persistent state variables, and applies it to study environments with persistent private information.[9] The concept may be viewed as an adaptation and strengthening of the concepts of internal consistency introduced by Bernheim and Ray (1989) and Farrell and Maskin (1989) to settings with persistent states.[10] Gromb (1994) studies a binary-state model of debt contracts and compares payoffs across the two states, similarly to the approach of this paper. Bassetto, Huo, and Rios-Rull (2020) study "organizational equilibria" in a dynamic game played a succession of time inconsistent agents, in which capital is a state variable. They impose a *weak separability* condition on agents' preferences, which stipulates that an agent's preference over sequences of actions is independent of the current level of capital. This weak separability condition allows comparisons of continuation equilibria across different levels of capitals. When comparing continuation equilibria, they impose a "no-restarting" condition, which means that there is no time t such that agents' would prefer to adopt an earlier continuation equilibrium over the current one. This condition may be viewed as a one-sided notion of internal consistency.[11] Ray (1994) proposes a reinforcement of internal consistency, which may be interpreted as follows: if parties are aware of all continuations of an equilibrium (as they should under internal consistency), they could build new equilibria recursively by using the set of continuation payoffs to incentivize current-period actions (see also Van Damme (1991)). Maestri (2017) studies explicit, dynamic renegotiation between a principal and an agent with a private type.[12]

The paper also builds on the literature on dynamic contracting with persistent private information initiated by Fernandes and Phelan (2000). Most modeling features of this paper are based on Williams (2011) who focuses on full commitment.[13] The model is related to optimal insurance and

[8]This feature is interesting, for instance, 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. The contracts characterized here give the agent the incentive to correctly book a nonrecurring loss on the firm's accounts.

[9]Hart and Tirole (1988), Laffont and Tirole (1990) Dewatripont (1989), Fudenberg and Tirole (1990), and Battaglini (2007) study contract renegotiation with private information in finite period models, which impose de facto constraints on the frequency of communication. Maestri (2017) studies an infinite horizon version of Hart and Tirole (1988) and finds that information is revealed arbitrarily quickly relative to the discount rate as parties become arbitrarily patient.

[10]See also Pearce (1987), Abreu and Pearce (1991), and Abreu et al. (1993).

[11]They also impose a no-delay condition, which has no equivalent in the present paper, and a notion of optimality which is similar to external stability in Bernheim and Ray (1989).

[12]Kranz and Ohlendorf (2013), Miller and Watson (2013), and Safronov and Strulovici (2018) consider explicit renegotiation in repeated games.

[13]See also Tchistyi (2006), Zhang (2009), and Kapička (2013). Doepke and Townsend (2006) and Fukushima and

taxation models studied by Green (1987), Thomas and Worrall (1990), Golosov et al. (2003), Farhi and Werning (2013). Golosov and Iovino (2020) study optimal insurance without commitment. Pavan, Segal, and Toikka (2014) study persistence in discrete time through "impulse response functions." Contracting with persistent private information also arises in delegated experimentation models (Bergemann and Hege (2005), Garfagnini (2011), Hörner and Samuelson (2013)), as an agent's private effort to learn about some technology may result in persistent superior information.[14]

The paper is organized as follows. Section 2 introduces illustrates the concept and implications of state consistency in a simple setting with discrete time and with a persistent but publicly observable state. Section 3 presents the contractual setting with persistent private information. Section 4 proposes a general formalization of state consistency. Section 6 characterizes truthful, renegotiation-proof contracts in the setting of Section 3. Section 7 studies how the sensitivity of the optimal contract to the agent's reports varies with persistence, discounting, and risk. Section 8 establishes an Observability Theorem and a Revelation Principle to address private information. Section 9 provides a necessary and sufficient condition for truthfulness on and off the equilibrium path. Section 10 discusses several extensions.

2 Discrete-Time Example with Persistent Public Information

2.1 Setting

An agent generates an output process $\{X_t\}_{t\in\mathbb{N}}$ such that

$$X_{t+1} = X_t + D(A_t, Z_t)$$

where $\{Z_t\}_{t\in\mathbb{N}}$ is a sequence of i.i.d. shocks, $A_t \in \mathcal{A}$ is the agent's effort at time t, D is a real-valued function, and $X_0 = x \in \mathbb{R}$ is given.[15] The agent incurs a cost of effort $h(A_t)$ at time t where h is some arbitrary function $h : \mathcal{A} \to \mathbb{R}$. A *contract* C maps each history $\{X_0, \ldots, X_t\}$ to a consumption level $C_t \in \mathbb{R}$. The principal and the agent have the same discount factor $\rho \in (0, 1)$. Given a contract C, the agent chooses an effort process A mapping histories into effort levels. Formally, an effort

Waki (2009) provide novel numerical methods and analysis.

[14] See also Guo (2016) and Halac, Kartik, and Liu (2016). Similarly in DeMarzo and Sannikov (2016), principal and agent both learn the agent's skill but the agent privately observes his effort. Sannikov (2014) considers a moral hazard problem in which the agent's action have long-term consequences. Garrett and Pavan (2012, 2015) study managerial compensation contracts when the type of the manager is persistent. These papers do not consider renegotiation.

[15] The example can be generalized to $X_{t+1} = \lambda X_t + D(A_t, Z_t)$ where $\lambda \in (0, 1]$. In this case, the transformation Φ_δ defined by equation (1) must be changed to $\Phi_\delta(C)(X'_0, \ldots, X'_t) = C(X'_0 - \delta, X'_1 - \delta\lambda \ldots, X'_t - \delta\lambda^t)$, with a similar change for the transformation of the agent's effort strategy, and the term $x/(1-\rho)$ in the payoff formulas appearing in Propositions 2 and 3 must be replaced by $x/(1-\rho\lambda)$. For expositional simplicity, we focus here on $\lambda = 1$.

process A is *admissible* if it is adapted to the filtration generated by the output process X. We let $\mathbf{A}$ denote the set of admissible effort processes.

Given a contract C and an admissible strategy A, the agent's expected discounted utility is

$$V(C, A; x) = E\left[\sum_{t \in \mathbb{N}} \rho^t (u(C_t) - h(A_t)) | X_0 = x\right]$$

for some real-valued utility function u and the principal's expected discounted payoff is

$$\Pi(C, A; x) = E\left[\sum_{t \in \mathbb{N}} \rho^t (X_t - C_t) | X_0 = x\right]$$

A pair (C, A) is a *contractual equilibrium* if A maximizes $V(C, \tilde{A}; x)$ over all admissible effort processes $\tilde{A} \in \mathbf{A}$. A contractual equilibrium (C, A) *starts from* (x, w) if $X_0 = x$ and $V(C, A; x) = w$.

2.2 Internal Consistency

In the context repeated games, two well-known and essentially identical concepts of renegotiation are *internal consistency* (Bernheim and Ray (1989)) and *weakly renegotiation-proofness* (Farrell and Maskin (1989)). According to these concepts, an equilibrium is renegotiation-proof if it does not have any continuation equilibrium that is Pareto dominated by another continuation equilibrium.

Although the concept internal consistency was originally introduced in the context of repeated games, it can be embedded into dynamic games with non-trivial state variables: say that an equilibrium is internally consistent if there do not exist two histories *leading to the same state* such that the players' continuation payoffs following the first history Pareto dominate those following the second history.

In the context of Section 2.1, a contractual equilibrium (C, A) is *internally consistent* if there do not exist two histories $\{X_s : s \leq t\}$ and $\{X'_s : s \leq t'\}$ leading to the same output level $X_t = X'_{t'} = x$ and the same continuation utility w for the agent such that principal gets a strictly higher continuation payoff after history $\{X'_s : s \leq t'\}$ than after $\{X_s : s \leq t\}$.

Intuitively, if a contractual equilibrium violates internal consistency and the dominated history $\{X_s : s \leq t\}$ is realized, the principal would strictly benefit from switching to the equilibrium continuation following $\{X'_s : s \leq t'\}$ and the agent would accept this switch, with a strict incentive to do so if the principal shares even an arbitrarily small fraction of the surplus gained from the switch. Internal consistency implies the following result, whose proof is straightforward and omitted. Given a contractual equilibrium (C, A) and history up to time t, let W_t denote the agent's continuation utility at time t.

Proposition 1 *If (C, A) is internally consistent, there exists a function $\mathbf{\Pi} : \mathbb{R}^2 \to \mathbb{R}$ such that the principal's continuation payoff process satisfies $\Pi_t = \mathbf{\Pi}(X_t, W_t)$ for all t.*

2.3 State Consistency

Beyond Internal Consistency: To see how internal consistency may be strengthened, it is useful to revisit about its rationale. Internal consistency presumes that, after observing some history, the principal is able to recognize that she could use the continuation equilibrium following another history to achieve a higher payoff. This cognitive ability should extend to other natural comparisons.

One such comparison, often used in economics, is convexification: if two distinct continuation equilibria give the same expected utility to the agent, and the agent has a concave utility, it is natural for the principal to consider convex combinations of these continuation equilibria, as they may achieve the same expected utility for the agent at a lower cost for the principal. This idea will be exploited in Section 6.2 to narrow down the set of renegotiation-proof equilibria. For now, we focus on comparisons across states, which constitute the main conceptual innovation of the paper.

Transformations Across Outputs: Fix a contractual equilibrium (C, A) starting from (x, w) and another output level $x' \in \mathbb{R}$, and let $\delta = x' - x \in \mathbb{R}$.

We define a new contract $\Phi_\delta(C)$ by

$$\Phi_\delta(C)(X'_0, \dots, X'_t) = C(X'_0 - \delta, \dots, X'_t - \delta). \tag{1}$$

We also define a new effort strategy $\Phi_\delta(A)$ by

$$\Phi_\delta(A)(X'_0, \dots, X'_t) = A(X'_0 - \delta, \dots, X'_t - \delta) \tag{2}$$

for any history $\{X'_0, \dots X'_t\}$.

The pair $(\Phi_\delta(C), \Phi_\delta(A))$ forms a contractual equilibrium starting from (x', w), which may be see as follows. Consider any strategy A' starting from x'. Given the contract $\Phi_\delta(C)$ and initial condition x', A' generates a stochastic output path $X'_0, X'_1, \dots$ with $X'_0 = x'$ and results in a consumption path given by (1).

Consider now the effort strategy $\Phi_{-\delta}(A')$ obtained by applying transformation (2) to strategy A' and shift parameter $-\delta$ rather than δ: for any history $(X_0, \dots X_t)$, we have $\Phi_{-\delta}(A')_t = A'(X_0 + \delta, \dots, X_t + \delta)$. Notice that for any strategy A and $\delta \in \mathbb{R}$, we have $\Phi_{-\delta}(\Phi_\delta(A)) = A$.

For any realization of the exogenous uncertainty $(Z_0, Z_1, \ldots)$, a simple forward induction on time t, starting from $t = 0$, shows the following: the strategies A' and $\Phi_{-\delta}(A')$, starting respectively from x' and x are such that $A'_t = \Phi_{-\delta}(A')_t$ and $X'_t = X_t + \delta$ for all t.

From (1), this implies that the agent receives the same consumption for all t in both of these cases. Since the agent's effort and consumption is the same across both contracts for each realization of the shocks $Z = (Z_0, \ldots)$, we conclude that

$$V(\Phi_\delta(C), A'|x + \delta) = V(C, \Phi_{-\delta}(A'); x)$$

for all A'. Equivalently,

$$V(\Phi_\delta(C), \Phi_\delta(A)|x + \delta) = V_x(C, A|x)$$

for all A. Therefore, we conclude that A is optimal given C and x if and only if $\Phi_\delta(A)$ is optimal given $\Phi(C)$ and x'.

Continuation of a Contractual Equilibrium: In order to compare continuations of contractual equilibria, it is useful to treat formally each such continuation as a contractual equilibrium with time reset to 0.

Formally, given any contractual equilibrium (C, A) and any history $\{X_0, \ldots, X_t\}$ leading to some continuation utility $W_t = w$ and output $X_t = x$, the continuation $(\hat{C}, \hat{A})$ of (C, A) following history $\{X_0, \ldots, X_t\}$ is a contractual equilibrium starting from $\hat{X}_0 = x$ and $\hat{W}_0 = w$, and such that for any $\tau \geq 0$ and $(\hat{X}_0, \ldots . \hat{X}_\tau)$, we have

$$\hat{C}(\hat{X}_0, \ldots, \hat{X}_\tau) = C(X_0, \ldots, X_{t-1}, X_t = \hat{X}_0, \hat{X}_1, \hat{X}_2, \ldots \hat{X}_\tau)$$

and

$$\hat{A}(\hat{X}_0, \ldots, \hat{X}_\tau) = A(X_0, \ldots, X_{t-1}, X_t = \hat{X}_0, \hat{X}_1, \hat{X}_2, \ldots \hat{X}_\tau),$$

where we indicated for clarity that $X_t = \hat{X}_0$ to emphasize that the end of history $\{X_0, \ldots, X_t\}$ is the beginning of $(\hat{C}, \hat{A})$'s history.

Consistency Across Outputs:

A contractual equilibrium (C, A) is *consistent across outputs* if the following holds for any two histories $\{X_0, \ldots, X_t\}$ and $\{X'_0, \ldots, X'_{t'}\}$ leading to the same continuation promised utility W_t for the agent. Denoting by $X_t = x$ and $X'_{t'} = x + \delta$ the output levels after these respective histories, and letting $(\hat{C}, \hat{A})$ denote continuation of (C, A) after $\{X_0, \ldots, X_t\}$, the principal's continuation payoff after $\{X'_0, \ldots, X'_{t'}\}$ is weakly greater than her payoff under the transformation $(\Phi_\delta(\hat{C}), \Phi_\delta(\hat{A}))$ of the continuation contract $(\hat{C}, \hat{A})$.

Intuitively, consistency across outputs means that the principal can think, after history $\{X'_0, \ldots, X'_{t'}\}$, of using the transformation of the continuation following $\{X_0, \ldots, X_t\}$ that makes this continuation compatible with output level $x+\delta$, and switch to this transformation if it is profitable. Consistency across outputs implies internal consistency, as can be seen by setting $\delta = 0$ in the definition.

PROPOSITION 2 *If (C, A) is consistent across outputs, then (i) the principal's payoff function $\mathbf{\Pi}$ is additively separable: there exist functions f and g such that $\mathbf{\Pi}(x, w) = f(x) - g(w)$ and (ii) f is linear in x: $f(x) = \frac{x}{1-\rho}$.*

Proof. Fix $w \in \mathbb{R}$ and consider any $x \neq x'$ such that the states (x, w) and (x', w) can both be reached by some histories. Given a continuation (C, A) starting from (x, w), the transformation of (C, A) constructed above for $\delta = x' - x$ starts from (x', w) and gives the principal an expected payoff equal to $\mathbf{\Pi}(x, w) + \frac{x'-x}{1-\rho}$, because the expected output level at all future periods is uniformly translated by $x' - x$. Consistency across outputs then implies that $\mathbf{\Pi}(x', w) \geq \mathbf{\Pi}(x, w) + \frac{x'-x}{1-\rho}$. The reverse transformation implies the reverse inequality. Hence, $\mathbf{\Pi}(x', w) = \mathbf{\Pi}(x, w) + \frac{x'-x}{1-\rho}$ for all (x, w) and (x', w) reachable by some histories. This implies that there exists some real number $g(w)$ such that $\mathbf{\Pi}(x, w) = \frac{x}{1-\rho} - g(w)$ for all x such that (x, w) can be reached. For x'' such that (x'', w) is not reached by any history, we can define the function $\mathbf{\Pi}(x'', w)$ using the same formula. Since this transformation works for each w, we get the desired result.

Transformation Across Promised Utilities: We now specialize the setting further, to create transformations across promised utilities of the agent. We now assume that consumption and effort levels are nonnegative at all times, that the agent's utility and cost functions are given by $u(c) = c^{\gamma_1}$ and $h(a) = a^{\gamma_2}$ for some parameters $0 < \gamma_1 < 1 < \gamma_2$, and that the output shocks take the following multiplicative form:[16]

$$X_{t+1} = X_t + A_t Z_t.$$

Under these assumptions, the agent's utility in any continuation of any contractual equilibrium must be nonnegative, because the agent can always abstain from putting any effort and receive at least zero utility from consumption in each period.

Given any contractual equilibrium (C, A) starting from state (x, w) and any $\beta > 0$, we now construct a contractual equilibrium (C', A') starting from $(x, \beta w)$.

We start with the following observation: since the initial condition x is fixed and known to the principal, observing the output process $(X_0, X_1, \ldots)$ is equivalent to observing the increments $(D_0 = A_0 Z_0, D_1 = A_1 Z_1, \ldots)$. We can therefore express C and A as functions

$$C_t(D_0, \ldots, D_{t-1})$$

[16] The example extends easily to the dynamics $X_{t+1} = X_t + dA_t Z_t$ where d is a positive constant. See also footnote 15.

and

$$A_t(D_0,\ldots,D_{t-1})$$

for all t.

We define a new contract starting from $(x,\beta w)$ as follows. Given observed increments $D'_0, D'_1, \ldots D'_{t-1}$, let

$$\phi_\beta(C)_t(D'_0, D'_1, \ldots, D'_{t-1}) = \beta^{1/\gamma_1} C_t(D'_0\beta^{-1/\gamma_2}, \ldots, D'_{t-1}\beta^{-1/\gamma_2}).$$

We also define a new strategy $\phi_\beta(A)$ for the agent, by

$$\phi_\beta(A)_t(D'_0, D'_1, \ldots, D'_{t-1}) = \beta^{1/\gamma_2} A_t(D'_0\beta^{-1/\gamma_2}, \ldots, D'_{t-1}\beta^{-1/\gamma_2}).$$

We now show that $(\phi_\beta(C), \phi_\beta(A))$ is a contractual equilibrium starting from $(x,\beta w)$ if and only if (C,A) is a contractual equilibrium starting from (x,w).

Consider any strategy A' starting from initial condition $(x,\beta w)$ and any realization of the exogenous uncertainty $(Z_0, Z_1, \ldots)$, the strategy A' defines an increment process $D' = (D'_0 = A'_0 Z_0, D'_1 = A'_1 Z_1, \ldots)$. Consider the strategy $\phi_{1/\beta}(A')$ defined by

$$\phi_{1/\beta}(A')_t(D_0, \ldots D_{t-1}) = \beta^{-1/\gamma_2} A'_t(D_0\beta^{1/\gamma_2}, \ldots, D_{t-1}\beta^{1/\gamma_2}) \tag{3}$$

for all t. Notice that for all A and $\beta > 0$, we have $\phi_{1/\beta}(\phi_\beta(A)) = A$.[17]

By construction, the strategies A' and $\phi_{1/\beta}(A')$ generate increment processes D' and D such that

$$D'_t = \beta^{1/\gamma_2} D_t.$$

for all $t \geq 0$. Moreover, by definition of the contract $\phi_\beta(C)$, the agent's consumptions at time t under $(C, \phi_{1/\beta})(A))$ and $(\phi_\beta(C), A')$ are related by

$$\phi_\beta(C)_t = C_t \beta^{1/\gamma_1} \tag{4}$$

for all t. This implies that the agent's utility from consumption and cost of effort satisfy $u(\phi_\beta(C)_t) = \beta u(C_t)$ by (4) and $h(A'_t) = \beta h(\phi_{1/\beta}(A')_t)$ by (3) at all times. Since these equalities hold for all realizations of the exogenous shocks, we conclude that

$$V(\phi_\beta(C), A'; x) = \beta V(C, \phi_{1/\beta}(A'); x)$$

for all A'. Equivalently,

$$V(\phi_\beta(C), \phi_\beta(A); x) = \beta V(C, A); x).$$

[17]Formally, this means that the map $\phi : (\beta, A) \mapsto \phi_\beta(A)$ defines an action of the group $((0,+\infty),\times)$ over the set $\mathbf{A}$ of admissible strategies. This point is developed in Section 4.

for all A. This shows that (C, A) is a contractual equilibrium starting from (x, w) if and only if $(\phi_\beta(C), \phi_\beta(A))$ is contractual equilibrium starting from $(x, \beta w)$.

State Consistency: We define state consistency in two steps. First, say that a contractual equilibrium (C, A) is *consistent across promised utilities* if the following holds for any two histories $\{X_0, \ldots, X_t\}$ and $\{X'_0, \ldots, X'_{t'}\}$ leading to the same output level $X_t = X'_{t'} = x$: denoting by w and βw the continuation utility levels right after these histories, the principal's continuation payoff after $\{X'_0, \ldots, X'_{t'}\}$ is weakly greater than her continuation payoff under the transformation $(\phi_\beta(\hat{C}), \phi_\beta(\hat{A}))$ of the continuation $(\hat{C}, \hat{A})$ of (C, A) after $\{X_0, \ldots, X_t\}$.

Second, say that a contractual equilibrium is *state consistent* if it is consistent across outputs and across promised utilities.

Proposition 3 *Suppose that the contractual equilibrium (C, A) is state consistent. Then, there exist positive constants α_1, α_2 such that the principal's payoff function $\mathbf{\Pi}$ satisfies*

$$\mathbf{\Pi}(x, w) = \frac{x}{1-\rho} + \alpha_2 w^{1/\gamma_2} - \alpha_1 w^{1/\gamma_1}.$$

The next proposition shows that state-consistent contractual equilibria and can be constructed explicitly.

Proposition 4 *There exists a continuum of state-consistent equilibria, which have the following form:*[18] *letting W_t denote the agent's continuation utility at time t, we have*

$$A_t = a^{1/\gamma_2} W_t^{1/\gamma_2}$$

$$C_t = c^{1/\gamma_1} W_t^{1/\gamma_1}$$

where a and c are positive constants related by

$$(\gamma_2 - 1)a = 1 - c. \tag{5}$$

In these equilibria, the effort and consumption levels at any time are only a function of the agent's continuation utility at that time, and these functions are pinned down by a one dimensional parameter (c, say), which fully determines, together with (5), each contract. To fully characterize the contract, one has to determine how the promised utility W_t depends on observed output. This is done in the proof of Proposition 4 and this dependence is also pinned down by c.

[18]The statement concerns a specific class of state consistent equilibria. It does not rule out the existence of other state consistent equilibria.

3 Persistent Private Information

We now turn to the main analysis of this paper, which includes persistent private information and formalizes the construction of transformation groups introduced in the example.

An agent generates cash flow $X_t \in \mathbb{R}$ at time $t \geq 0$, which evolves according to the dynamic equation[19]

$$dX_t = [(\xi - \lambda X_t)]\, dt + \sigma dB_t \tag{6}$$

where B is the standard Brownian motion and $X_0 = x$ is given and commonly known. 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. λ is the rate at which a shock in the current cash flow decays over time. Uncertainty is modeled with a probability space $(\Omega, \mathcal{F}, P)$ satisfying the usual conditions and whose outcomes ω are identified with the paths of B.

The agent reports and transfers to the principal a cash flow Y_t that obeys the dynamic equation

$$dY_t = dX_t + L_t dt = [(\xi - \lambda X_t) + L_t]\, dt + \sigma dB_t \tag{7}$$

where L_t is the rate at which the agent *lies* about the increment dX_t of the true cash flow.[20]

The *gap* $G_t = Y_t - X_t$ between reported and actual cash flows satisfies[21]

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

The agent's initial cash flow is assumed to be known by the principal.

Assumption 1 *(i) The principal observes the report process $\{Y_t\}_{t>0}$ but not the actual cash flow process $\{X_t\}_{t>0}$. (ii) $Y_0 = X_0$.*

Contract: A contract is a stochastic process $C = \{C_t\}_{t\geq 0}$ that is adapted to the filtration $\mathbb{F}^Y = \{\mathcal{F}^Y\}_{t\geq 0}$ generated by the report process Y. The process C, discounted at rate r, is assumed to be integrable when the agent tells the truth: $E[\int_0^\infty e^{-rt}|C_t|dt \ |L \equiv 0] < \infty$.

[19]The model is due to Williams (2011). An earlier version of this paper (Strulovici (2011)) includes moral hazard: the agent privately chooses some effort level $A_t \in \mathbb{R}$ that affects the cash flow dynamics according to $dX_t = (A_t + \xi - \lambda X_t)dt + \sigma dB_t$. This addition does not affect the main results, as explained in Section 10.2.

[20]The rate at which the agent lies may be unbounded. To address this, Section 9 allows jumps in the agent's report process and proposes a natural extension of the contract to this case.

[21]One could impose constraints on the lies that the agent can make, e.g., by requiring that $Y_t \leq X_t$ (the agent cannot transfer more than he earns) or that $G_t \leq \bar{g}$ for some arbitrary $\bar{g} > 0$ (the agent cannot overreport more than a fixed amount). The contracts derived in this paper remain truthful in the presence of such constraints and for any $\bar{g} > 0$, no matter how small, the necessary conditions for truth-telling are the same as if G is unconstrained.

A contract C requires the principal to make a transfer $C_t \in \mathbb{R}$ to the agent at time t: the agent gives Y_t to the principal and the principal gives C_t to the agent. An alternative interpretation is that the agent reports and keeps Y_t and the principal gives a (possibly negative) subsidy $C_t - Y_t$ to the agent. These interpretations are formally equivalent and the former is used throughout the paper for consistency.

We rule out private savings and assume that the agent immediately consumes the transfer that he receives from the principal, as well as any difference between his real and reported cash flows:[22]

Assumption 2 *The agent's consumption at time t is equal to $C_t + (X_t - Y_t) = C_t - G_t$.*

The agent's strategy consists of a lying process L adapted to the agent's information, which corresponds to the filtration $\mathbb{F}^X = \{\mathcal{F}_t^X\}_{t\geq 0}$ generated by X.[23]

Given an initial condition x, a contract C, and a strategy L, the agent's expected discounted utility is[24]

$$V_0(C, L; x) = E\left[\int_0^\infty e^{-rt}\left(u(C_t + X_t - Y_t)\right) dt\right] \tag{8}$$

where u is strictly concave and X, Y evolve according to (6) and (7) and the initial conditions $X_0 = Y_0 = x$. In computations to follow, the agent will be assumed to have exponential utility $u(c) = -\exp(-\theta c)$ for some risk-aversion coefficient $\theta > 0$.

A strategy L is *admissible* if the following transversality or "No-Ponzi" condition is satisfied. Let $V_T(L)$ denote the agent's continuation utility at time T when control L is applied until time T and truthtelling is used thereafter.[25]

$$\lim_{T\to+\infty} E\left[e^{-rT}|V_T(L)|\right] = 0 \quad a.s. \tag{9}$$

Definition 1 *Given an initial condition x, a* contractual equilibrium (C, L) *consists of a contract C and a strategy L for the agent that solves*

$$\sup_{L'} V_0(C, L'; x)$$

over all admissible strategies.

[22]See Bloedel, Krishna, and Strulovici (2020) for the case of hidden savings.

[23]The agent also observes Y. However, since X determines Y, given the agent's strategy, X is a sufficient statistic for the agent's information.

[24]The strategy L affects the probability measure over the paths of Y, which affects the expectation.

[25]Admissibility rules out, for instance, strategies in which the agent continually underreports his cash flow, which permits him to get a higher immediate transfer from the principal but leads to an ever-decreasing continuation utility, with a decrease rate that exceeds r. Condition (9) is standard. It is used to take limits in verification arguments, to check that the solution of an HJB equation is equal to the value function of the agent's optimization problem.

A contractual equilibrium requires that the strategy L be incentive compatible given the contract C and initial condition x. The contract C is not required to satisfy any incentive compatibility condition for the principal.

Objective of the Principal: Given an initial condition x and a contractual equilibrium (C, L), the principal's expected payoff is

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

where Y is given by (7), X follows the dynamic equation (6), and $X_0 = Y_0 = x$. The objective of the principal is to maximize her expected payoff subject to giving the agent some minimal expected lifetime utility w (i.e., $V_0 \geq w$) and to renegotiation-proofness constraints that are the focus of the next section.

4 Concept of Renegotiation

Any model of renegotiation entails a comparison between some current agreement and alternative agreements. The key is to determine the set of challengers that parties may consider as valid alternatives to the current agreement. This section proposes an approach to do this when there are persistent state variables.

4.1 Cash Flow and Continuation Utility as State Variables

In our setting, one of the state variables is the current cash flow X_t generated by the agent. Since this cash flow is persistent, it affects the set of achievable payoffs for the players at any given time.

In addition, it is useful to treat the agent's continuation utility V_t as a state. We take the perspective of the principal and consider whether there exist contractual equilibria that give the agent the same utility V_t as the current equilibrium but increase the principal's payoff Π_t given the current cash flow X_t.

4.2 Transformation Groups

We formalize state-consistency using the language of group theory. This formalization makes it easy to envision the application of the concept of state consistency beyond the specific applications considered in the present paper.

Let

- $\mathcal{S}$ denote the state space. In our application, $\mathcal{S}$ consists of all pairs (v, x) of continuation utility and cash flow over the relevant domain.
- $\mathbf{G}$ denote a group—in the algebraic sense—that exerts a *left action* on $\mathcal{S}$. This means that for any $g \in \mathbf{G}$ and any $(v, x) \in \mathcal{S}$, one can associate an element $(\tilde{v}, \tilde{x}) = g \circ (v, x) \in \mathcal{S}$ that respects the group's structure, as follows:[26]
  - $e \circ (v, x) = (v, x)$ for the identity element e of $\mathbf{G}$
  - $(gg') \circ (v, x) = g \circ (g' \circ (v, x))$ for any $g, g' \in \mathbf{G}$.

Now consider the set $\mathcal{E}$ of all contractual equilibria (C, L). Each contractual equilibrium is associated with an initial state (v, x).

To define a left action group on $\mathcal{E}$, we associate for each $g \in \mathbf{G}$ and $(C, L) \in \mathcal{E}$, a new contractual equilibrium $\Phi_g(C, L)$ that respects the structure of the group, as follows:

Definition 2 *The mapping from $\Phi : \mathbf{G} \times \mathcal{E} \to \mathcal{E}$ defined by $(g, (C, L)) \mapsto \Phi_g(C, L)$ is a* transformation group *if it satisfies the following axioms:*

> **Axiom 1**. *For any $(C, L) \in \mathcal{E}$, $\Phi_e(C, L) = (C, L)$.*
>
> **Axiom 2**. *For all $g, g' \in \mathbf{G}$ and $(C, L) \in \mathcal{E}$, $\Phi_{g'}(\Phi_g(C, L)) = \Phi_{g'g}(C, L)$.*
>
> **Axiom 3**. *For any $g \in \mathbf{G}$ and $(C, L) \in \mathcal{E}$ starting from state (v, x), the contractual equilibrium $\Phi_g(C, L)$ starts in state $(\tilde{v}, \tilde{x}) = g \circ (v, x)$.*

Intuitively, a transformation group describes analogies that a principal can use to compare equilibria across states. The definition applies equally well to equilibria to discrete-time and continuous-time settings.

Axioms 1 and 2 are the defining properties of a left group action: Axiom 1 says that if the state is unchanged (i.e., the group identity e is applied), then the transformation is the identity mapping over $\mathcal{E}$, and Axiom 2 is an associativity axiom that will be used shortly, applied to group elements that are inverse of each other. Axiom 3 requires that the transformed contractual equilibria be consistent with the effect that the group has on the state: in the translation example above, for instance, it means that if g transforms the state v into $\tilde{v}$ when g operates on the state space $\mathcal{S}$, then g should also transform any contractual equilibrium that gives expected utility v to the agent into one that gives him expected utility $\tilde{v}$.

[26] These two conditions define the left action. The group $\mathbf{G}$ must satisfy the standard definition of a group, which include the existence of an identity element and an inverse for each element of the group.

4.3 State Consistency for Truthful Contracts

We begin the analysis by studying state consistency for truthful equilibria, defined as follows.

Definition 3 *A contractual equilibrium* (C, L) *is* truthful *if* $L \equiv 0$*. A contract* C *is* truthful *if* $(C, L \equiv 0)$ *is a contractual equilibrium.*

Truthful contracts are a good starting point for the analysis of renegotiation because the principal unambiguously knows the current cash-flow level and the agent's continuation utility, and can use this information to renegotiate the contract. Section 8 shows that the focus on truthful contracts is without loss of generality as long as the strategy of the agent satisfies one of the following regularity properties: it is either locally Lipschitz with respect to the true cash-flow process (a property satisfied, e.g., by truthtelling), or it responds with an arbitrarily small lag to news arrival.

We define a concept of state consistency, according to which challengers are obtained through the transformation group Φ. Given any contractual equilibrium (C, L), let $\Pi(C, L)$ denote the principal's expected payoff. Although we focus for now on truthful equilibria, we emphasize that the concepts in this section apply to all contractual equilibria.

Definition 4 *A contractual equilibrium* (C, L) *starting from state* (v, x) *is* state consistent *with respect to* Φ *if, after any history leading to some state* $(\tilde{v}, \tilde{x}) = g \circ (v, x)$ *and continuation contractual equilibrium* $(\tilde{C}, \tilde{L})$*, we have*

$$\Pi(\tilde{C}, \tilde{L}) \geq \Pi\left(\Phi_g(C, L)\right),$$

and, reciprocally,

$$\Pi(C, L) \geq \Pi\left(\Phi_{g^{-1}}(\tilde{C}, \tilde{L})\right).$$

State consistency entails two requirements: first, each continuation of (C, L) must sustain the comparison with the transformation of (C, L) to the state corresponding to this continuation.[27] Second, the initial contractual equilibrium must sustain the comparison with each continuation equilibrium transformed back to the initial state, i.e., corresponding to the group element g^{-1}, since $g^{-1} \circ (g \circ (v, x)) = e \circ (v, x) = (v, x)$.

A first observation is that a state-consistent contractual equilibrium is internally consistent as long as Φ satisfies the following monotonicity condition.

Definition 5 *The transformation group* Φ *is* monotone *if for any two contractual equilibria* (C, L) *and* (C', L') *starting from some common state* (v, x) *and giving payoffs* $\Pi(C, L) \leq (<) \Pi(C', L')$ *to*

[27] This definition does not distinguish between on-path and off-path continuations. This distinction is irrelevant for the setting of Section 3 since the agent's lies are absolutely continuous and the report process is always on path.

the principal, and any $g \in \mathbf{G}$*, the principal's payoffs in the transformed contractual equilibria* $(\tilde{C}, \tilde{L}) = \Phi_g(C, L)$ *and* $(\tilde{C}', \tilde{L}') = \Phi_g(C', L')$ *satisfy* $\Pi(\tilde{C}, \tilde{L}) \leq (<)\Pi(\tilde{C}', \tilde{L}')$*.*

Monotonicity means that the ranking of the principal's payoffs is preserved under transformations to a different state.

Proposition 5 *Suppose that* Φ *is monotone and that the contractual equilibrium* (C, L) *starting from state* (v, x) *is state consistent with respect to* Φ*. Then, after any finite history leading to state* $(\tilde{v}, \tilde{x}) = g \circ (v, x)$*, the continuation payoff for the principal is equal to her initial payoff in the contractual equilibrium* $\Phi_g(C, L)$*.*

This proposition, proved in the Appendix, implies the following result.

Corollary 1 *If* Φ *is monotone, any state-consistent equilibrium is internally consistent.*

The proof is immediate: if (C, L) is state consistent, Proposition 5 implies that any two histories leading to the same state $(\tilde{v}, \tilde{x}) = g \circ (v, x)$ give the same continuation payoff $\tilde{\pi}$ to the principal and, hence that the players' continuation payoffs $(\tilde{v}, \tilde{\pi})$ of both parties are not Pareto ranked.

Strength, Stability, and Uniqueness of the State-Consistency Concept: In principle, there may be several transformations groups to consider. The group $\mathbf{G}$ operating on the state space is *transitive* (in the group theoretic sense) if for any states (v, x) and (v', x') there exists some $g \in \mathbf{G}$ such that $(v', x') = g \circ (v, x)$. This property is satisfied by the group constructed in our main application (Section 5). If Φ is defined with respect to a transitive group $\mathbf{G}$, it means that we can compare contractual equilibria across any two states.

The theory of state consistency does not rely on transitivity: it is well defined for any group $\mathbf{G}$. This observation is useful if, for instance, in settings for which it is hard to establish comparisons across all pairs. The next proposition establishes two results pertaining to non-transitive groups. Given a group $\mathbf{G}$ operating on $\mathcal{S}$ and a transformation group Φ defined with respect to $\mathbf{G}$, we can consider, for any subgroup $\mathbf{G}'$ of $\mathbf{G}$ the transformation group Φ' that is the restriction of Φ with respect to $\mathbf{G}'$. A particular subgroup is the *trivial* group that contains only the identity element e. With respect to the trivial group, Φ reduces to the identity transformation over contractual equilibria. We have the following result.

Proposition 6 *(i) Suppose that* Φ *is a transformation group with respect to* $\mathbf{G}$ *and that* $\mathbf{G}'$ *is a subgroup of* $\mathbf{G}$*. A contractual equilibrium is state consistent with respect to* Φ *only if it is state consistent with respect to* Φ'*. (ii) If* $\mathbf{G}'$ *is the trivial group, a contractual equilibrium is state consistent with respect to* Φ' *if and only if it is internally consistent.*

Part (i), whose proof is straightforward, establishes a nestedness condition for concepts of state-consistency. Regarding Part (ii), we note that the identity transformation group is clearly monotonic, so internal consistency follows from Corollary 1. The reverse direction follows immediately from the definition of state consistency.

From Proposition 6, the stronger concepts of state consistency are obtained for transitive groups. Fixing a group $\mathbf{G}$, there may a priori exists multiple transformation groups with respect to $\mathbf{G}$, which could potentially lead to different solution concepts for state consistency. However, the next result shows that, for a contractual equilibrium to sustain comparisons with respect to several monotone transformation groups, these groups must, taken individually, yield the *same* concept of state consistency. Consider two monotone transformation groups Φ and $\tilde{\Phi}$. Given a contractual equilibrium (C, L) starting from (v, x), suppose that there is a state $(v', x') = g \circ (v, x)$ for which the groups yield different payoffs: $\Pi(\Phi_g(C, L)) > \Pi(\tilde{\Phi}_g(C, L))$. Applying the transformations again from (v', x') to $(v, x) = g^{-1}(v', x')$ and using monotonicity, we get

$$\Pi(\tilde{\Phi}_{g^{-1}}(\Phi_g(C, L))) > \Pi(C, L).$$

Therefore, (C, L) is dominated (from the principal's perspective) by a simple composition of elements in the two groups. In such a case, we will say that (C, L) is *unstable* with respect to $(\Phi, \tilde{\Phi})$, otherwise, (C, L) is called *stable* with respect to $(\Phi, \tilde{\Phi})$. Intuitively, instability implies that the principal's ability to consider transformations that yield different payoffs after a given history prevents the existence of a state-consistent contractual equilibrium.[28]

Proposition 7 (Concept Equivalence for Stable Contractual Equilibria) *Let Φ and $\tilde{\Phi}$ be two monotone transformation groups with respect $\mathbf{G}$ and (C, L) be a contract equilibrium starting from state (v, x) that is stable with respect to $(\Phi, \tilde{\Phi})$. Then, (C, L) is state consistent with respect to Φ if and only if it is state consistent with respect to $\tilde{\Phi}$. Moreover, the principal's payoffs under the Φ and $\tilde{\Phi}$ transformations of (C, L) to any other state (v', x') are identical.*

5 Transformation Group: Explicit Construction

This section constructs a transitive transformation group Φ explicitly for the setting of Section 3.

[28] For the setting of Section 3, it is possible to build somewhat "pathological" transformation groups different from those constructed in Section 5, for which state consistent contracts (according to the original transformation group of that section) is not stable. However, unlike the original transformation group of Section 5, these transformation groups have two potential issues: (i) they are not in the spirit of internal consistency, because they transform a given contract into contracts that are not easily obtained from the original contract, and (ii) there need exist any truthful contract that is state consistent with respect to these other transformation groups.

5.1 Transforming Contracts Across Cash Flow Levels

Given a contactual equilibrium (C, L) starting from state (v, x), and any $\hat{x} \neq x$, we construct a new contractual equilibrium that starts from $(v, \hat{x})$. Since C is adapted to the filtration $\mathbb{F}^Y$ of the agent's report process Y, it can be expressed as $C_t = \mathcal{C}(Y_s : s \leq t)$ for some functional $\mathcal{C}$ (for notational simplicity, $\mathcal{C}$ is defined on paths of Y of varying time lengths, which implicitly allows $\mathcal{C}$ to depend t).

Now suppose that the initial cash flow level is $\hat{x}$ and that the agent uses strategy L. This strategy generates a report process $\hat{Y}$, which is given by equation (7) and initial condition $\hat{Y}_0 = \hat{x}$ and the true cash flow process $\hat{X}$ satisfies (6) subject to initial condition $\hat{X}_0 = \hat{x}$. In this situation, the principal can choose to give the agent the consumption process $\hat{C}_t = \mathcal{C}(\tilde{Y}_s : s \leq t)$, where $\tilde{Y}$ is constructed from $\hat{Y}$ as follows: $\tilde{Y}_0 = x$ and

$$d\tilde{Y}_t = d\hat{Y}_t - (\xi - \lambda \hat{Y}_t)dt + (\xi - \lambda \tilde{Y}_t)dt. \tag{10}$$

Intuitively, the process $\tilde{Y}$ is the report process that the agent *would have produced* if the initial cash flow had been x instead of $\hat{x}$ and the agent had followed the *same strategy* L that generates $\hat{Y}$ when starting from $\hat{x}$. We call $\tilde{Y}$ the *virtual report process.* Notice that the principal can always compute the virtual report process $\tilde{Y}$ from observing the report process $\hat{Y}$, thanks to equation (10).

Proposing contract $\hat{C}$, which is based on $\tilde{Y}$, allows the principal to make the consumption process independent of the initial cash flow level $\hat{x}$. By insulating the agent's compensation process from the initial cash flow, the principal separates the agent's lying incentives from the initial cash flow. This separation works when the report process Y has a linear dynamic, as shown the next result.

Proposition 8 *Suppose that the contract C defined by $C_t = \mathcal{C}(Y_s : s \leq t)$ together with strategy L forms a contractual equilibrium starting from (v, x). Then, the contract $\hat{C}$ defined by $\hat{C}_t = \mathcal{C}(\tilde{Y}_s : s \leq t)$ together with strategy L forms a contractual equilibrium starting from $(v, \hat{x})$.*

Proposition 8 allows us define the first part of the transformation Φ: we set for any $\hat{x} \in \mathbb{R}$

$$\Phi_{(1,\hat{x}-x)}(C, L) = (\hat{C}, L).$$

The reason for the "1" subscript has to do with how the group $\mathbf{G}$ operating on the state space will be defined and will become clear in the next subsection.

We emphasize that this transformation applies for general utility functions of the agent, as was in the case in Proposition 2. Nothing in the proof of Proposition 8 requires that the agent have an exponential utility function. The key to providing Proposition 8 is that the cash flow has a linear dynamic equation (specifically, the drift of the cash flow is affine in its current value).

5.2 Transforming Contracts Across Promised-Utility Levels

From now on, we assume that the agent's utility function is given by $u(c) = -\exp(-\theta c)$ for some risk aversion parameter $\theta > 0$. In particular, the agent's flow utility $u(C_t - G_t)$ is always negative and, hence, so are his promised and continuation utilities at all times.

Consider a contractual equilibrium (C, L) starting from state (v_0, x), with $v_0 < 0$, and consider any alternative utility level $v_1 = \beta v_0$ for some $\beta \in (0, \infty)$. We define a new contract $\hat{C}$ as follows:

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

for all t. We have the following result:

PROPOSITION 9 *(C, L) is a contractual equilibrium starting from (v_0, x) if and only if $(\hat{C}, L)$ is a contractual equilibrium starting from (v_1, x).*

Proof. Let $V(C, L)$ denote the agent's expected utility when he follows strategy L given contract C. Since L is optimal for the agent, we have

$$V(C, L) = v_0 \geq v(C, L')$$

for all L'. For any $\hat{L}'$, let $L' = \hat{L}'$. Then, it is straightforward to check from (8) that

$$V(C, \hat{L}') = \beta V(C, L') \leq \beta V(C, L) = \beta v_0 = v_1$$

and the inequality is tight if $\hat{L}' \equiv L$, which shows that L is optimal given $\hat{C}$ and provides utility βv_0 to the agent . ■

For any contractual equilibrium (C, L) and $\beta > 0$, we set

$$\Phi_{(\beta,0)}(C, L) = (\hat{C}, L).$$

5.3 Action Group

Since the agent's flow utility is negative, the relevant state space for the agent is $\mathcal{S} = (-\infty, 0) \times \mathbb{R}$: it consists of all possible pairs (v, x) of expected utility and cash flows.

Let $\mathbf{G} = (0, +\infty) \times \mathbb{R}$ denote the algebraic group defined by the binary operation $(\beta, \delta)(\beta', \delta') = (\beta\beta', \delta + \delta')$ for all $(\beta, \delta), (\beta', \delta') \in \mathbf{G}$. It is straightforward to check that $\mathbf{G}$ is an Abelian group with identity element $(1, 0)$.

Moreover, the operation $(\beta, \delta) \circ (v, x) = (\beta v, x + \delta)$ defines a left action of $\mathbf{G}$ on the state space $\mathcal{S}$: the associativity and identity axioms are straightforward to check.

Finally, the group $\mathbf{G}$ defines a left action on the set of contractual equilibria, as follows:

1. Sections 5.1 and 5.2 define $\Phi_{(\beta,0)}(C,L)$ and $\Phi_{(1,\delta)}(C,L)$ for all $\beta > 0$, $\delta \in \mathbb{R}$ and contractual equilibrium (C,L).

2. For any $(\beta,\delta) \in \mathbf{G}$, we define $\Phi_{(\beta,\delta)}(C,L)$ by

$$\Phi_{(\beta,\delta)}(C,L) = \Phi_{(\beta,0)}(\Phi_{(1,\delta)}(C,L))$$

3. It is straightforward to check that this definition is consistent with 1. and that Φ satisfies Axioms 1,2,3 of Definition 2. Hence, Φ is a transformation group.[29]

4. Finally, it is also straightforward to check that Φ is monotone: if $\Pi(C,L) \geq \Pi(C',L')$ for any two contractual equilibria starting from the same state, then the constructions of Sections 5.1 and 5.2 preserve this payoff relation.

These observations are summarized by the following proposition.

Proposition 10 *The mapping $\Phi : \mathbf{G} \times \mathcal{E} \to \mathcal{E}$ is a monotone transformation group.*

6 Characterization of Truthful Renegotiation-Proof Contracts

This section derives a functional form for the payoff of the principal in any truthful state-consistent contract. It then introduces a mild strengthening of state consistency, in which the set of challengers is convexified. Renegotiation-proof contracts are those that sustain the comparison with this convexified set of challengers. The convexification, together with the strict concavity of the agent's utility function, implies that the contractual variables themselves, and not just the continuation payoff of the principal, are Markovian, and leads to a complete characterization of truthful, renegotiation-proof contracts. The class of truthful, renegotiation-proof contracts is parametrized by a single real number, which may be interpreted as the sensitivity of the agent's continuation utility to his reports.

Truthful Contracts and Promised Utility:

Given a truthful contract C, the filtrations $\mathbb{F}^X$ and $\mathbb{F}^Y$ are identical and the agent's continuation utility at time t is

$$W_t = E\left[\int_t^\infty e^{-r(\tau-t)}\left(u(C_\tau)\right)d\tau \,\middle|\, \mathcal{F}_t^Y\right].$$

[29] To prove associativity, one must show that consumption translation commutes with the virtual cash flow construction of Section 5.1, which is immediate from the construction.

The quantity W_t is the agent's *promised utility* at time t : it is the agent's expected discounted utility if he truthfully reports his cash flows. Since the promised utility process $W = \{W_t\}_{t\geq 0}$ is adapted to the filtration $\mathbb{F}^Y$, the Martingale Representation Theorem implies that W satisfies the dynamic equation[30]

$$dW_t = (rW_t - u(C_t))dt + \sigma S_t d\tilde{B}_t \tag{11}$$

where the process S_t is $\mathbb{F}^Y$-adapted and $\tilde{B}_t$ is an $\mathbb{F}^Y$-adapted Brownian motion under the probability measure in which the agent reports truthfully, i.e.,

$$d\tilde{B}_t = \frac{dY_t - (\xi - \lambda Y_t)dt}{\sigma}.$$

If the agent does *not* report his cash flow truthfully, $\tilde{B}$ is no longer a martingale: plugging Equation (7) into the previous equation yields

$$d\tilde{B}_t = \frac{\lambda(Y_t - X_t)dt + L_t dt + \sigma dB_t}{\sigma}. \tag{12}$$

Intuitively, the cash flow shocks reported by the agent's are biased upward if (i) the agent over-reports his cash flow ($L_t > 0$) or (ii) the actual cash flow is lower than the reported cash flow ($X_t < Y_t$). Effect (i) is straightforward. Effect (ii) comes from mean-reversion: if $X_t < Y_t$, then X_t has a higher mean than what the principal expects based on Y_t. Even if the agent does not produce any additional lie ($L_t = 0$), reported increments are positively biased.

It follows from (12) that when the agent lies, his promised utility evolves as

$$dW_t = (rW_t - u(C_t))dt + S_t(L_t dt + \lambda G_t dt + \sigma dB_t), \tag{13}$$

where $\{B_t\}_{t\geq 0}$ is the standard Brownian motion. The coefficient S_t is the *sensitivity* of the agent's promised utility to the agent's report increment at time t. It will be treated as a choice variable of the principal in the recursive formulation of the problem.

Agent's Incentives: Persistence of the agent's private information implies that past lies can have a long term impact on incentives: if the agent has lied even for a short period before time t, he has affected the report history $Y^t = \{Y_s\}_{s\leq t}$ and therefore also affected his future consumption flow C and his future incentives to report the truth. Intuitively, if the agent *underreports* his cash flow increment dX_t at time t, he affects his utility through two channels. First, this *reduces his promised utility*, which depends on the report dY_t by a sensitivity factor S_t, and thus reduces the consumption stream that the principal aims to give the agent. Second, to deliver a *given level of promised utility*, the principal *must provide higher transfers to the agent* if the agent generates lower cash flows. The first channel incentivizes the agent to report higher cash flows and the second channel incentivizes

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

him to report lower cash flows. For the contract to be truthful, these two incentives must balance each other, at least on the equilibrium path. The agent's reporting incentives are the subject of Section 9.

6.1 Structure of the Principal's Payoff

Consider a truthful, state-consistent contract C for the transformation group Φ constructed in Section 5. Since the contract is truthful, there is no difference on the equilibrium path between X_t and Y_t, or between V_t and W_t.

We will use (w, y) instead of (v, x) to denote the state, to make these equalities clear. For any state $(w, y) \in (-\infty, 0) \times \mathbb{R}$, let $\Pi(w, y)$ denote the principal's expected payoff for the transformation of $(C, 0)$ that corresponds to state (w, y). From Proposition 5, $\Pi(w, y)$ is also the principal's continuation payoff after any history leading to state (w, y). We now derive the functional form of $\Pi(w, y)$.

First, consider the dependence on the cash flow level y. As noted in Section 5.1, the only difference in continuation payoffs for the principal between histories starting from state (w, y) and (w, y') concerns the expected discounted transfer from the agent $\Upsilon(y) = E[\int_0^\infty e^{-rt} Y_t dt | w, y]$, because the processes C have the same distribution independently of y.

To compute $\Upsilon(y)$, note that since the contract is truthful, we have $Y_t \equiv X_t$. Moreover, the cash flow process X_t has an explicit formula, which is computed in the Appendix (Equation (50)), the expected discounted transfer from the agent satisfies

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

After simplification, this yields

$$\Upsilon(y) = \frac{y}{r + \lambda} + \frac{\xi}{r(r + \lambda)}. \tag{14}$$

This yields the following result:

$$\Pi(w, y) = \Pi(w, 0) + \frac{y}{r + \lambda}. \tag{15}$$

Next, consider two different promised utility levels w_0 and $w_1 = \beta w_0$ for the agent, where $\beta > 0$. The translation performed in Section 5.2 implies that

$$\Pi(\beta w_0, y) = \frac{\log(\beta)}{\theta r} + \Pi(w_0, y).$$

The previous analysis yields the following result.

Proposition 11 *For any truthful, state consistent contract C, the principal's payoff satisfies the following relation:*

$$\Pi(w', y') = \frac{(y' - y)}{r + \lambda} + \frac{\log(w'/w)}{r\theta} + \Pi(w, y).$$

In particular, the principal's continuation payoff in state (w, y) satisfies

$$\Pi(w, y) = \frac{y}{r + \lambda} + \frac{\log(-w)}{\theta r} + \Pi(-1, 0). \tag{16}$$

Proposition 11 has an immediate but important corollary:

Corollary 2 *If C is truthful and state consistent, the principal's continuation payoff takes the form $\Pi_t = \Pi(W_t, Y_t)$ where $\Pi(\cdot, \cdot)$ is twice continuously differentiable.*

6.2 Convexification and Renegotiation-Proof Contracts

Proposition 11 established that the continuation payoff of the principal in any truthful, state-consistent contract is Markov in the state variables (W_t, Y_t). In principle, the contract itself need not be Markov: there could be multiple continuation contracts starting from a given state (w, y) that give the principal the same payoff, and all of these continuations could in principle be used after various histories.

In order to guarantee the contract is Markov, we impose a mild strengthening of state consistency, in which the set of challengers to a given continuation is the convex hull of the relevant set of contract transformations.

When the agent's utility is concave and a contractual equilibrium has multiple continuations giving the agent the same expected utility, it is natural for the principal to consider convexifications of these continuations, as they may be cheaper.

Specifically, consider two truthful contracts C^1, C^2 starting from the same cash flow x and giving the same utility v to the agent. For any $\lambda \in [0, 1]$, consider the contract C^λ defined as follows: for any t and report history $\{Y_s : s \leq t\}$,

$$u(C_t^\lambda) = \lambda u(C_t^1) + (1 - \lambda) u(C_t^2). \tag{17}$$

The consumption level C_t^λ is uniquely defined if u is strictly increasing and continuous, which is the case for exponential utility functions. In words, for any t and report history $\{Y_s\}_{s \leq t}$, C_t^λ provides a flow utility to the agent that is a fixed convex combination of the flow utilities that he gets under C_t^1 and C_t^2. If C^λ is truthful, this implies that it provides expected utility v to the agent because (i) C^1 and C^2 both give expected utility v to the agent, and (ii) all three contracts generate the same distribution for the report process Y. The next result shows that C^λ is indeed truthful and formalizes this result.

Proposition 12 *Suppose that $u(c) = -\exp(-\theta c)$ for some parameter $\theta > 0$. For any $\lambda \in [0,1]$ and initial cash flow x, the contract C^λ is truthful and gives the agent expected utility v.*

Given a truthful contract C and state (v,x), let $\mathcal{K}(C)$ denote the set of all continuation contracts of C, and let

$$\mathcal{K}(v,x) = \left\{\Phi_g(\tilde{C}) : (g,\tilde{C}) \in \mathbf{G} \times \mathcal{K}(C) \quad \text{s.t. } g \circ (v_{\tilde{C}}, x_{\tilde{C}}) = (v,x)\right\}$$

where $(v_{\tilde{C}}, x_{\tilde{C}})$ denotes the starting state of contract $\tilde{C}$. In words, $\mathcal{K}(v,x)$ consists of all contracts starting from state (v,x) that are transformations of continuation contracts of C.

Finally, let

$$\text{Conv}(\mathcal{K}(v,x)) = \{C^\lambda : \lambda \in [0,1], C^1, C^2 \in \mathcal{K}(v,x)\}$$

denote the set of convex combinations of contracts in $\mathcal{K}(v,x)$.

The set $\text{Conv}(\mathcal{K}(v,x))$ defines the class of all challengers that the principal can consider, after any history leading up to state (v,x), to replace the current continuation contract.

Definition 6 *A truthful contract C is* renegotiation-proof *if (i) the pair $(C, L \equiv 0)$ forms a state-consistent contractual equilibrium and (ii) for any history leading to state (v,x), the continuation contract $\tilde{C}$ satisfies*

$$\Pi(\tilde{C}) \geq \Pi(C') \tag{18}$$

for all $C' \in \text{Conv}(\mathcal{K}(v,x))$.

This definition captures the following intuition: if the principal can come up with some challengers of a contract, possibly by comparing contract continuations across states, then she can also consider the convex combination of these challengers. Convexification is used in Lemma 1 to pin down contractual variables as a function of the state (Proposition 13).

6.3 Contractual Variables

We now derive the functional form of contractual variables C_t, S_t for all renegotiation-proof contracts.

Proposition 13 *For any truthful, renegotiation-proof contract, there exist constants $c_1, \bar{s}$ such that*

$$C_t = c_1 - \frac{\log(-W_t)}{\theta} \quad a.s. \tag{19}$$

$$S_t = -W_t \bar{s} \quad a.s. \tag{20}$$

for all $t \geq 0$.

The parameters c_1 and $\bar{s}$ are the consumption and sensitivity levels provided at time 0 by the transformation of the contract that starts with promised utility -1 (any cash flow level).

From Proposition 13, the parameters c_1 and $\bar{s}$ fully characterize contract C. We will see in Section 9 that a contract to be truthful, c_1 and $\bar{s}$ must satisfy an additional equation, which will further reduce the set of renegotiation-proof contracts to a single parameter family.

Proof. From (11), the statement of Proposition 13 is equivalent to W being a geometric Brownian motion or, equivalently, to the Itô process Z given by $Z_t = \log(-W_t)$ having constant drift and volatility.

To prove that the drift and volatility of Z are constant, we will use an intermediate result. Let (w, y) denote the initial state of the truthful renegotiation-proof contract C. Fix a time T and history up to time T and let $\tilde{C}$ and $(\tilde{w}, \tilde{y})$ denote the continuation contract and state following this history. Let $\hat{C}$ denote the transformation of $\tilde{C}$ that starts in state (w, y): letting g denote the element of $\mathbf{G}$ such that $(\tilde{w}, \tilde{y}) = g \circ (w, y)$ (which is always possible since $\mathbf{G}$ is transitive) we let $\hat{C} = \Phi_{g^{-1}}(\tilde{C})$. The contract $\tilde{C}$ is truthful as a continuation of C, which is truthful. Therefore, $\hat{C}$ is also truthful because Φ preserves truthfulness (Propositions 8 and 9). The contract $\hat{C}$ may be proposed by the principal at time 0, instead of contract C, as it starts from the same state (w, y). We now show that the stochastic processes $C = (C_t)_{t \geq 0}$ and $\hat{C} = (\hat{C}_t)_{t \geq 0}$ must be identical functions of the realized cash-flow process $Y = (Y_s)_{s \geq 0}$. This result, which is proved in Appendix B, is formalized as follows. Let $\mathbb{L}$ denote the vector space of processes Z such that $E[\int_0^\infty e^{-rt}|Z_t|dt] < \infty$ endowed with the norm $\|Z\|_{\mathbb{L}} = E[\int_0^\infty e^{-rt}|Z_t|dt]$. By assumption, any contract C belongs to $\mathbb{L}$ under truth-telling.

Lemma 1 *The processes C and $\hat{C}$ are identical in $\mathbb{L}$: $\|C - \hat{C}\|_{\mathbb{L}} = 0$.*

Let $\{\hat{W}_t\}_{t \geq 0}$ denote the agent's promised utility under contract $\hat{C}$. $\hat{W}$ is an $\mathbb{F}^Y$-adapted Itô process given by

$$\hat{W}_t = E\left[\int_t^\infty e^{-r(s-t)} u(\hat{C}_s) ds \quad \middle| Y_z : z \leq t\right]. \tag{21}$$

Lemma 1 implies that $\hat{W}$ is identical to W and that the process $\hat{Z} : \{\hat{Z}_t = \log(-\hat{W}_t)\}_{t \geq 0}$ is identical to the process $Z : \{Z_t = \log(-W_t)\}_{t \geq 0}$.

Let μ_t and σ_t denote the drift and volatility of Z at time t. Our objective is to show that μ_t and σ_t are constant. Since $\hat{Z}$ and Z are identical, we have:

$$d\hat{Z}_t = dZ_t = \mu_t dt + \sigma_t dB_t.$$

By definition of Φ, the contract $\hat{C} = \Phi_{g^{-1}}(\tilde{C})$ is obtained by translating the process $\tilde{C} = \{C_{T+t}\}_{t\geq 0}$ by the constant $\frac{1}{\theta}(\log(-W_T) - \log(-w))$ (setting $\beta = \log(-W_t/-w)$ in Section 5.2) and making it independent of the initial condition Y_T, as explained in Section 5.1. The latter condition is equivalent to the requirement that $\hat{C}$ depends on the Brownian path $\{B_t\}_{t\geq 0}$ in the same way that $\tilde{C}$ depends on the Brownian path $\{B_{T+t} - B_T\}_{t\geq 0}$.[31]

From (21) and the fact that $u(c) = -\exp(-\theta c)$, this implies that the process $\{\hat{Z}_t\}_{t\geq 0}$ is identical in law to the process $\bar{Z} = \{Z_{T+t} - Z_T + Z_0\}_{t\geq 0}$ conditioned on $\mathcal{F}_T^Y$. The latter process has drift $\{\mu_{T+t}\}_{t\geq 0}$ and volatility $\{\sigma_{T+t}\}_{t\geq 0}$. Equality in law of $\hat{Z}$ and $\bar{Z}$ implies that the drift and volatility of $\hat{Z}$ at time zero is equal to the drift and volatility of $\bar{Z}$ at time $t = 0$, i.e.,

$$\mu_0 = \mu_T \quad \sigma_0 = \sigma_T.$$

Since the time T and history up to time T were arbitrary, this shows that μ_T and σ_T are constant and proves the proposition.

■

Proposition 13 immediately implies the following corollary:

Corollary 3 *Any truthful, renegotiation-proof contract is Markovian: there exist functions $c(\cdot)$ and $s(\cdot)$ such that $C_t = c(W_t, Y_t)$ and $S_t = s(W_t, Y_t)$ for all t.*

In fact, Proposition 13 shows a stronger result: the functions $c(\cdot,\cdot)$ and $s(\cdot,\cdot)$ are independent of Y_t: the agent's consumption at any time depends only on his current promised utility, and so does the sensitivity S_t.[32]

Proposition 13 provides a necessary condition for the form of truthful renegotiation-proof contracts. This necessary condition is strong, as it reduces the family of such contracts to a two-parameter family with parameters c_1 and $\bar{s}$. This reduction to two parameters is essentially due to state consistency and convexification.[33]

[31] With a truthful contract starting from cash flow y, the reported cash flow process Y has an explicit formula, given by $Y_t = e^{-\lambda t}y + \int_0^t e^{\lambda(s-t)}\xi ds + \int_0^t e^{\lambda(s-t)}\sigma dB_s$ (see Equation (50)). Therefore, the path of Y is entirely pinned down by the initial condition y and the path of B. Independence of the contract with respect to the initial condition y then shows the claim. Remark: although easier to understand with the closed form of Y, the claim holds as long as the cash flow process Y solves an SDE with noise B that has a unique solution (see, e.g., Cherny (2002, Theorem 3.2)).

[32] For readers familiar with the concept of symmetry in physics and Noether's theorem, this property may be viewed as a *conservation law* of the contract that stems from state consistency: state consistency with respect to Y_t may be viewed as a *symmetry* of the contract expressed in terms of the left group action $C \mapsto \Phi_{(1,\delta y)}(C)$ and invariance of the contract with respect to Y_t is a consequence of this symmetry.

[33] From an algebraic perspective, each truthful Markovian contract generates an orbit in the set $\mathcal{T}$ of all truthful

In fact, the requirement that truthtelling be incentive compatible for the agent imposes a relationship between the variables c_1 and $\bar{s}$, as follows. Let $u_1 = u(c_1)$. The following result is shown in Section 9.

THEOREM 1 *The set of truthful renegotiation-proof contracts is characterized by Equations* (11), (19), *and* (20), *and* u_1 *and* $\bar{s} \in (0, \theta)$ *satisfy the following relation:*

$$u_1 = -\frac{\bar{s}\lambda}{\theta - \bar{s}}. \tag{22}$$

Theorem 1 shows that each truthful renegotiation-proof contract can be reduced to a single parameter, $\bar{s}$. To optimize the principal's payoff with respect to $\bar{s}$, the next proposition computes the payoff explicitly.

PROPOSITION 14 *The principal's payoff is given by:*

$$\frac{1}{\theta r}\left[\log(-u_1(\bar{s})) + \frac{1}{r}\left(r + u_1(\bar{s}) - \frac{1}{2}\sigma^2\bar{s}^2\right) + D(w, y)\right] \tag{23}$$

where $u_1(\bar{s}) = -\frac{\bar{s}\lambda}{\theta - \bar{s}}$ *and*

$$D(w, y) = \log(-w) + \frac{\theta}{r + \lambda}(ry + \xi).$$

7 Comparative Statics

The following comparative statics obtain by maximizing (23) with respect to $\bar{s}$.

7.1 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 15 *The optimal sensitivity* $\bar{s}$ *is decreasing in* λ*.*

Intuitively, when the cash flow is more persistent (λ is *lower*), the agent's current report has more bearing on expected future cash flows. This raises the magnitude of the agency problem and requires a higher a sensitivity to induce truthful reporting.

Markovian contracts through the left action of group Φ. Proposition 13 says that the quotient of $\mathcal{T}$ with respect to these orbits (or equivalence classes) is two-dimensional.

The proof of Proposition 15 is straightforward. Ignoring terms and factors in the principal's objective (23) that do not influence the choice of $\bar{s}$, the optimal choice of $\bar{s}$ is the solution to

$$\log(-u_1(\bar{s})) + \frac{1}{r}\left(u_1(\bar{s}) - \frac{1}{2}\sigma^2\bar{s}^2\right), \tag{24}$$

where

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

The function $\bar{s} \mapsto u_1(\bar{s})$ is submodular in $(\bar{s}, \lambda)$, as is easily checked. Therefore, the second term of (24) 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 (32):

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

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

7.2 Impact of Volatility and Discount Rate

It is easy to check from (23) that the principal's payoff is decreasing in the volatility of the agent's cash flow. The next result shows that the optimal sensitivity coefficient is also decreasing in cash flow volatility.

PROPOSITION 16 *The optimal sensitivity $\bar{s}$ is decreasing in σ.*

Proof. Expression (23) is submodular in σ and $\bar{s}$. The result follows from Topkis (1978).

This result is intuitive: A higher volatility exposes the agent to more risk, other things equal. Since the agent is risk averse, it is optimal for the principal to offset this by reducing the agent's exposure to exogenous cash flow fluctuations, by reducing $\bar{s}$. Perhaps more surprisingly, the sensitivity is decreasing in the patience of the players, as shown in the Appendix.

PROPOSITION 17 *The optimal sensitivity $\bar{s}$ is increasing in r.*

A possible intuition is that a less patient agent is more concerned about getting an immediate subsidy (or tax reduction) than about his future utility. Inducing truth-telling thus requires a higher sensitivity of the agent's future utility to his current report.

7.3 Long-Run Properties

Bloedel, Krishna, and Strulovici (2020) show that the agent's utility always has a negative drift under the optimal contract, but that the long-run properties of consumption and promised utility depend on the parameters of the agent's cash flow process. Fixing the volatility parameter, the optimal contract induces immiserisation when persistence is low, but sends the agent to bliss when persistence is high. Fixing the degree of persistence, the optimal contract induces immiserisation when volatility is high, but sends the agent to bliss when volatility is low.

8 Observability Theorem and Revelation Principle

Reconstructing the Agent's True Cash Flow: An Observability Theorem

We focused up to now on truthful equilibria. Consider instead a contractual equilibrium in which the agent lies on path about his cash flow. We will see that, given an equilibrium strategy L that is $\mathbb{F}^X$ adapted, the principal can under mild regularity conditions on L reconstruct the true cash flow X_t and the agent's true continuation utility V_t from the agent's report process Y. To this end, notice that the agent's lying process is $\mathbb{F}^X$-adapted and may thus be expressed as a functional

$$L_t = \mathcal{L}(t, X_s : s \le t), \tag{25}$$

which depends, at each time t on the path of X until time t. Given the functionals $\{\mathcal{L}(t,\cdot)\}_{t\ge 0}$—which the principal knows in equilibrium—and the report process Y, we can construct a process $\tilde{X}_t$ defined by the equation

$$d\tilde{X}_t = dY_t - \mathcal{L}(t, \tilde{X}_s : s \le t)dt \tag{26}$$

and the initial condition $\tilde{X}_0 = Y_0 = X_0$. Equation (26) is a functional stochastic differential equation in which the unknown is the process $\tilde{X}$ and the primitive is the process Y. From (7), the true cash flow process X is a weak solution of (26). Moreover, any strong solution of (26) is, by definition, adapted to Y. If (26) has a unique solution, the principal can back out the process X by observing Y and using her knowledge of the agent's equilibrium strategy. This intuition is formalized by Theorem 2, which is proved in Appendix C.[34]

Theorem 2 (Observability Theorem) *Consider a contractual equilibrium (C, L) such that L is locally bounded: for each $T > 0$, there exists $M(T) > 0$ such that $\sup_{t\le T} |L_t| \le M(T)$ a.s. Then, the following holds:*

[34]While the model's horizon is infinite, Theorem 2 provides conditions under which, for each $T > 0$, the agent's reports up to time T reveal the agent's continuation value and the current-cash flow at time T. Thus, to prove Theorem 2, it suffices to show that it holds for any finite horizon $[0, T]$, where $T > 0$ is arbitrary.

1. *If Equation* (26) *has a strong solution, then this solution is unique, and* X *is* $\mathbb{F}^Y$*-adapted.*

2. *If* L *is locally Lipschitz continuous,*[35] *then Equation* (26) *has a unique strong solution and* X *is* $\mathbb{F}^Y$*-adapted.*

3. *If there exists some* $\varepsilon > 0$ *such that for each* $t > 0$*, the functional* $\mathcal{L}(t, \cdot)$ *is independent of* $\{X_s\}_{s\in[t-\varepsilon,t]}$*, then Equation* (26) *has a unique strong solution and* X *is* $\mathbb{F}^Y$*-adapted.*

Conditions given in 2. and 3. both include truthtelling as a feasible strategy. In particular, the truthful strategy $L \equiv 0$ is constant and thus clearly Lipschitz. The lag ε appearing in 3. may be viewed as small behavioral constraint or frictions which requires that agent's response time to new information is no exactly zero, but bounded below by some arbitrarily small constant.[36]

To interpret Theorem 2, notice that the principal's belief about the agent's cash flow may be quite different from the report made by the agent. Given some contract C, the agent can affect his transfers by overreporting or underreporting his cash flows: these reports mechanically feed into the transfer specified by C, regardless of their effect on the principal's belief. Theorem 2 says that the principal's belief process matches the true cash flow as long as the principal's *belief equation* (26) has a strong solution (Part 1), which is guaranteed if the agent's strategy exhibits some arbitrarily small lag (Part 3) or if it is Lipschitz for some arbitrary Lipschitz constant (Part 2).

This distinction is conceptually important in the context of renegotiation: it means that in equilibrium the principal knows the agent's continuation utility under the current contract when renegotiating with the agent, even if the agent has fed inaccurate reports into the contract C. Thus, it is legitimate to treat the actual cash flow, X_t, and agent continuation value V_t as state variables for the purpose of analyzing renegotiation, as long as the agent's strategy satisfies the assumptions of Theorem 2, which we impose as part of the admissibility requirement for L.

Focus on Adapted Strategies

The assumption that the agent's lying process is adapted to $\mathbb{F}^X$ plays a key role in Theorem 2. While the focus on adapted strategies is quasi-universal in economic models with diffusion processes, this focus is generally not innocuous. In particular, it rules out mixed strategies and it is a priori *with* loss of generality. In diffusion models, there have been very few attempts to model mixed

[35]This means that for each $T > 0$, there exists $\bar{L}(T) > 0$ such that all functionals $\{\mathcal{L}(t, \cdot)\}_{t\le T}$ are $\bar{L}(T)$-Lipschitz continuous over their respective domain.

[36]The truthful renegotiation-proof contracts characterized by Theorem 1 are truthful even when the agent's strategy is allowed to violate both conditions 2. and 3. of Theorem 2, since none of these conditions is imposed in the verification that truthful behavior is optimal in Section 9.2. Imposing condition 2. or 3. is useful only to guarantee that the focus on truthful contracts is without loss of generality, by Theorem 3.

strategies in general as this raises difficult technical problems.[37] Focusing on adapted strategies is a priori restrictive, but for diffusion models there are compelling reasons for this focus, that are independent of any concern for renegotiation or private information. For example, the Martingale Representation Theorem, which is used by DeMarzo and Sannikov (2006), Sannikov (2008) and numerous subsequent papers, relies on this assumption. Williams (2011), on which the main application of the present paper is based, assumes that agent's reporting strategy is adapted to the observed process (cf. last paragraph of p. 1239 in that paper).

There are separate reasons to think that mixing would be of little help in the present setting. In continuous-time models, the principal has infinitely many opportunities to extract information from the agent, and one may expect any uncertainty caused by mixing would be quickly resolved as the principal would renegotiate contracts with arbitrarily high frequency. For example, Maestri (2017) and Strulovici (2017) study related protocols in discrete time and found that as the frequency of renegotiation becomes arbitrarily large, the private information of the agent is revealed arbitrarily quickly. This justification lies outside the scope of the present paper as there is no explicit protocol of negotiation in the present paper.

8.1 Revelation Principle

From Theorem 2, the principal knows the agent's true output and continuation value even in contractual equilibria that are not truthful. Therefore, it is legitimate for the principal to consider challengers based on the true continuation state (V_t, X_t), even if the contractual equilibrium is not truthful.

As noted in Section 4, the transformation group and the concept of state consistency defined in that section applies equally well to contractual equilibria as it does to truthful contractual equilibria, and the comparisons are legitimate due to Theorem 2.

Since the concept of a state-consistent contractual equilibrium is well defined, we can ask whether such equilibrium is outcome equivalent to a *truthful* contractual equilibrium, i.e., whether a version of the Revelation Principle holds in this setting.

In the absence of commitment, it is well known that the Revelation Principle need not hold.[38] However, a version of this principle holds in our setting, as described next.

Say that two contractual equilibria starting from the same state are *outcome equivalent* if they induce the same net transfer after any history $\{X_s\}_{s\leq t}$ of the real output.

[37] One such attempt is due to Bernard (2016)'s PhD thesis, which focuses on this question.

[38] See, e.g., Bester and Strausz (2001).

Theorem 3 *Suppose that Φ is monotone. Then, any state-consistent contractual equilibrium is outcome-equivalent to a truthful state-consistent contractual equilibrium.*

Proof. Consider any contractual equilibrium (C, L). As with the standard Revelation Principle, we can replace the contract C by a contract $\hat{C}$ in which the agent truthfully reports his cash flow, and the principal simulates the lies that the agent would have made in the initial equilibrium and makes the same net transfer to the agent that she would have made in the initial equilibrium. By construction, the truthful contractual equilibrium $(\hat{C}, \hat{L} \equiv 0)$ is outcome-equivalent to the initial contractual equilibrium.

We must show that if (C, L) is state consistent with respect to Φ, then so is $(\hat{C}, 0)$. Let (v, x) denote the initial state (which is the same for both contractual equilibria) and consider any real output history $\{X_s\}_{s \leq t}$ leading to some state $(v', x') = g \circ (v, x)$. Let $(\hat{C}', 0)$ denote the continuation equilibrium of $(\hat{C}, 0)$ following this history and (C', L') denote the continuation equilibrium of (C, L) following this corresponding history.

By outcome equivalence, we have

$$\Pi(\hat{C}', 0) = \Pi(C', L') \tag{27}$$

and

$$\Pi(C, L) = \Pi(\hat{C}, 0).$$

Since Φ is monotone, the previous equality implies that

$$\Pi(\Phi_g(C, L)) = \Pi(\Phi_g(\hat{C}, 0)). \tag{28}$$

Comparing (27) and (28) and using the fact that (C, L) is state consistent then shows that $(\hat{C}', 0)$ also satisfies the definition of state consistency (Definition 4). ■

9 Necessary and Sufficient Conditions for Incentive Compatibility

This section proves Theorem 1: it shows that (22) is necessary and sufficient for the contracts of Proposition 13 to be incentive compatible (i.e., truthful). The argument used to prove sufficiency illustrates a strategy, potentially useful in other problems, to deal with an unbounded reporting domain by expending the strategy space of the agent.

9.1 Necessity

Consider any contract that has the form of Proposition 13, which is characterized by parameters $c_1 \in \mathbb{R}$ and $\bar{s} \geq 0$, such that $C_t = c(W_t)$ and $S_t = s(W_t)$ for all t, where the functions $c(\cdot)$ and $s(\cdot)$

are given by $c(w) = c_1 - \log(-w)/\theta$ and $s(w) = \bar{s}(-w)$. The promised utility of the agent evolves as

$$dW_t = (r + u_1)W_t dt + s(W_t)(dG_t + \lambda G_t dt + \sigma dB_t) \tag{29}$$

where $G_t = Y_t - X_t$ where $u_1 = u(c_1)$. Moreover, the agent's actual consumption is $c(W_t) - G_t$. Therefore, the agent cares about Y_t and X_t only through their difference G_t. The agent's optimization problem reduces to

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

subject to $dG_t = L_t dt$, $G_0 = g$, and (29). The HJB equation for this problem is

$$\begin{aligned} 0 = \sup_l \Big\{ & u(c(w) - g) - rv(w, g) \\ & + v_w\left(rw - u(c(w)) + s(w)(l + \lambda g)\right) + v_g l + \frac{1}{2}(s(w))^2\sigma^2 v_{ww}\Big\}. \end{aligned} \tag{30}$$

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 $f(0) = 1$, which means that the agent cannot do better than get his promised utility w. The Bellman equation becomes, after simplification and dividing throughout by $(-w)$,

$$0 = \sup_\ell \left\{u_1 \exp(\theta g) + f(g)(-u_1 + \bar{s}(\ell + \lambda g)) - f'(g)\ell\right\}. \tag{31}$$

A priori, the function f defining the agent's value function need not be everywhere differentiable, but it is a viscosity supersolution of (31) (see, e.g., Øksendal and Sulem (2004)), which means in this context that the supremum in (31) must be less than or equal to 0. 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, for $g = 0$, that it is optimal for the agent not to increase the gap above zero, which requires that

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

and it is optimal not to decrease it below zero, which requires that

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

Therefore, incentive compatibility implies that $f_l(0) \leq \bar{s} \leq f_r(0)$.

We now prove that these inequalities are tight. Suppose by contradiction 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. We approximate the agent's optimal strategy by considering an arbitrarily negative lying rate $-K$ over some small time window, and taking the limit as $K \to +\infty$. Specifically, suppose that at time t the state is (w, g) for some fixed $g \in (0, \eta)$, so that the agent expected utility at time t is $v(w, g) = wf(g)$, and suppose that the agent lies at an arbitrarily large rate $-K$ between times t and $t_K = t + g/K$, which brings the gap to 0 at time t_K. The agent's promised utility over this time interval obeys the dynamic equation[39]

$$dW_t = W_t(r + u_1 + K\bar{s})dt + \bar{s}(-W_t)\sigma dB_t,$$

which integrates to $W_{t_K} = \hat{W}_{t_K} \exp(\bar{s}K)$, where $\hat{W}_{t_K}$ is the agent's promised utility at time t_K if he doesn't lie during this period. The agent's expected utility at time t_K is therefore $v(W_{t_K}, G_{t_K}) = \hat{W}_{t_K} \exp(\bar{s}g)f(0)$ since we chose t_K so that $G_{t_K} = 0$. As $K \to +\infty$, we have $t_K \to t$, $\hat{W}_{t_K} \to w$, and $v(W_{t_K}, G_{t_K}) \leq v(w, g)$ since lying at the arbitrarily large rate $-K$ is only one of many possible strategies for the agent. Moreover, incentive compatibility requires that $f(0) = 1$. Taken together, these observations yield $wf(g) \geq w \exp(\bar{s}g)$. Dividing by $w < 0$ and noting that the argument works for all $g \in (0, \eta)$, we get

$$f(g) \leq \exp(\bar{s}g)$$

for all $g \in (0, \eta)$. Since the functions on both sides of the equation are equal to 1 for $g = 0$, we get

$$f'_r(0) \leq \left.\frac{d[\exp(\bar{s}g)]}{dg}\right|_{g=0} = \bar{s},$$

which yields the desired contradiction. By a similar argument, $f_l(0) = \bar{s}$.

In conclusion, f is differentiable at 0, with derivative $\bar{s}$. This implies that the value function $\bar{v}(w, x, y) = v(w, y - x)$ of the agent is differentiable with respect to x and to y whenever $x = y$. In Appendix D (Lemma 2), this differentiability is used in an envelope argument to show that, necessarily,

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

The right-hand side is computed explicitly in Appendix D (Lemma 3), yielding

$$\bar{s} = \frac{\theta(-u_1)}{\lambda - u_1} \in (0, \theta). \tag{32}$$

Notice that $\bar{s} < \theta$ for all $u_1 < 0$ and $\lambda > 0$. 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 effect of earning less than the reported cash flow would reduce the agent's flow utility at rate θ, which would be dominated by the increase in promised utility governed by the sensitivity parameter $\bar{s}$. Incentive compatibility rules this case out.

[39]The equation is exact up to a second-order term, which becomes negligible as $K \to +\infty$.

9.2 Sufficiency

We have established that (22) is necessary for incentive compatibility. We now show that it is sufficient. The objective (31) 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 account for this, we now allow the agent to report jumps in his cash flows, which expands the agent's strategy space.[40]

Let $W_{t_+}(\Delta L)$ denote the agent's promised utility if he reports a jump ΔL in his cash flow at time t. For contracts with a fixed sensitivity parameter, as considered here, a natural closure of the contract to report jumps is to stipulate that

$$W_{t_+}(\Delta L) = \exp(-\bar{s}\Delta L)W_t. \tag{33}$$

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 (see the analysis of the necessary condition in Section 9.1 for a formal argument), the dynamic equation $dW_t = \bar{s}(-W_t)Kdt$. 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 (33).

Report jumps amount to impulse controls on the part of the agent (see for example, Øksendal and Sulem (2004)). The HJB equation (31) becomes

$$0 = \max\left\{\sup_{\ell\in\mathbb{R}}\left\{u_1\exp(\theta g) + f(g)(-u_1 + \bar{s}(\ell+\lambda g)) - f'(g)\ell\right\},\right.$$
$$\left.\sup_{\Delta L\in\mathbb{R}\setminus\{0\}}\left\{\exp(-\bar{s}\Delta L)f(g+\Delta L) - f(g)\right\}\right\}. \tag{34}$$

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

$$\sup_{\ell}\left\{u_1\exp(\theta g) + \exp(\bar{s}g)\left(-u_1 + \bar{s}\left(\ell+\lambda g\right)\right) - \bar{s}\exp(\bar{s}g)\ell\right\} \le 0 \tag{35}$$

for all g. The terms involving ℓ cancel out and (35) reduces to

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

Convexity of the exponential function implies that for all $g \neq 0$,

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

[40] Admissibility is unchanged: it requires that the transversality condition (9) holds. This condition is used to check in a standard verification argument that the solution of the HJB equation must be equal to the value function.

Since $u_1 < 0$, this implies that (35) is satisfied if

$$\exp(\bar{s}g)\,(u_1(\theta - \bar{s})g + \bar{s}\lambda g) \leq 0.$$

The second factor is zero, from (22), 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 therefore 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 holds only if $g = 0$.

10 Discussion

10.1 Summary: Concept of Renegotiation

This paper introduces a concept of renegotiation that strengthens internal consistency by exploiting comparisons across states. These comparisons are weaker than those imposed by strongly renegotiation-proof equilibria (Farrell and Maskin (1989)) and could be applied to study stable "social norms" (Asheim (1991)), with the following cognitive interpretation: players may recognize relations between equilibria across different states, just as they recognize, in a repeated game, whether continuation payoffs across different histories are Pareto ranked. This requires players to think by analogy rather than come up with radically new equilibria. From a modeling perspective, economic analysis often relies on specific payoff or dynamic structures (e.g., CARA or CRRA preferences, linear or geometric growth). For such models, the present approach provides a natural way to model renegotiation.

The concept is well-defined in any of the following situations: (i) there is no private information (time could be discrete or continuous), (ii) there is private information but we focus on truthful contracts, or (iii) there is private information, and the setting is a diffusion model and strategies are adapted to the underlying Brownian process, as in the main application of this paper. In this last setting, the paper provides an Observability Theorem and a Revelation Principle that address informational asymmetries.

10.2 Robustness and Extensions

State Consistency in Games: The concept of state consistency developed in this paper has focused on bilateral contractual relationships. State consistency could in principle be extended to dynamic games with an infinite horizon. The idea is as follows. Consider a dynamic game with n agents, whose environment at time t is described by some state variable $X_t \in \mathcal{X}$. Suppose that given (i) any subgame-perfect equilibrium (SPE) E starting from state $X_0 = x$ and (ii) any other state $x' \neq x$ in $\mathcal{X}$, it is possible to construct another SPE $E' = \Phi_{x,x'}(E)$ starting from x'. Let $V' = (V_1', \ldots, V_n')$ denote the vector of payoffs of all players from equilibrium E'. Suppose that, under the original SPE E, the game reaches a history at which the state is x', and the continuation equilibrium is $\bar{E}'$. Let $\bar{V}' = (\bar{V}_1', \ldots, \bar{V}_n')$ denote players' continuation payoffs following this history. Then, if $V' > V$, players could leave $\bar{E}'$ and agree to move to E' instead. State consistency would generalize internal consistency by allowing comparisons in such dynamic games.

One may envision many such games. For example, the stage game could entail players taking actions that determine how they split a pie, whose size evolves dynamically, either as an exogenous Markov process, or a controlled process that is affected by players' actions in the last period. The state variable X_t would then specify the size of the pie at time t, and one could imagine scenarios in which any continuation equilibrium starting from state x could be scaled into a different continuation equilibrium starting from another pie size $x' \neq x$.

This extension of state consistency to dynamic games, possibly including privately observed states is a major direction to develop the ideas put forth in this paper.

Effort to Improve Real Cash Flow: The agent could affect the real cash flow by putting some privately observed effort. This extension is realistic in many settings (e.g., to study labor incentives and taxation) and connects the model with numerous paper on dynamic moral hazard. An earlier version of this paper (Strulovici (2011)) studies this extension: the agent produces privately observed effort to affect his cash flow: $dX_t = (\xi - \lambda X_t + A_t)dt + \sigma dB_t$, where $A = \{A_t\}_{t \geq 0}$ is the agent's effort process, adapted to $\mathbb{F}^X$, and the agent incurs an additively separable cost $\phi(A_t)$ at time t from effort A_t. Although they are more complicated, all the arguments and results of this paper continue to go through:[41] assuming that ϕ takes the exponential form $\phi(a) = \bar{\phi} \exp(\chi a)$ for some positive parameters $\bar{\phi}, \chi$, the definition of state consistency and the form of renegotiation-proof contacts (consumption and sensitivity) are unchanged and the agent's cost of effort is now

[41] The only significant change concerns the convexification argument. Convexifying the cost of effort across two contracts leads to different paths for the cash flow process and affects consumption under each path. This issue can be addressed by backing out the underlying Brownian path from the agent's report process and expressing the contracts in these terms.

proportional to the negative of the agent's utility: $\phi(A_t) = \phi_1(-W_t)$ where $\phi_1 = \bar{s}/\chi > 0$ and the incentive compatibility condition (22) is replaced by the condition: $u_1(\bar{s}) = \bar{s}(\chi\lambda + \bar{s})/(\chi(\theta - \bar{s})$. The comparative statics also go through (some are harder to prove—see Strulovici (2011)) and new comparative statics with respect to the cost of effort obtain.

Reporting Constraints Truthful renegotiation-proof contracts clearly remain truthful if the agent faces additional constraints on cash flow reports and transfers. For example, the agent could be unable from over-reporting his cash flows beyond some upper bound. Such constraints reduce the set of possible deviations and facilitate truth-telling.

Self-Insurance and Hidden Savings Bloedel, Krishna, and Strulovici (2020) consider an alternative implementation of the model, in which the principal allows the agent to self insure by investing in some asset whose return is predetermined by the principal. They show that equilibria of the self-insurance problem are equivalent to the renegotiation-proof equilibria studied here. They also study the possibility of hidden savings and show that this addition pins down the sensitivity parameter of the contract and yields the contract studied by Williams (2011).

Appendix A: Proofs for Section 2

Proof of Proposition 3

Fix any $x \in \mathbb{R}$ and consider any w_0, w positive such that (x, w_0) and (x, w) can both be reached by some histories. Take some history $(X_0, \ldots, X_t)$ leading to state $(X_t = x, W_t = w_0)$. Also let, for $\tau \geq 1$, $\Delta_\tau = \sum_{s=t}^{t+\tau-1} A_s Z_s$ denote the (random) output increment between time t and $t + \tau$. The continuation payoff of the principal following that history has the form

$$\mathbf{\Pi}(x, w_0) = \frac{x}{1-\rho} + \bar{\alpha}_2 - \bar{\alpha}_1$$

where $\bar{\alpha}_2 = E_t[\sum_{\tau\geq 1} \rho^\tau \Delta_\tau]$ and $\bar{\alpha}_1 = E_t[\sum_{\tau\geq 0} \rho^\tau C_\tau]$ and $E_t[\cdot]$ is the expectation conditional on history $(X_0, \ldots, X_t)$.

Let $(\hat{C}, \hat{A})$ denote the continuation of (C, A) following history $(X_0, \ldots, X_t)$. Letting $\beta = w/w_0$, the transformation $(\phi_\beta(\hat{C}), \phi_\beta(\hat{A}))$ constructed in Section 2 is a contractual equilibrium starting from (x, w). Compared to $(\hat{C}, \hat{A})$ it has the effect of scaling the effort process $\{\hat{A}_\tau = A_{t+\tau}\}_{\tau\geq 0}$ uniformly by the factor β^{1/γ_2} and the consumption process $\{\hat{C}_\tau = C_{t+\tau}\}_{\tau\geq 0}$ uniformly by the factor β^{1/γ_1}, and thus give the principal an expected payoff of

$$\frac{x}{1-\rho} + \beta^{1/\gamma_2}\bar{\alpha}_2 - \beta^{1/\gamma_1}\bar{\alpha}_1.$$

Consistency across promised utilities then implies that $\mathbf{\Pi}(x, \beta w_0) \geq \frac{x}{1-\rho} + \beta^{1/\gamma_2}\bar{\alpha}_2 - \beta^{1/\gamma_1}\bar{\alpha}_1$. The reverse transformation implies the reverse inequality. Hence,

$$\mathbf{\Pi}(x, w) = \frac{x}{1-\rho} + \alpha_2 w^{1/\gamma_2} - \alpha_1 w^{1/\gamma_1} \tag{36}$$

where

$$\alpha_1 = \bar{\alpha}_1 / w_0^{1/\gamma_1} \qquad \alpha_2 = \bar{\alpha}_2 / w_0^{1/\gamma_2}.$$

Equation (36) shows that the conclusion of Proposition 3 holds for all w such that (x, w) can be reached by some history. For w such that (x, w) is not reached by any history, we can complete the function $\mathbf{\Pi}(x, w)$ using the same formula. Moreover, notice that by consistency across outputs, Proposition 2 implies that the coefficients α_1 and α_2 are independent of x, which shows that the conclusion of Proposition 3 holds for all histories. ■

Proof of Proposition 4

In any equilibrium, the agent's promised utility $w = W_0$ at time 0 satisfies:

$$w = u(C_0) - h(A_0) + \rho E[W_1]. \tag{37}$$

We construct contractual equilibria of the following form:

$$A_t = a^{1/\gamma_2} W_t^{1/\gamma_2} \tag{38}$$

$$C_t = c^{1/\gamma_1} W_t^{1/\gamma_1} \tag{39}$$

where a and c are positive constants. In such equilibria, (38) and (39) imply that $u(C_0) = cw$ and $h(A_0) = aw$. Combining this with (37) yields

$$E[W_1] = \frac{1 + a - c}{\rho} w.$$

To analyze incentive compatibility of the agent's strategy, we must specify how the agent's continuation utility evolves as a function of the observed increment $D_0 = A_0 Z_0$. Since A_0 is proportional to w^{1/γ_2} and $E[W_1]$ is proportional to w, we conjecture that

$$W_1 = w^{1-1/\gamma_2} d D_0 \tag{40}$$

for some positive constant d. Under this conjecture, we have

$$E[W_1] = w d z a^{1/\gamma_2}$$

where $z = E[Z_0]$ is a primitive of the model. This yields the relation

$$d z a^{1/\gamma_2} = \frac{1 + a - c}{\rho}. \tag{41}$$

Now consider the agent's incentive to put effort at time 0: for any effort level A, the cost of effort is A^{γ_2} and the expected benefit in terms of promised utility is, from (40) discounted to time 0, $\rho w^{1-1/\gamma_2} Azd$. Therefore, the optimal effort level A_0 solves

$$\max_A \rho w^{1-1/\gamma_2} Azd - A^{\gamma_2}.$$

Since $\gamma_2 > 1$, the objective is strictly concave in A and the maximizer A_0 is determined by the first-order condition

$$\rho z d w^{1-1/\gamma_2} = \gamma_2 A_0^{\gamma_2 - 1},$$

which yields

$$A_0 = \left(\frac{zd\rho}{\gamma_2}\right)^{\frac{1}{\gamma_2 - 1}} w^{1/\gamma_2}. \tag{42}$$

Comparing (38) and (42) yields

$$\gamma_2^{\gamma_2} a^{\gamma_2 - 1} = (zd\rho)^{\gamma_2}. \tag{43}$$

Rearranging (41) we have

$$zd\rho = (1 + a - c)a^{-1/\gamma_2}$$

or

$$(zd\rho)^{\gamma_2} = (1 + a - c)^{\gamma_2} a^{-1}. \tag{44}$$

Comparing (43) and (44) then yields

$$(1 + a - c)^{\gamma_2} a^{-1} = a^{\gamma_2 - 1} \gamma_2^{\gamma_2},$$

which simplifies to

$$a(\gamma_2 - 1) = 1 - c. \tag{45}$$

This shows that there exists a one-dimensional family of state-consistent contractual equilibria, characterized by (38), (39), the dynamic equation $W_{t+1} = W_t^{1-1/\gamma_2} D_t d$, and the relation (45) between the parameters.

Equation (45) implies that $c \leq 1$. This is expected, because cw is the consumption utility given to the agent in period 0. If cw exceeded the promised utility w, the agent could just shirk and received in one period more than what is promised over his entire lifetime.

Also notice that for any give consumption parameter c, the agent's effort parameter a is decreasing in γ_2, which is consistent with the fact that γ_2 determines the agent's marginal cost of effort.

Finally, fixing γ_2, notice that a and c are inversely related. This may be interpreted as follows: $(1 - c)w$ is the amount of promised utility that is given to the agent in the form of expected continuation utility as opposed to immediate consumption utility cw. The higher this amount, the easier it is to incentivize the agent to work.

Appendix B: Proofs for Sections 4–7

Proof of Proposition 5

State consistency implies that, after any history leading up to some state $(\tilde{v}, \tilde{x}) = g \circ (v, x)$ and continuation equilibrium $(\tilde{C}, \tilde{L})$, we have

$$\Pi(\tilde{C}, \tilde{L}) \geq \Pi(\Phi_g(C, L)).$$

Suppose by contradiction that this inequality is strict. Monotonicity of Φ applied to g^{-1} implies that

$$\Pi(\Phi_{g^{-1}}(\tilde{C}, \tilde{L})) > \Pi(\Phi_{g^{-1}}(\Phi_g(C, L))) = \Pi(\Phi_{g^{-1}g}(C, L)) = \Pi(C, L),$$

which contradicts the state consistency of C. ■

Proof of Proposition 8

We prove that if a strategy L gives a higher expected utility to the agent than another strategy L' under the initial contract, then it also does so under the new contract $\hat{C}$. Suppose that L gives higher expected utility than L' under C. This means that

$$E\left[\int_0^\infty e^{-rt}\left(u\left(\mathcal{C}(Y_s : s \leq t) - \int_0^t L_s ds\right)\right) dt\right] \geq E\left[\int_0^\infty e^{-rt}\left(u\left(\mathcal{C}(Y'_s : s \leq t) - \int_0^t L'_s ds\right)\right) dt\right] \quad (46)$$

where Y and Y' satisfy the dynamic equations

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

and

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

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

Now consider another initial condition $\hat{x}$. The reporting processes under strategies L and L' are respectively

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

and

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

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

By construction, the virtual cash flow 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 $\tilde{Y}_0 = \tilde{Y}'_0 = x$. Combining the last four equations pairwise yields

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

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

subject to the conditions $\tilde{Y}_0 = \tilde{Y}'_0 = x$. By definition of the contract $\hat{C}$, the agent's expected utilities using strategies L and L' when facing contract $\hat{C}$ are identical to the expected utilities appearing in (46) and we have just shown that $\tilde{Y}$ and $\tilde{Y}'$ are subject to the same dynamic equations and initial condition as Y and Y'. This prove that strategy L also dominates L' under contract $\hat{C}$. ■

Proof of Proposition 12

Consider any reporting strategy L and let $V^\lambda(L)$ denote the agent's expected utility when using strategy L and facing contract C^λ. Since the agent consumes $C_t^\lambda - G_t$ at time t and the agent has utility $u(c) = -\exp(-\theta c)$, we have

$$V^\lambda(L) = E\left[\int_0^\infty e^{-rt} U_t^\lambda e^{\theta G_t^L} dt | L \right]$$

where $U_t^\lambda = u(C_t^\lambda)$ and $G_t^L = \int_0^t L_s ds$ is the gap. By construction of C^λ, we have

$$U_t^\lambda = \lambda U_t^1 + (1 - \lambda) U_t^2$$

for all t and report histories $\{Y_s : s \leq t\}$, where $U_t^i = u(C_t^i)$ for $i \in \{1, 2\}$. This implies that

$$V^\lambda(L) = \lambda V^1(L) + (1 - \lambda) V^2(L) \tag{47}$$

where $V^i(L)$ is the agent's expected utility with strategy L under contract C^i. By assumption, the contracts C^1 and C^2 are truthful, so for $i \in \{1, 2\}$, we have $V^i(L) \leq v$ for all L and $V^i(L \equiv 0) = v$. Combining these observations with (47) yields $V^\lambda(L) \leq v$, and the inequality is tight if $L \equiv 0$. ■

Proof of Lemma 1

Since C is truthful and state consistent, $\hat{C}$ is also truthful and we have $\Pi(C) = \Pi(\hat{C}) = \Pi(w, y)$ by Proposition 11. Fix any $\lambda \in (0, 1)$. From Proposition 12, the convexification C^λ of C and $\hat{C}$ is also truthful and gives the same promised utility w to the agent given y. Moreover, strict concavity of the utility function u together with the definition of C^λ (Equation (17)) implies that

$C_t^\lambda \leq \lambda C_t + (1-\lambda)\hat{C}_t)$ for all t with a strict inequality whenever $C_t \neq \hat{C}_t$. This implies that $\Pi(C^\lambda) < \lambda\Pi(C^1) + (1-\lambda)\Pi(C^2) = \Pi(w, y)$ and violates renegotiation-proofness of C (Equation (18)) unless $\|C - \hat{C}\|_{\mathbb{L}} = 0$. ■

Proof of Proposition 14

The principal's objective is to maximize $\Pi(w, x)$ with respect to u_1. From (14), and neglecting for now terms that are not affected by the contract choice, this is equivalent to solving

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

where

$$\alpha_t = \frac{E[\log(-W_t)]}{\chi} \qquad \text{and} \qquad E[C_t] = c_1 - \frac{E[\log(-W_t)]}{\theta}.$$

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

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

Therefore,

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

The principal's objective is to maximize

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

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

Proof of Proposition 17

The optimal value of $\bar{s}$ is unchanged if we multiply (23) by r and get rid of the term r in the factor $(r + u_1 - 1/2\sigma^2\bar{s}^2)$. The resulting objective function is equal to

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

The function $u_1(-\bar{s})$ is increasing in $\bar{s}$. Therefore, the objective is supermodular in $(\bar{s}, r)$ and the result follows (see, e.g., Topkis (1978)). ■

Appendix C: Proof of Theorem 2

To prove strong uniqueness, we can analyze the problem under any probability measure that is absolutely continuous with respect to P. Consider the probability measure Q^Y under which $\frac{1}{\sigma}Y$

is the standard Brownian motion. This measure is absolutely continuous because Y has constant quadratic variation equal to σ, and the drift of Y satisfies the standard integrability conditions.[42] Equation (26) is an SDE with constant volatility coefficient and functional drift $-\mathcal{L}(t,\cdot)$.

Part 1. The boundedness assumed in the premise of Theorem 2 implies that condition (4.169) in Liptser and Shiryaev (2001) is satisfied. Theorem 4.13 of Liptser and Shiryaev (2001) then implies that (26) has a unique (in law) weak solution. This, together with Theorem 3.2 in Cherny (2002) implies pathwise uniqueness, which implies that X is adapted to Y and the existence of a unique strong solution, by the Yamada-Watanabe principle.[43]

Part 2. Equation (26) satisfies assumptions (a1)–(a4) of Huang (2018): Assumption (a1) and (a2) trivially hold because volatility is constant, assumption (a4) holds because the term b in Huang's paper is identically equal to 0, and assumption (a3) is explicitly imposed in the premise of Theorem 2 and the premise of Part 2. Theorem 1.1 of Huang (2018) then implies pathwise uniqueness—which implies that X is adapted to Y—and the existence of a unique strong solution.[44]

Part 3. Equation (26) satisfies conditions (C1)–(C3) of Bachmann (2020): Condition (C1) comes from the boundedness assumed in the premise of Theorem 2. Condition (C2) holds trivially because the volatility coefficient is constant and positive, and Condition (C3) is explicitly imposed in the premise of Part 2. Bachmann's Theorem 1.5 then implies that pathwise uniqueness holds. This implies that X is adapted to Y and the existence of a strong solution. ■

Appendix D: Proofs for Section 9 – Necessary Conditions for Incentive Compatibility

LEMMA 2 *For any truthful, renegotiation-proof contract, we have*

$$\bar{v}_y(w,x,x) = -\bar{v}_x(w,x,x) = -\int_0^\infty e^{-(r+\lambda)t} u'(C_t + X_t - Y_t)dt. \tag{48}$$

Proof. The value function of the agent satisfies the optimization problem

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

[42]The drift of Y depends on two components: the drift of X, which is linear in X, and the lying process L which is assumed to be locally bounded.

[43]This principle was stated and proved by Yamada and Watanabe (1971) for standard (as opposed to functional) SDEs. However, their proof extends without any change to functional SDEs, a fact that is well-known and widely used in the literature on functional SDEs. See, e.g., Cherny (2002, Figure 1) and Huang (2018, Theorem 2.8). I am grateful to Nikolai Krylov and Stefan Bachmann for pointing this out.

[44]See Theorem 2.8 in Huang (2018).

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

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

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

$$dY_t = [(\xi - \lambda Y_t) + \bar{L}_t]dt + \sigma dB_t. \tag{49}$$

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

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

subject to $Y_0 = y$, $X_0 = x$, (49), and

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

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

If the contract is incentive compatible, it is optimal to set $\bar{L}_t = 0$ whenever initial conditions are such that $y = x$. Section 9 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 $\bar{v}_x(w, x, x)$ 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, and $\mathcal{C}$ are independent of the initial condition x, and

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

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 defined by the dynamic equation (6) is a generalization of Ornstein-Uhlenbeck processes, which can be explicitly integrated:

$$X_t = e^{-\lambda t}x + \int_0^t e^{\lambda(s-t)}\xi ds + \int_0^t e^{\lambda(s-t)}\sigma dB_s \tag{50}$$

This implies that $\frac{dX_t}{dx} = e^{-\lambda t}$ and yields the formula of Lemma 2. ■

Lemma 3 *The following equality holds for any truthful renegotiation-proof contract:*

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

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

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

Moreover,

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

With the exponential utility specification, we have $u(C_t) = (-W_t)u_1$. Letting $\vartheta(t) = E[W_t]$ ($\vartheta(0) = w$), this implies that

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

and hence that

$$E[W_t] = e^{(r+u_1)t}w. \tag{51}$$

Integrating this expression over time yields the result. ■

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.

Bachmann, S. (2020) "On the Strong Feller Property and Well-Posedness for SDEs with Functional, Locally Unbounded Drift," Working Paper, https://arxiv.org/abs/1808.05629.

Bassetto, M., Huo, Z., and J.-V. Ríos-Rull (2020) "Institution Building without Commitment," Working Paper, University of Pennsylvania.

Battaglini, M. (2007) "Optimality and Renegotiation in Dynamic Contracting," *Games and Economic Behavior*, Vol. 60, pp. 213–246.

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.

Bernard, B. (2016) *Continuous-Time Repeated Games with Imperfect Information: Folk Theorems and Explicit Results,* Doctor of Philosophy Thesis, Department of Mathematics and Statistical Sciences, University of Alberta.

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

Bloedel, A., Krishna, R. V., and O. Leukhina (2020) "Insurance and Inequality with Persistent Private Information," *Working paper*, Stanford University and Federal Reserve Bank of St. Louis.

Bloedel, A., Krishna, R. V., and B. Strulovici (2020) "Persistent Private Information: Comment," *Working Paper*.

Cherny, A. S. (2002) "On the Uniqueness in Law and the Pathwise Uniqueness for Stochastic Differential Equations," *Theory of Probability and its Applications*, Vol. 46, pp. 406–419.

DeMarzo, P., Sannikov, Y. (2006) "Optimal Security Design and Dynamic Capital Structure in a Continuous-Time Agency Model," *Journal of Finance*, Vol. 61, pp. 2681–2724.

DeMarzo, P., Sannikov, Y. (2016) "Learning, Termination, and Payout Policy in Dynamic Incentive Contracts," *Review of Economic Studies*, Vol. 84, pp. 182–236.

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. (2013) "Insurance and Taxation over the Life Cycle," *Review of Economic Studies*, Vol. 80, pp. 596–635.

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.

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. (2011) “Delegated Experimentation,” Working Paper, Northwestern University.

Garrett, D., Pavan, A. (2015) “Dynamic Managerial Compensation: A Variational Approach,” *Journal of Economic Theory*, Vol. 159, pp. 775–818.

Garrett, D., Pavan, A. (2012) “Managerial Compensation in a Changing World,” *Journal of Political Economy*, Vol. 120, pp. 879–925.

Golosov, M., Kocherlakota, N., and A. Tsyvinski (2003) “Optimal Indirect and Capital Taxation,” *Review of Economic Studies* Vol. 70, pp. 569–587.

Golosov, M., Iovino, L. (2020)“Social Insurance, Information Revelation, and Lack of Commitment” Working Paper, Bocconi University and University of Chicago.

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.

Guo, Y. (2016) “Dynamic Delegation of Experimentation,” *American Economic Review*, Vol. 106, pp. 1969–2008.

Hart, O., Tirole, J. (1988) “Contract Renegotiation and Coasian Dynamics,” *Review of Economic Studies*, Vol. 55, pp. 509–540.

Hörner, J., Samuelson, L. (2013) “Incentives for Experimenting Agents,” *RAND Journal of Economics*, Vol. 44, pp. 632–663.

Huang, X. (2018) “Strong Solutions for Functional SDEs with Singular Drift,” *Stochastics and Dynamics*, Vol. 18, No. 2.

Kapička, M. (2013) “Efficient Allocations in Dynamic Private Information Economies with Persistent Shocks: A First-Order Approach,” *Review of Economic Studies*, Vol. 80, pp. 1027–1054.

Karatzas, I., Shreve, S. (1991) *Brownian Motion and Stochastic Calculus*, Springer, New York, N.Y.

Kranz, S., Ohlendorf, S. (2013) *Renegotiation-Proof Relational Contracts*, *Games and Economic Behavior*, Vol. 80, pp. 157–178.

Laffont, J.-J., Tirole, J. (1990) "Adverse Selection and Renegotiation in Procurement," *Review of Economic Studies*, Vol. 57, No. 4, pp. 597–625.

Liptser, R., Shiryaev, A. (2001) *Statistics of random processes: I. General theory (Vol. 1)*, Springer.

Maestri, L. (2017) "Dynamic Contracting under Adverse Selection and Renegotiation," *Journal of Economic Theory*, Vol. 171, pp. 136–173.

Milgrom, P., Segal, I. (2002) "Envelope Theorems for Arbitrary Choice Sets," *Econometrica*, Vol. 70, No. 2, pp. 583–601.

Miller, D., Watson, J. (2013) "A Theory of Disagreement in Repeated Games with Bargaining," *Econometrica*, Vol. 81, pp. 2303–2350.

Øksendal, B., Sulem, A. (2004) *Applied Stochastic Control of Jump Diffusions*, Springer Verlag, Berlin, Germany.

Pavan, A., Segal, I., and J. Toikka (2014) "Dynamic Mechanism Design: A Myersonian Approach," *Econometrica*, Vol. 82, pp. 601–653.

Pearce, D. (1987) "Renegotiation-Proof Equilibria: Collective Rationality and Intertemporal Cooperation," *Cowles Foundation Discussion Paper*, No. 855, Yale University.

Ray, D. (1994) "Internally Renegotiation-proof Equilibrium Sets: Limit Behavior with Low Discounting," *Games and Economic Behavior*, Vol. 6, pp. 162–177.

Safronov, M., Strulovici, B. (2018) "Contestable Norms," Working Paper, Cambridge University and Northwestern University.

Sannikov, Y. (2008) "A Continuous-Time Version of the Principal–Agent Problem," *Review of Economic Studies*, Vol. 75, No. 3, pp. 957—984.

Sannikov, Y. (2014) "Moral Hazard and Long-Run Incentives," Working Paper, Princeton University.

Strulovici, B. (2011) "Renegotiation-Proof Contracts with Moral Hazard and Persistent Private Information," Discussion Paper, Center for Mathematical Studies in Economics and Management Science, Northwestern University.

Strulovici, B. (2017) "Contract Negotiation and the Coase Conjecture: A Strategic Foundation for Renegotiation-Proof Contracts," *Econometrica*, Vol. 85, pp. 585–616.

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.

Van Damme, E. (1991) *Stability and perfection of Nash equilibria*, Springer-Verlag, Berlin.

Williams, N. (2011) "Persistent Private Information," *Econometrica*, Vol. 79, pp. 1233–1275.

Yamada. T., Watanabe, S. (1971) "On the Uniqueness of Solutions of Stochastic Differential Equations," *Journal of Mathatics of Kyoto University*, Vol. 11, pp. 155–167.

Zhang, Y. (2009) "Dynamic Contracting with Persistent Shocks," *Journal of Economic Theory*, Vol. 144, pp. 635–675.